A Bayesian approach to analyzing replicated preference tests

Suzanne R. Dubnicka

Published Version Information

Citation: Dubnicka, S. R. (2013). A Bayesian approach to analyzing replicated preference tests. Journal of Sensory Studies, 28(3), 171-187.

Digital Object Identifier (DOI): doi:10.1111/joss.12033

Publisher's Link: http://onlinelibrary.wiley.com/doi/10.1111/joss.12033/full

A Bayesian Approach to Analyzing Replicated Preference Tests

Runnning Title: Bayesian Analysis of Replicated Preference Tests

Suzanne R. Dubnicka
Kansas State University
Department of Statistics
Dickens Hall
Manhattan, KS 66506
dubnicka@k-state.edu
voice: 785-532-0513
fax: 785-532-7736

July 17, 2012

Abstract

Replicated or multiple preference tests have become important tools for assessing consistency in consumer preferences in repeated tests as well as overall consumer preference. Replicated preference tests can also provide the means for separating discriminators from non-discriminators. Despite the increasing popularity of multiple preference tests, there are few statistical tools for their analysis, especially when one is interested in assessing the consistency of consumer preferences. This paper presents flexible Bayesian methods for examining overall consumer preference and consistency of consumer preference in replicated preference tests. In particular, this paper presents Bayesian methods for forced-choice preference testing with two tests and then extends this methodology to include forced-choice preference testing with more than two tests and replicated preference testing with a no-preference or no-choice option. The methods produce intuitive and easily interpreted probabilities. These methods are applied to various replicated preferences test data from the literature.

Practical Applications

The Bayesian methods presented in this paper will help sensory scientists and statisticians working with sensory data to explore data from replicated preference tests more thoroughly. Currently-used methodology only allows scientists to assess the overall preference for a particular product. The methodology in this paper allows for a wider array of questions to be answered. In particular, this paper focuses on consumers' ability to consistently choose the same product. Initiallly, the methods apply to forced-choice tests, but they are later extended to include a no-preference option. Allowing for a no-preference options is another important contribution of this work.

Keywords: forced-choice preference test; McNemar's test; multiple preference tests; no-preference option; statistical analysis

1 Introduction

Preference testing is commonly used to determine consumer preference in comparisons of two, or sometimes more, products. Typically, consumers complete one preference test. However, replicated or multiple preference tests serve several purposes. Greenberg and Collins (1966) discuss the use of replicated preference tests to distinguish between discriminators and non-disciminators. Wilke, Cochrane, and Chambers IV (2006) make several observations regarding consumer preferences. For example, consumers may not be consistent in their product preferences across tests, and the percentage of consumers preferring a particular product do not necessarily stay the same across tests.

Despite the observations of Wilke, Cochrane, and Chambers IV (2006), there are relatively few statistical tools for anaylzing replicated preference tests, especially when one is concerned with the consistency of consumer preference across tests. Cochrane, Dubnicka, and Loughin (2005) compare various non-Bayesian methods for analyzing replicated preference tests, in particular, for determining overall product preference. Ennis and Bi (1998) employ a beta-binomial model to account for variability among tests and provide maximum likelihood estimates for the parameters of this model. However, they still only provide one overall estimate for product preference; they do not attempt to assess changes in product preference across tests. Bi (2003) specifies a Bayesian model but assumes that the probability of preferring a particular product is constant across tests and among consumers. In this paper, we propose a Bayesian model that allows the probability of preferring a particular product to vary across tests.

McNemar's test is a commonly used frequentist, i.e., non-Bayesian, test for determining the difference in proportions for binary matched pairs data. In the context of replicated forced-choice preference testing with tests on two occasions, McNemar's test would be used to determine if the proportions of consumers preferring product A differed on those two occasions. Altham (1971) developed a Bayesian approach for analyzing binary matched pairs data, essentially providing a Bayesian version of McNemar's test. More generally, Agresti and Hitchcock (2005) discuss Bayesian approaches for categorical data with the work of Altham (1971) among them.

In this paper, we first apply the methodology of Altham (1971) to the case of replicated preference testing with forced-choice preference tests with two tests, i.e., with tests on two occasions. In the context of replicated preference testing, this Bayesian version of McNemar's test provides us with the posterior probability that product A is preferred more on the second occasion (test) than on the first. This posterior probability can tell us, in some sense, if preference for product A increases over time. Furthermore, this posterior probability can also help us determine if switching from another product, say product B, on the first occasion to product A on the second occasion is more likely than the reverse. Furthermore, we extend this methodology to answer other questions relevant to replicated preference testing with two tests. In particular, under the same assumptions and notation provided by Altham (1971), we determine whether consumers are more likely to chose the same product on both occasions than to switch product preference and whether product A is preferred more than product B on both occasions, among other questions. In general, we focus on questions regarding overall preference for product A as well as consistency in product preference over time.

After considering questions related to the replicated preference tests on two occasions, we extend the Bayesian methodology in two ways. First, we consider the extension to more than two occasions. In the statistics literature, this is referred to as binary repeated measures data; however, in preference testing, we attempt to answer questions that are somewhat different that those typically considered in the statistics literature. Nevertheless, the Bayesian methodology for two occasions extends very naturally to accommodate more than two occasions. We again focus on questions regarding the changes in product preference over time and the consistency of individual consumer preferences. Second, we consider the extension to non-forced-choice replicated preference tests, that is, tests in which a "no choice" or "no preference" option is allowed. The proposed Bayesian methodology also extends quite easily to handle this option.

In Section 2, we present the Bayesian approach of Altham (1971) in the context of replicated forced-choice preference tests with tests on two occasions, and we extend that approach to answer various other questions of interest in preference testing. We also discuss Bayesian methodology in general with a particular emphasis on choosing the prior distribution in this context, and we apply the methodology to real data. In Section 3, we extend this Bayesian methodology to replicated forced-choice preference tests with tests on more than two occasions and again apply it to real data. Section 4 provides details for our Bayesian approach to replicated preference testing with a "no preference" option. An example is also provided. All methods are implemented in R, and the R code is included in the Appendix.

2 Matched Pairs

To set notation, let the binary variable y_{ijk} equal 1 if consumer k prefers product i on occasion 1 and prefers product j on occasion 2, where $i, j = 1, 2$ with 1 denoting product A and 2 denoting product B, and $k = 1, \ldots, n$ with n denoting the total number of consumers. For example, $y_{127} = 1$ if consumer 7 prefers product A on the first occasion and product B on the second. Furthermore, let $n_{ij} = \sum_{k=1}^{n} y_{ijk}$ be the number of consumers who prefer product i on occasion 1 and product j on occasion 2. For example, n_{12} is the number of conumsers who prefer A on the first occasion and B on the second. Also, let θ_{ij} denote the probability that i is preferred on occasion 1 and j is preferred on occasion 2, that is,

$$P(y_{ijk} = 1) = \theta_{ij}. \tag{1}$$

We are making a necessary assumption that the probability in (1) is the same for all consumers. Then $\theta_{i\cdot} = \theta_{i1} + \theta_{i2}$ is the (marginal) probability that product i is preferred on occasion 1, and $\theta_{\cdot j} = \theta_{1j} + \theta_{2j}$ is the (marginal) probability that product j is preferred on occasion 2. Finally, $\theta_{11} + \theta_{12} + \theta_{21} + \theta_{22} = 1$.

Under the Bayesian paradigm, parameters are random quantities. Therefore, Bayesian methods require that we specify a likelihood, that is, a probability distribution for the data, and prior distributions, which are probability distributions for the parameters. Based on the assumption that consumers' preferences are independent of one another and that (1) holds, the likelihood is given, up to a proportionality constant, by

$$p(\boldsymbol{n}|\boldsymbol{\theta}) \propto \prod_{k=1}^{n} \prod_{i=1}^{2} \prod_{j=1}^{2} \theta_{ij}^{y_{ijk}} = \prod_{i=1}^{2} \prod_{j=1}^{2} \theta_{ij}^{n_{ij}}, \tag{2}$$

where $\boldsymbol{n} = (n_{11}, n_{12}, n_{21}, n_{22})$ is the data vector, and $\boldsymbol{\theta} = (\theta_{11}, \theta_{12}, \theta_{21}, \theta_{22})$ is the vector of parameters of interest. This likelihood is essentially a multinomial distribution.

In Bayesian analyses with the binomial distribution as the likelihood, it is common to use the beta distribution as the prior distribution. This combination of binomial likelihood and beta prior results in a posterior distribution that is also a beta distribution. The multinomial distribution is simply a multivariate extension of the binomial distribution, and the common prior distribution in this case is the Dirichlet distribution, which is a multivariate extension of the beta distribution. Therefore, we choose our prior distribution to be the Dirichlet distribution given by

$$p(\boldsymbol{\theta}) = \Gamma(\mu_0) \prod_{i=1}^{2} \prod_{j=1}^{2} \frac{1}{\Gamma(\mu_{ij})} \theta_{ij}^{\mu_{ij}-1}, \tag{3}$$

where $0 < \theta_{ij} < 1$ for $i, j = 1, 2$ such that $\sum_{i=1}^{2} \sum_{j=1}^{2} \theta_{ij} = 1$ and $\mu_0 = \sum_{i=1}^{2} \sum_{j=1}^{2} \mu_{ij}$ with $\mu_{ij} > 0$, $i, j = 1, 2$. $\Gamma(\cdot)$ is the gamma function, and $\Gamma(x) = (x-1)!$ if x is a positive integer. The values of parameters $\mu_{ij} > 0$ of the Dirichlet distribution need to be set prior to analysis. The choice of these parameter values is discussed below.

Inference in Bayesian analyses is based on the posterior distribution, which is the distribution of the parameters conditional on the data. Essentially, the posterior distribution can be thought of as an update of the prior distribution based on the observed data. Together the likelihood in (2) and the prior in (3) lead to a Dirichlet posterior distribution given, up to a proportionality constant, by

$$p(\boldsymbol{\theta}|\boldsymbol{n}) \propto \prod_{i=1}^{2} \prod_{j=1}^{2} \theta_{ij}^{n_{ij}+\mu_{ij}-1}. \tag{4}$$

That is, the posterior distribution of $\boldsymbol{\theta}$, given the data $\boldsymbol{n}$, is a Dirichlet distribution with parameters $\boldsymbol{\nu} = (\nu_{11}, \nu_{12}, \nu_{21}, \nu_{22})$ where $\nu_{ij} = n_{ij} + \mu_{ij}$.

2.1 Choosing Prior Parameters

The choice of prior parameters should reflect one's belief about the values of the θ_{ij} prior to conducting the current preference tests as well as one's certainty about those values. In choosing the values of the parameters of the Dirichlet prior distribution, μ_{ij}, it helpful to understand some characteristics of the Dirichlet distribution. These characteristics are similar to those of the beta distribution.

- The (prior) mean of θ_{ij} is given by μ_{ij}/μ_0, where $\mu_0 = \sum_{i=1}^{2}\sum_{j=1}^{2}\mu_{ij}$.

- If the μ_{ij} are equal to one another, then the Dirichlet distribution is symmetric.

- If $\mu_{ij} = 1$, $i, j = 1, 2$, then the Dirichlet distribution is flat, essentially a multivariate uniform distribution.

- If $\mu_{ij} > 1$, $i, j = 1, 2$, then the Dirichlet distribution is unimodal. Larger values of the μ_{ij} produce a less variable Dirichlet distribution, suggesting that values far from the mode are less likely.

In particular, the prior means can be chosen to reflect the prior belief about the approximate values of the θ_{ij}, and the exact values of the μ_{ij} can then be selected, keeping in mind that larger values of the μ_{ij} reflect great certainty in the choice of the means. We suggest the following guidelines in choosing prior parameters in the replicated preference testing setting.

- The weight of the prior information is given by $\mu_0 = \mu_{11} + \mu_{12} + \mu_{21} + \mu_{22}$. The value of μ_0 relative to the total sample size $n = n_{11} + n_{12} + n_{21} + n_{22}$ should be considered carefully in selecting the values of the μ_{ij}. Specifically, values of μ_0 close to n suggest that one is willing to give just as much weight to prior beliefs as to the study data. However, if one wants the analysis to be driven more by the data, then μ_0, and hence the μ_{ij}, should be much smaller than n.

- The prior weight given to switching product preference from occasion 1 to occasion 2 should also be considered. In particular, $\mu_{11} + \mu_{22} > \mu_{12} + \mu_{21}$ suggests a prior belief that consumers are more likely to choose the same product twice than to switch products, while $\mu_{11} + \mu_{22} < \mu_{12} + \mu_{21}$ suggests a belief that consumers are more likely to switch.

- Furthermore, the values of μ_{12} and μ_{21}, relative to one another, imply two related ideas. First, $\mu_{21} > \mu_{21}$ suggests a prior belief that switching from product B to product A will happen more often than switching from product A to product B. However, $\mu_{21} > \mu_{12}$ also implies a prior belief that product A will be chosen more on occasion 2 than it was on occasion 1.

- Finally, $\mu_{11} + \mu_{12} > \mu_{21} + \mu_{22}$ implies prior belief that product A is preferred more than product B at time 1, while $\mu_{11} + \mu_{21} > \mu_{12} + \mu_{22}$ implies a prior belief that A is preferred more than B on occasion 2.

To illustrate more concretely, suppose we believe that, based on previous preference tests with these or similar products, the percentage of consumers who will prefer product A on both occasions is about 45%, and the percentage who prefer product B on both occasions is about 25%. Further, suppose prior studies suggest that the percentage who will switch from A to B is about the same as the percentage who will switch from B to A. Thus, this leads us to believe, a priori, that the means of $(\theta_{11}, \theta_{12}, \theta_{21}, \theta_{22})$ are $(0.45, 0.15, 0.15, 0.25)$, respectively. There are an infinite number of choices for the μ_{ij} depending on $\mu_0 = \mu_{11} + \mu_{12} + \mu_{21} + \mu_{22}$. The value of μ_0 should be chosen to reflect our certainty regarding the values of the means that we have chosen. This certainty should be measured relative to the number of consumers in the replicated preference test that we plan to conduct. For example, if we believe there is more information in our planned replicated preference test with 50 consumers than in our prior beliefs, we should choose μ_0 to be smaller than 50. If we think there is as much or more information in our prior beliefs than in our study data, we should choose μ_0 greater than or equal to 50; however, this choice is not often justifiable. Suppose we want to put relatively little weight on our prior beliefs. Then we might choose $\mu_0 = 10$ which leads to $\mu_{11} = 4.5$, $\mu_{12} = 1.5$, $\mu_{21} = 1.5$, and $\mu_{22} = 2.5$. Note that this choice of prior parameters also reflects a belief that product A is preferred more than B on both occasions. If one has no prior information regarding the products, it is reasonable to choose $\mu_{11} = \mu_{12} = \mu_{21} = \mu_{22} = 1$, resulting in a noninformative prior, and the data will drive the analysis rather than prior beliefs. For illustration, several choices of prior parameters are compared in the example in Section 2.3.

2.2 Details of Analysis

In a Bayesian analysis, the posterior distribution, which is the distribution of the parameters conditional on the data, is used to answer all questions of interest. Recall that the posterior distribution can be thought of as an update of our prior knowledge based on the study data. In our case, we will use certain posterior probabilities, that is, probabilities regarding the parameters conditional on the data, to answer the questions outlined in Section 1. For example, we may be interested in knowing if product A is more likely to be preferred on occasion 2 than on occasion 1, implying that consumers may grow to like product A more over time. In terms of the parameters θ_{ij}, we are interested in the posterior probability that $\theta_{\cdot 1} = \theta_{11} + \theta_{21}$ is greater than $\theta_{1\cdot} = \theta_{12} + \theta_{22}$, that is, the probability that product A is preferred more on occasion 2 than occasion 1 given the data:

$$P_1 = P(\theta_{\cdot 1} > \theta_{1\cdot}|\boldsymbol{n}) = P(\theta_{11} + \theta_{21} > \theta_{11} + \theta_{12}|\boldsymbol{n}) = P\left(\frac{\theta_{12}}{\theta_{12} + \theta_{21}} < \frac{1}{2}\middle|\boldsymbol{n}\right). \tag{5}$$

It can be shown that, conditional on the data $\boldsymbol{n}$ from the replicated preference test, the distribution of $\theta_{12}/(\theta_{12} + \theta_{21})$ is a beta distribution. In particular,

$$\left.\frac{\theta_{12}}{\theta_{12} + \theta_{21}}\right|\boldsymbol{n} \sim \text{Beta}(\nu_{12}, \nu_{21}), \tag{6}$$

where $\nu_{ij} = n_{ij} + \mu_{ij}$. Standard statistical software, such as R, can be used to compute the probability in (5).

If ν_{12} and ν_{21} are both positive integers, then the desired probability, $P(\theta_{\cdot 1} > \theta_{1\cdot}|\boldsymbol{n})$, may be evaluated as a binomial tail:

$$P_1 = P(\theta_{\cdot 1} > \theta_{1\cdot}|\boldsymbol{n}) = \sum_{r=0}^{\nu_{21}-1} \binom{\nu_{12} + \nu_{21} - 1}{r} \left(\frac{1}{2}\right)^{\nu_{12}+\nu_{21}-1}. \tag{7}$$

As Altham (1971) notes, this is comparable to the p-value obtained in McNemar's test:

$$\text{p-value} = \sum_{r=0}^{n_{21}} \binom{n_{12} + n_{21}}{r} \left(\frac{1}{2}\right)^{n_{12}+n_{21}}. \tag{8}$$

This is the p-value calculated when testing the null hypothesis of no difference in preference on the two occasions against the alternative that product A is preferred more on occasion 1 than on occasion 2. In Bayesian analyses, one does not specify null and alternative hypotheses. Rather one can compute posterior probabilities as in equation (5) or Bayes factors. We will focus on posterior probabilities in this paper. However, the posterior probability P_1 will equal the p-value from McNemar's test if, and only if, $\mu_{12} = 1$ and $\mu_{21} = 0$. This particular choice of prior parameters corresponds to a prior belief that A is preferred more often on occasion 1 than on occasion 2. Note, however, that the posterior probability P_1 and McNemar's p-value, although equal in value, have quite different meanings.

We may also be interested in consistency of product choice, that is, whether consumers are more likely to choose the same product on both occasions than to switch. In terms of the parameters, we wish to find the posterior probability that $\theta_{11}+\theta_{22}$ is greater than $\theta_{12}+\theta_{21}$, that is,

$$P_2 = P(\theta_{11} + \theta_{22} > \theta_{12} + \theta_{21}|\boldsymbol{n}) = P(\theta_{11} + \theta_{22} > 1/2|\boldsymbol{n}). \tag{9}$$

Given the preference test data $\boldsymbol{n}$, the distribution of $\theta_{11} + \theta_{22}$ follows a beta distribution:

$$\theta_{11} + \theta_{22}|\boldsymbol{n} \sim \text{Beta}(\nu_{11} + \nu_{22}, \nu_0 - \nu_{11} - \nu_{22}), \tag{10}$$

where $\nu_0 = \sum_{i=1}^{2} \sum_{j=1}^{2} \nu_{ij}$ with $\nu_{ij} = n_{ij} + \mu_{ij}$. In fact,

$$\theta_{ij} + \theta_{i'j'}|\boldsymbol{n} \sim \text{Beta}(\nu_{ij} + \nu_{i'j'}, \nu_0 - \nu_{ij} - \nu_{i'j\prime}) \tag{11}$$

so that a variety of questions may be answered by using a beta distribution.

Some questions, however, do not result in probabilities which can be computed directly from known distributions. For example, we can consider ways of assessing changes in preference, such as the posterior probability that product B is preferred more on occasion 1 but product A is preferred more on occasion 2. Thus, we are interested computing

$$P_3 = P(\theta_{11} + \theta_{12} < 1/2 \text{ and } \theta_{11} + \theta_{21} > 1/2|\boldsymbol{n}) = P(\theta_{11} + \theta_{12} < 1/2 < \theta_{11} + \theta_{21}|\boldsymbol{n}). \tag{12}$$

Similarly, we may be interested in determining if, in some sense, A is preferred more that B. This question may be framed as the posterior probability that A is preferred on both occasions:

$$P_4 = P(\theta_{1\cdot} > 1/2 \text{ and } \theta_{\cdot 1} > 1/2|\boldsymbol{n}) = P(\theta_{11} + \theta_{12} > 1/2 \text{ and } \theta_{11} + \theta_{21} > 1/2|\boldsymbol{n}) \tag{13}$$

Neither of these probabilities can be easily computed analytically from the exact posterior distribution. In cases such as these, we will use Monte Carlo simulation to estimate the desired probabilities. In particular, we will randomly generate parameter values from the Dirichlet posterior distribution in equation (4) and use those simulated values to approximate the posterior distribution of the quantity of interest. This, in turn, enables us to approximate the posterior probability of interest.

For example, suppose we want to compute the probability in equation (12). First, we randomly generate a very large number of values, say R values, of $\boldsymbol{\theta} = (\theta_{11}, \theta_{12}, \theta_{21}, \theta_{22})$ from the Dirichlet posterior distribution given by equation (4). The probability of interest (12) is then be estimated by counting the number of generated values of $\boldsymbol{\theta}$ for which the event of interest, namely, $\theta_{11} + \theta_{12} < 1/2 < \theta_{11} + \theta_{21}$, occurs and dividing that by R. More specifically, we use the following algorithm to estimating the probability in (12).

Algorithm 1 *For $r = 1, \ldots, R$, do the following:*

1. *Randomly sample $\boldsymbol{\theta}^{(r)} = (\theta_{11}^{(r)}, \theta_{12}^{(r)}, \theta_{21}^{(r)}, \theta_{22}^{(r)})$ from a Dirichlet distribution with parameters $\boldsymbol{\nu} = (\nu_{11}, \nu_{12}, \nu_{21}, \nu_{22})$, where $\nu_{ij} = n_{ij} + \mu_{ij}$.*

2. *For that sample, let $s_r = 1$ if $\theta_{11}^{(r)} + \theta_{12}^{(r)} < 1/2 < \theta_{11}^{(r)} + \theta_{21}^{(r)}$, and let $s_r = 0$ otherwise.*

Then the probability in (12) is estimated by the proportion of times $\theta_{11} + \theta_{12} < 1/2 < \theta_{11} + \theta_{21}$ out of the R randomly generate samples:

$$\hat{P}_3 = \frac{1}{R} \sum_{r=1}^{R} s_r. \tag{14}$$

Larger values of R produce more precise estimates of the desired probability. The algorithm for computing the probability in equation (13) is similar. The above algorithm was implemented in R. The R code for estimating the probabilities in equations (12) and (13) is given in the Appendix.

2.3 Cola Data

Wilke, Cochrane, and Chambers IV (2006) conduct two different forced-choice replicated preference tests: one comparing two brands of raisin bran, the other comparing two colas. In both cases, preferences tests were conducted on four occasions; that is, four tests were conducted. In this section, we will use the data from the cola preference tests on the third and fourth occasions to illustrate the methods. The raisin bran data will be used to illustrate an extension of this methodology to more than two occasions in Section 3.

The cola preference tests were conducted on 296 consumers, 18 years of age or older, who were self-reported acceptors of cola products. Forced-choice preference tests were conducted on four occasions. More details regarding the procedures can be found in Wilke et al. (2006). Table 1 summarizes the data collected from the third and fourth occasions.

[Table 1 about here.]

For the purpose of this example, Test 3 is the first occasion, and Test 4 is the second occasion. Using the earlier notation, we have $n_{11} = 120$, $n_{12} = 62$, $n_{21} = 56$, and $n_{22} = 58$. Prior parameters μ_{ij} should be chosen prior to viewing the data. However, it is instructive to compare different prior choices and their effects on the analyses and conclusions. Recall that choosing $\mu_{11} = \mu_{12} = \mu_{21} = \mu_{22} = 1$ results in a noninformative prior, suggesting that we have little prior information regarding the parameters θ_{ij}. Also, note that the values of the μ_{ij} are quite small compared to the values of the n_{ij}. Practically speaking, this means that the data will drive the analysis, and our prior information will play a minimal roll. Additionally, choosing $\mu_{11} = 1$, $\mu_{12} = 0$, $\mu_{21} = 1$, and $\mu_{22} = 1$ allows us to compute the Bayesian McNemar's test of Altham (1971). As the values of μ_{ij} for this prior are also very small relative to the observed counts, this is also a noninformative prior. To distinguish the latter prior, we refer to it as the McNemar prior.

If we feel our prior information should play a greater roll in the analysis, the values of the μ_{ij} should be larger. For example, we may be willing to give the prior about one-third of the weight of the data, say, $\mu_0 = 100$. Furthermore, prior experience with similar products may suggest that consumers are more likely to stay with the same product on both occasions than to switch ($\mu_{11} + \mu_{22} > \mu_{12} + \mu_{21}$) and that consumers who switch are more likely to switch from product B to product A than the reverse ($\mu_{21} > \mu_{12}$). One set of prior parameters that reflects these beliefs, subject to $\mu_0 = 100$, is $\mu_{11} = 35$, $\mu_{12} = 5$, $\mu_{21} = 35$, and $\mu_{22} = 35$. Notice that this prior puts the same weight on preferring product A on both occasions as preferring product B on both occasions. There are many other priors that would meet these criteria.

Table 2 describes several different priors that were considered for this study. Other than the noninformative prior, all priors are based on $\mu_0 = 100$, making them informative priors. In many cases, these priors are more informative than one might be willing to impose. These priors, described in the table, have a variety of characteristics. Some priors are similar to the data in that the prior proportions are similar to the data proportions, while others are quite different from the data. Priors were also chosen to favor the different posterior probabilities under consideration: P_1, P_2, P_3, and P_4. The intent is not to provide an exhaustive list of priors but to show how the prior parameter selection impacts the resulting values of the posterior probabilities of interest.

Recall that the posterior probabilities of interest are

$$\begin{aligned}
P_1 &= P(\theta_{\cdot 1} > \theta_{1\cdot}|\boldsymbol{n}) \;=\; P(\theta_{11} + \theta_{21} > \theta_{11} + \theta_{12}|\boldsymbol{n}) \;=\; P\left(\frac{\theta_{12}}{\theta_{12} + \theta_{21}} < \frac{1}{2}\middle|\boldsymbol{n}\right), \\
P_2 &= P(\theta_{11} + \theta_{22} > \theta_{12} + \theta_{21}|\boldsymbol{n}) \;=\; P(\theta_{11} + \theta_{22} > 1/2|\boldsymbol{n}), \\
P_3 &= P(\theta_{11} + \theta_{12} < 1/2 \text{ and } \theta_{11} + \theta_{21} > 1/2|\boldsymbol{n}) \;=\; P(\theta_{11} + \theta_{12} < 1/2 < \theta_{11} + \theta_{21}|\boldsymbol{n}), \\
P_4 &= P(\theta_{1\cdot} > 1/2 \text{ and } \theta_{\cdot 1} > 1/2|\boldsymbol{n}) = P(\theta_{11} + \theta_{12} > 1/2 \text{ and } \theta_{11} + \theta_{21} > 1/2|\boldsymbol{n}).
\end{aligned}$$

Thus, P_1 is the posterior probability that product A is preferred more on the second occasion (time 4) than it was on the first occasion (time 3). P_2 is the posterior probability that a consumer is more likely to prefer the same product on both occasions than to switch. Therefore, $1 - P_2$ is the probability that a consumer is more likely to switch product preference than to stay with the same product. P_3 is the posterior probability that product B is preferred more on the first occasion (time 3) but product A is preferred more on the second occasion (time 4). A high value of P_3 would indicate an overall change in preference among consumers. P_4 is the probability that product A is preferred more on both occasions. A high value of P_4 would be one indication that product A is preferred more overall.

The resulting values of these posterior probabilities are given in Table 3. Recall that, when noninformative priors are used, the data are driving the analyses. Therefore, to see the influence of the choice of prior distribution on the posterior probabilities of interest, one can compare the posterior probabilities under the various informative priors to the corresponding probabilities under the noninformative prior. For example, prior B essentially results as the same posterior probabilities as the noninformative prior because the prior parameter values reflect the proportions in the data. Prior C has its greatest influence on posterior probability P_2, lowering its value from that of the noninformative prior, because the prior parameters put a greater weight on switching products than the data reflects.

To conclude this example, we compare the p-value for McNemar's test with the value of the posterior probability P_1 under the McNemar prior, i.e., the Dirichlet prior with $\mu_{11} = 1$, $\mu_{12} = 0$, $\mu_{21} = 1$, and $\mu_{22} = 1$. Recall that this p-value and posterior probability will be the same numerically but will differ in interpretation. In fact, we find that this common value is 0.323. That is, the p-value of McNemar's test is 0.323 which means that we fail to reject the null hypothesis that preference for product A differs on the two occasions. As this p-value was computed for the one-sided alternative, we cannot conclude that product A is preferred more on occasion 1 than on occasion 2. For the Bayesian analysis, using the McNemar prior, the posterior probability P_1 is given by 0.323. That is, the probability that product A is preferred more on occasion 2 than occasion 1 is 0.323. It is more likely that product A is preferred more on occasion 1 than occasion 2, with probability equal to $1 - 0.323 = 0.677$. Notice that there is a similarity in the conclusions, but the underlying meaning differs.

[Table 2 about here.]

[Table 3 about here.]

3 More than Two Occasions

Suppose that forced-choice preference tests are conducted more than twice, i.e., on more than two occasions. The result is binary repeated measures data, and we can easily extend the models in Section 2 to accommodate this situation. For concreteness, suppose that four preference tests are conducted. Let the binary variable y_{ijklm} equal 1 if consumer m, where $m = 1, \dots, n$, prefers products (i, j, k, l) on occasions $(1, 2, 3, 4)$, respectively. Again $i, j, k, l = 1, 2$ where 1 is product A and 2 is product B. For example, $y_{11217} = 1$ if consumer 7 prefers product A on occasions 1, 2, and 4 and product B on occasion 3. As above, let $n_{ijkl} = \sum_{m=1}^{n} y_{ijklm}$ be the number of consumers who prefer products (i, j, k, l) on occasions $(1, 2, 3, 4)$, respectively. For example, n_{1121} is the number of consumers who prefer product A on occasions 1, 2, and 4 and product B on occasion 3. Finally, let θ_{ijkl} be the probability that products (i, j, k, l) are preferred on occasions $(1, 2, 3, 4)$, respectively.

As in the case of binary matched pairs, i.e., two preference tests, the likelihood for binary repeated measures data, i.e., more than two tests, is given by the multinomial distribution:

$$p(\boldsymbol{n}|\boldsymbol{\theta}) \propto \prod_{i=1}^{2} \prod_{j=1}^{2} \prod_{k=1}^{2} \prod_{l=1}^{2} \theta_{ijkl}^{n_{ijkl}}, \tag{15}$$

where $0 < \theta_{ijkl} < 1$ such that $\sum_{i=1}^{2} \sum_{j=1}^{2} \sum_{k=1}^{2} \sum_{l=1}^{2} \theta_{ijkl} = 1$. The Dirichlet distribution is again an appropriate choice for a prior distribution:

$$p(\boldsymbol{\theta}) \propto \prod_{i=1}^{2} \prod_{j=1}^{2} \prod_{k=1}^{2} \prod_{l=1}^{2} \theta_{ijkl}^{\mu_{ijkl}-1}, \tag{16}$$

where $\mu_{ijkl} > 0$ are specified prior to analysis using the guidelines described in Section 2.1. This leads to a Dirichlet posterior distribution:

$$p(\boldsymbol{\theta}|\boldsymbol{n}) \propto \prod_{i=1}^{2} \prod_{j=1}^{2} \prod_{k=1}^{2} \prod_{l=1}^{2} \theta_{ijkl}^{\nu_{ijkl}-1}, \tag{17}$$

where $\nu_{ijkl} = n_{ijkl} + \mu_{ijkl}$ are the parameters of the posterior distribution.

3.1 Details of Analysis

In the case of more than two occasions, the questions of interest tend to be more complex, leading to more complex analyses. Often an exact distribution cannot be specified to compute posterior probabilities of interest and computational methods must be employed. For example, we may want to know if preference for product A is likely to increase over time. Thus, interest lies in the posterior probability

$$P_5 = P(\theta_{\cdots 1} > \theta_{\cdot\cdot 1\cdot} > \theta_{\cdot 1\cdot\cdot} > \theta_{1\cdots}|\boldsymbol{n}), \tag{18}$$

where

$$\theta_{\cdots 1} = \sum_{i=1}^{2}\sum_{j=1}^{2}\sum_{k=1}^{2}\theta_{ijk1} \tag{19}$$

is the marginal probability that product A is preferred, over product B, on occasion 4, $\theta_{\cdot\cdot 1\cdot}$ is the probability that A is preferred on occasion 3, $\theta_{\cdot 1\cdot\cdot}$ is the that A is preferred on occasion 2, and $\theta_{1\cdots}$ is the probability that A is preferred on occasion 1. We may also be interested, as in the case of only two occasions, in determining if consumers are more likely to repeat their preference than to switch:

$$P_6 = P(\theta_{1111} + \theta_{2222} > 1/2|\boldsymbol{n}). \tag{20}$$

This is the posterior probability that consumers are more likely to choose the same product on all four occasions than to switch even one time. A looser definition of consistency in product choice would allow consumers to switch product preference at most one time. Therefore, we may be interested in the posterior probability that consumers are more likely to switch at most one time than to switch more than one time:

$$P_7 = P(\theta_{1111} + \theta_{2222} + \theta_{2111} + \theta_{1222} + \theta_{2221} + \theta_{1112} + \theta_{1122} + \theta_{2211} > 1/2|\boldsymbol{n}). \tag{21}$$

There are many other questions than can be posed in the case of more than two occasions in a replicated preference test, but the exact analysis, that is, a Bayesian analysis using the exact posterior distribution, is complicated at best. For these cases, we suggest a computational approach which involves taking random samples from the Dirichlet posterior distribution. The exact steps are comparable to those of Algorithm 1. For specificity, the following algorithm can be used to compute the posterior probability in (21).

Algorithm 2 *For $r = 1, \dots, R$, do the following:*

1. *Randomly sample probabilities $\boldsymbol{\theta}^{(r)} = \{\theta_{ijkl}^{(r)} : i, j, k, l = 1, 2\}$ from a Dirichlet distribution with parameters $\boldsymbol{\nu}^{(r)} = \{\nu_{ijkl}^{(r)} : i, j, k, l = 1, 2\}$, where $\nu_{ijkl} = n_{ijkl} + \mu_{ijkl}$.*

2. *For that sample, let $s_r = 1$ if $\theta_{1111}^{(r)} + \theta_{2222}^{(r)} + \theta_{2111}^{(r)} + \theta_{1222}^{(r)} + \theta_{2221}^{(r)} + \theta_{1112}^{(r)} + \theta_{1122}^{(r)} + \theta_{2211}^{(r)} > 1/2$, and let $s_r = 0$ otherwise.*

Then the probability in (21) is estimated by the proportion of times $\theta_{1111} + \theta_{2222} + \theta_{2111} + \theta_{1222} + \theta_{2221} + \theta_{1112} + \theta_{1122} + \theta_{2211} > 1/2$ out of the R randomly generate samples:

$$\hat{P}_7 = \frac{1}{R} \sum_{r=1}^{R} s_r. \tag{22}$$

Note that both $\boldsymbol{\theta}^{(b)}$ and $\boldsymbol{\nu}^{(b)}$ are vectors of length 16 in this case. That is, there are 16 probabilities θ_{ijkl} and 16 posterior parameters ν_{ijkl}.

Because the Dirichlet distribution is related to the gamma distribution, drawing random samples from the Dirichlet distribution is relatively easy using standard statistical software which may have random number generators for the gamma distribution, if not the Dirichlet distribution itself. R code for the examples in the next section can be found in the Appendix. The event of interest is given by the inequalities specified in equations (18), (20), and (21), or any other question of interest. At least $R = 1000$ random samples is recommended to achieve reasonable precision in approximating the posterior probability of interest. As the random sampling process is computationally simple, taking more samples will not substantially increase the computation time but will increase the precision of the approximations.

3.2 Raisin Bran Data

For the raisin bran data, we applied our methodology with a noninformative prior, that is, the Dirichlet distribution with all $\mu_{ijkl} = 1$. The posterior probability that preference for product

A increases over time, equation (18), is $P_5 = 0.512$. In addition, we can use the draws from the posterior distribution to estimate the values of $\theta_{\cdots 1}$, $\theta_{\cdot\cdot 1\cdot}$, $\theta_{\cdot 1\cdot\cdot}$, and $\theta_{1\cdots}$. To estimate the value of $\theta_{\cdots 1}$, for example, we start by computing $\theta_{\cdots 1}^{(b)}$ for each of the B random draws from the posterior distribution. As this was already done in the computation of P_5, no additional work is required. These B values $\theta_{\cdots 1}^{(1)}, \ldots, \theta_{\cdots 1}^{(B)}$ approximate the posterior distribution of $\theta_{\cdots 1}$ given the data $\boldsymbol{n}$. An estimate of $\theta_{\cdots 1}$ is then computed by taking the mean or median of this posterior distribution, or simply the sample mean or sample median of the B values $\theta_{\cdots 1}^{(1)}, \ldots, \theta_{\cdots 1}^{(B)}$. We estimated $\theta_{\cdots 1}$, $\theta_{\cdot\cdot 1\cdot}$, $\theta_{\cdot 1\cdot\cdot}$, and $\theta_{1\cdots}$ using sample medians to be $\hat{\theta}_{\cdots 1} = 0.676$, $\hat{\theta}_{\cdot\cdot 1\cdot} = 0.690$, $\hat{\theta}_{\cdot 1\cdot\cdot} = 0.732$, and $\hat{\theta}_{1\cdots} = 0.760$. This reinforces that preference for product A does appear to increase over time.

In addition, the posterior probability that consumers are more likely to repeat their preference than to switch at all, equation (20), is $P_6 = 0.264$. Note that, in implementing Algorithm 2, step 1 allows us to obtain an approximation of the posterior distribution of $\theta_{1111} + \theta_{2222}$, given $\boldsymbol{n}$. A histogram or density estimate of these values can be constructed to illustrate this posterior distribution. Figure 5 shows the density estimte of this posterior distribution. The vertical line marks 0.5 so that the area to the right of the vertical line is the posterior probability of interest, P_6.

[Figure 1 about here.]

When carrying out more than two tests, that is, tests on more that two occasions, we may be interested in whether a consumer switches product preference infrequently or in whether the consumer "settles into" preferring a particular product. These are vague terms that may have different meaning to different evaluators. For the sake of illustration, we define switching infrequently as switching preference at most once, and we are interested in knowing if consumers are more likely to switch product preference infrequently:

$$P_7 = P(\theta_{1111} + \theta_{2222} + \theta_{2111} + \theta_{1222} + \theta_{1112} + \theta_{2221} + \theta_{1122} + \theta_{2211} > 1/2|\boldsymbol{n}),$$

given in Equation (21). Using a noninformative prior, and implementing an algorithm comparable to Algorithm 2, this probability is essentially $P_7 = 1$. Thus, switching product preference at most once is more likely than switching product preference more than once. Further, we define "settling into" a particular product as choosing the same product on all four occasions or on the last three of the four occasions. We are interested then in the probability that consumers are more likely to settle into a product than not:

$$P_8 = P(\theta_{1111} + \theta_{2222} + \theta_{2111} + \theta_{1222} > 1/2|\boldsymbol{n}). \tag{23}$$

Again, using a noninformative prior, so that the data drives the computation, this probability is also essentially $P_8 = 1$. In both cases, graphs of the posterior distributions show that the entire distribution is to the right of 0.5.

4 No Preference Option

Suppose that two products are compared on each of two occasions and that a no-preference option is available. That is, a consumer may choose product A or product B or may specify no preference. Alternatively, suppose that three (or more) products are compared on each occasion. Each of these can be analyzed in a similar manner as in the previous section using the multinomial-Dirichlet model.

Consider a replicated preference test in which consumers are asked to specify their preference for one of two products or to specify no preference on two different occasions. Let the binary variable y_{ijk} equal 1 if consumer k chooses i on occasion 1 and chooses j on occasion 2, where $i, j = 1, 2, 3$ with 1 denoting preference for product A, 2 denoting product B, and 3 denoting no preference, and n is the total number of consumers. Thus, $y_{328} = 1$ if conumer 8 indicated no preference on the first occasion and preference for product B on the second. Also, let $n_{ij} = \sum_{k=1}^{n} y_{ijk}$ be the number of consumers who choose i on occasion 1 and j on occasion 2. For example, n_{32} is the number of consumers specifying no preference on the first occasion and preference for product B on the second. Finally, let θ_{ij} denote the probability that i is chosen on occasion 1 and j is chosen on occasion 2. This notation is a simple extension of that used in the forced-choice case in which the indices accommodate the no-preference option. The same notation can be used for forced-choice replicated preference tests with three products and two occasions. We could further extend this notation to include more than two preference tests as in Section 3.

As in the previous two sections, the likelihood is given by a multinomial distribution:

$$p(\boldsymbol{n}|\boldsymbol{\theta}) \propto \prod_{i=1}^{3}\prod_{j=1}^{3}\prod_{k=1}^{n}\theta_{ij}^{y_{ijk}} = \prod_{i=1}^{3}\prod_{j=1}^{3}\theta_{ij}^{n_{ij}}, . \tag{24}$$

The primary difference between this likelihood and the one for forced-choice preference tests is the number of parameters, θ_{ij}. In this case of two occasions and two products with a no-preference options, there are nine parameters instead of the four parameters in the forced-choice scenario. The Dirichlet prior distribution is similarly given by

$$p(\boldsymbol{\theta}) \propto \prod_{i=1}^{3}\prod_{j=1}^{3}\theta_{ij}^{\mu_{ij}-1}, \tag{25}$$

where $0 < \theta_{ij} < 1$ for $i, j = 1, 2, 3$ such that $\sum_{i=1}^{3}\sum_{j=1}^{3}\theta_{ij} = 1$. This again leads to a Dirichlet posterior distribution:

$$p(\boldsymbol{\theta}|\boldsymbol{n}) \propto \prod_{i=1}^{3}\prod_{j=1}^{3}\theta_{ij}^{\nu_{ij}-1}, \tag{26}$$

where $\nu_{ij} = n_{ij} + \mu_{ij}$.

4.1 Details of Analysis

As in Section 3.1, the analysis involves defining the questions of interest and using random sampling from the posterior distribution (26). For example, we may want to know if consumers are more likely to be consistent in their preferences than not or if consumers are more likely to prefer product A more than product B on the second occasion, regardless of their choices on the first occasion. We may also be interested in assessing the probability that consumers can even distinguish between the two products. Answering this question is facilitated by the use of the no-preference option.

Whatever the question of interest is, in terms of the parameters θ_{ij}, we proceed as in Section 3.1. That is, we take many random samples from the appropriate Dirichlet posterior distribution and estimate the posterior probability of interest by the proportion of samples that satisfy the event of interest. R code is again included in the Appendix for the examples that follow.

4.2 Beer Data

In an effort to distinguish between discriminators and non-discriminators, Greenberg and Collins (1966) conducted double preference tests, that is, replicated preference tests on two occasions, both with and without the no-preference option, as well as a triangle test. Here we are interested in applying our Bayesian methodology to the double preference test with the no-preference option. In their study, a total of $n = 617$ male beer drinkers in the New York area were given a beer preference taste test on two occasions. On each occasion, the men were asked to specify their preference for beer A or beer B or to indicate no preference. The data, given in Table 4, are observed counts computed from the percentages given in the Greenberg and Collins (1966) paper.

[Table 4 about here.]

Greenberg and Collins (1966) define non-discriminators, as observed from the double preference test, as those with inconsistent preferences in two tests or no preference in one or both tests. Potential discriminators are then defined to be those with consistent preference in both tests, that is, choosing product A both times or choosing product B both times. Recall that θ_{11} is the probability of choosing product A on both occasions and θ_{22} is the probability of choosing product B on both occasions. To determine if consumers are more likely to be discriminators than not, we want to compute the posterior probability of choosing consistently more often than not, that is,

$$P_9 = P(\theta_{11} + \theta_{22} > 1/2|\boldsymbol{n}). \tag{27}$$

The algorithm for estimating this probability follows. The R code can be found in the Appendix.

Algorithm 3 *For $r = 1, \ldots, R$, do the following:*

1. *Randomly sample probabilities $\boldsymbol{\theta}^{(r)} = \{\theta_{ij}^{(r)} : i, j = 1, 2, 3\}$ from a Dirichlet distribution with parameters $\boldsymbol{\nu}^{(r)} = \{\nu_{ij}^{(r)} : i, j = 1, 2, 3\}$, where $\nu_{ij} = n_{ij} + \mu_{ij}$.*

2. *For that sample, let $s_r = 1$ if $\theta_{11}^{(r)} + \theta_{22}^{(r)} > 1/2$, and let $s_r = 0$ otherwise.*

Then the probability in (27) is estimated by the proportion of times $\theta_{11} + \theta_{22} > 1/2$ *out of the* R *randomly generate samples:*

$$\hat{P}_9 = \frac{1}{R} \sum_{r=1}^{R} s_r. \tag{28}$$

Choosing a non-informative Dirichlet prior, i.e., $\nu_{ij} = 1$ for $i, j = 1, 2, 3$, we estimate (27) to be $P_9 = 0.006$. Thus, consumers are not likely to be discriminators, defining discriminators as Greenberg and Collins (1966). A plot of the posterior distribution, created from the 1000 values of $\theta_{11}^{(b)} + \theta_{22}^{(b)}$ obtained via random sample as described in Algorithm 3, is show in Figure 2. The shaded area represents the posterior probability of interest, P_9.

[Figure 2 about here.]

5 Conclusion

This paper proposes a straightforward Bayesian approach to analyzing data from replicated preference tests. The methods discussed extend those of Altham (1971) who provided a Bayesian version of McNemar's test for binary matched pairs data. First, by considering different prior parameters than Altham (1971), our methods extend the Bayesian methodology to answer questions specific to forced-choice replicated preference testing with two tests or occasions. Furthermore, we broaden the scope of the methodology by considering more than two tests or occasions. Finally, we allow for the no-preference option to be included in replicated preference testing.

Posterior probabilities, which are very easy to interpret, were employed to answer various questions of interest regarding the preferences of consumers. In some cases, exact posterior distributions could be specified, but simple computational methods could be implemented when necessary. It should be noted that Bayes factors, which allow for Bayesian hypothesis testing, could also be used in this setting to compare two different models. The same computational issues would be encountered as when computing posterior probabilities. In addition, the conclusions provided by Bayes factors are limited, and their interpretation is less clear. We, therefore, prefer the computation of posterior probabilities in this setting.

Some questions go unanswered with this work. For example, the Bayesian methods proposed here are limited in their ability to examine changes in preference over time. Bayesian methods that focus on changes over time are currently being developed. This work may also prove useful in developing Bayesian methods for replicated difference tests.

Appendix

R Code for Section 2 Example: Cola Data

This is the R code for entering the cola data from Wilke, Cochrane, and Chambers IV (2006), for computing the posterior probability P_1, and for implementing Algorithm 1 to compute posterior probability P_3. The noninformative prior is used.

```
# Entering cola data.
# 1 = cola A; 0 = cola B
cola<-matrix(c(rep(c(1,1,1,1),65),
          rep(c(1,1,1,0),17),
          rep(c(1,1,0,1),24),
          rep(c(1,0,1,1),19),
          rep(c(0,1,1,1),19),
          rep(c(1,1,0,0),16),
          rep(c(1,0,1,0),11),
          rep(c(0,1,1,0),14),
          rep(c(1,0,0,1),15),
          rep(c(0,1,0,1),9),
          rep(c(0,0,1,1),17),
          rep(c(1,0,0,0),9),
          rep(c(0,1,0,0),12),
          rep(c(0,0,1,0),20),
          rep(c(0,0,0,1),8),
```

```
          rep(c(0,0,0,0),21)),ncol=4,byrow=T)
# only need the third and fourth tests
y<-cola[,c(3,4)]
dimnames(y)<-list(NULL,c("time3","time4"))

# Recode data so that 1 = cola A, 2 = cola B
y[y==0]<-2

# Specifiy parameters for prior distribution.  Noninformative prior used here.
mu11<-1; mu12<-1; mu21<-1; mu22<-1

# Summarize data and compute parameters for posterior distribution.
n<-table(as.data.frame(y))
n11<-n[1,1]; n12<-n[1,2]; n21<-n[2,1]; n22<-n[2,2]
v11<-mu11+n11; v12<-mu12+n12; v21<-mu21+n21; v22<-mu22+n22
v<-matrix(c(v11,v12,v21,v22),2,2,byrow=T)

# Compute posterior probability P1.
p1<-pbeta(0.5,v12,v21)

# Compute posterior probability P3 via Algorithm 1.
R<-10000
z<-t(replicate(R,rgamma(4,v,1)))
sum.z<-apply(z,1,sum)
draws<-t(z/sum.z)          # matrix with B samples from Dirichlet posterior
t3<-rep(1:2,2)             # product labels for test 3
t4<-sort(t1)               # product labels for test 4
theta.At3<-apply(draws[t3==1,],2,sum) # compute theta11+theta12 for each sample
```

```
theta.At4<-apply(draws[t4==1,],2,sum) # compute theta11+theta21 for each sample
p3<-mean((theta.At3<0.5)&(theta.At4>0.5))
```

R Code for Section 3 Example: Raisin Bran Data

This is the R code for entering the raisin bran data from Wilke, Cochrane, and Chambers IV (2006) and implementing Algorithm 2. The noninformative prior is used.

```
# Entering raisin bran data.
# 1 = raisin bran A; 0 = raisin bran B
rb<-matrix(c(rep(c(1,1,1,1),139),
          rep(c(1,1,1,0),6),
          rep(c(1,1,0,1),13),
          rep(c(1,0,1,1),16),
          rep(c(0,1,1,1),28),
          rep(c(1,1,0,0),5),
          rep(c(1,0,1,0),10),
          rep(c(0,1,1,0),8),
          rep(c(1,0,0,1),9),
          rep(c(0,1,0,1),6),
          rep(c(0,0,1,1),13),
          rep(c(1,0,0,0),11),
          rep(c(0,1,0,0),8),
          rep(c(0,0,1,0),7),
          rep(c(0,0,0,1),12),
          rep(c(0,0,0,0),14)),ncol=4,byrow=T)
y<-rb
colnames(y)<-c("time1","time2","time3","time4")

# Recode data so that 1 = raisin bran A, 2 = raisin bran B
```

```
y<-ifelse(y==0,2,1)

# Specifiy parameters for prior distribution.  Noninformative prior used here.
mu<-array(1,dim=c(2,2,2,2))

# Summarize data and compute parameters for posterior distribution.
n<-table(as.data.frame(y))
v<-mu+n

# Compute posterior probability P7 via Algorithm 2.
R<-1000
z<-t(replicate(R,rgamma(16,v,1)))
sum.z<-apply(z,1,sum)
draws<-t(z/sum.z)           # matrix with B samples from Dirichlet posterior

t1<-rep(1:2,8)                    # product labels for test 1
t2<-rep(c(1,1,2,2),4)             # product labels for test 3
t3<-rep(c(1,1,1,1,2,2,2,2),2)     # product labels for test 3
t4<-c(rep(1,8),rep(2,8))          # product labels for test 4

lowswitch<-draws[(t1==1)&(t2==1)&(t3==1)&(t4==1),]
          +draws[(t1==2)&(t2==2)&(t3==2)&(t4==2),]
          +draws[(t1==2)&(t2==1)&(t3==1)&(t4==1),]
          +draws[(t1==1)&(t2==2)&(t3==2)&(t4==2),]
          +draws[(t1==1)&(t2==1)&(t3==1)&(t4==2),]
          +draws[(t1==2)&(t2==2)&(t3==2)&(t4==1),]
          +draws[(t1==1)&(t2==1)&(t3==2)&(t4==2),]
          +draws[(t1==2)&(t2==2)&(t3==1)&(t4==1),]
```

```
p7<-mean(lowswitch>0.5)
```

R Code for Section 4 Example: Beer Data

This is the R code for entering the beer data from Greenberg and Collins (1966) and implementing Algorithm 3. The noninformative prior is used.

```
# Entering beer data.
# 1 = beer A, 2 = beer B, 3 = no preference
beer<-matrix(c(rep(c(1,1),148),
             rep(c(1,2),123),
             rep(c(1,3),12),
             rep(c(2,1),142),
             rep(c(2,2),130),
             rep(c(2,3),20),
             rep(c(3,1),12),
             rep(c(3,2),12),
             rep(c(3,3),20)),ncol=2,byrow=T)
y<-beer

# Specifiy parameters for prior distribution.  Noninformative prior used here.
mu11<-1; mu12<-1; mu13<-1; mu21<-1; mu22<-1; mu23<-1; mu31<-1; mu32<-1; mu33<-1

# Summarize data and compute parameters for posterior distribution.
n<-t(table(as.data.frame(y)))
n11<-n[1,1]; n12<-n[1,2]; n13<-n[1,3]
n21<-n[2,1]; n22<-n[2,2]; n23<-n[2,3]
n31<-n[3,1]; n32<-n[3,2]; n33<-n[3,3]
v11<-mu11+n11; v12<-mu12+n12; v13<-mu13+n13
v21<-mu21+n21; v22<-mu22+n22; v23<-mu23+n23
```

```
v31<-mu31+n31; v32<-mu32+n32; v33<-mu33+n33
v<-c(v11,v12,v13,v21,v22,v23,v31,v32,v33)

# Compute posterior probability P9 via Algorithm 3.
R<-1000
z<-t(replicate(R,rgamma(9,v,1)))
sum.z<-apply(z,1,sum)
draws<-t(z/sum.z)      # matrix with B samples from Dirichlet posterior
t1<-sort(rep(1:3,3))   # product labels for test 1
t2<-rep(1:3,3)         # product labels for test 2
noswitch<-draws[(t1==1)&(t2==1),]+draws[(t1==2)&(t2==2),]
p9<-mean(noswitch>0.5)
```

References

AGRESTI, A. and HITCHCOCK, D. B. 2005. Bayesian inference for categorical data analysis. Stat. Method Appl. *14*, 297–330.

ALTHAM, P. M. 1971. The analysis of matched proportions. Biometrika. *58*, 561–576.

BI, J. 2003. Difficulties and a way out: a Bayesian approach for sensory difference and preference tests. J. Sensory Stud. *18*, 1–18.

COCHRANE, C.-Y. C., DUBNICKA, S., and LOUGHIN, T. 2005. Comparison of methods for analyzing replicated preference tests. J. Sensory Stud. *20*, 484–502.

ENNIS, D. M. and BI, J. 1998. The beta-binomial model: accounting for inter-trial variation in replicated difference and preference tests. J. Sensory Stud. *13*, 389–412.

GREENBERG, A. and COLLINS, S. 1966. Paired comparison taste tests: some food for thought. J. Marketing Res. *3*, 76–80.

WILKE, K. D., COCHRANE, C.-Y. C. and CHAMBERS IV, E. 2006. Multiple preference tests can provide more information on consumer preferences. J. Sensory Stud. *21*, 612–625.

List of Figures

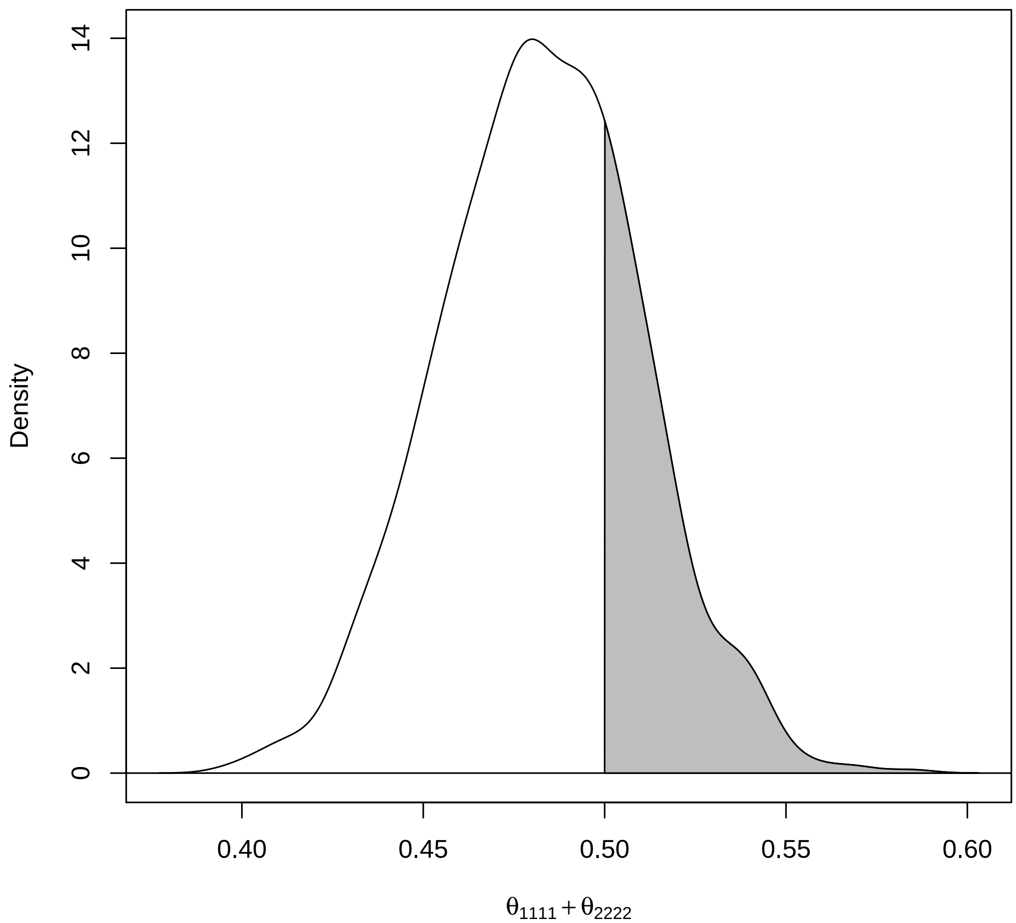

Figure 1: Posterior Distribution of $\theta_{1111} + \theta_{2222}$ for Raisin Bran Example

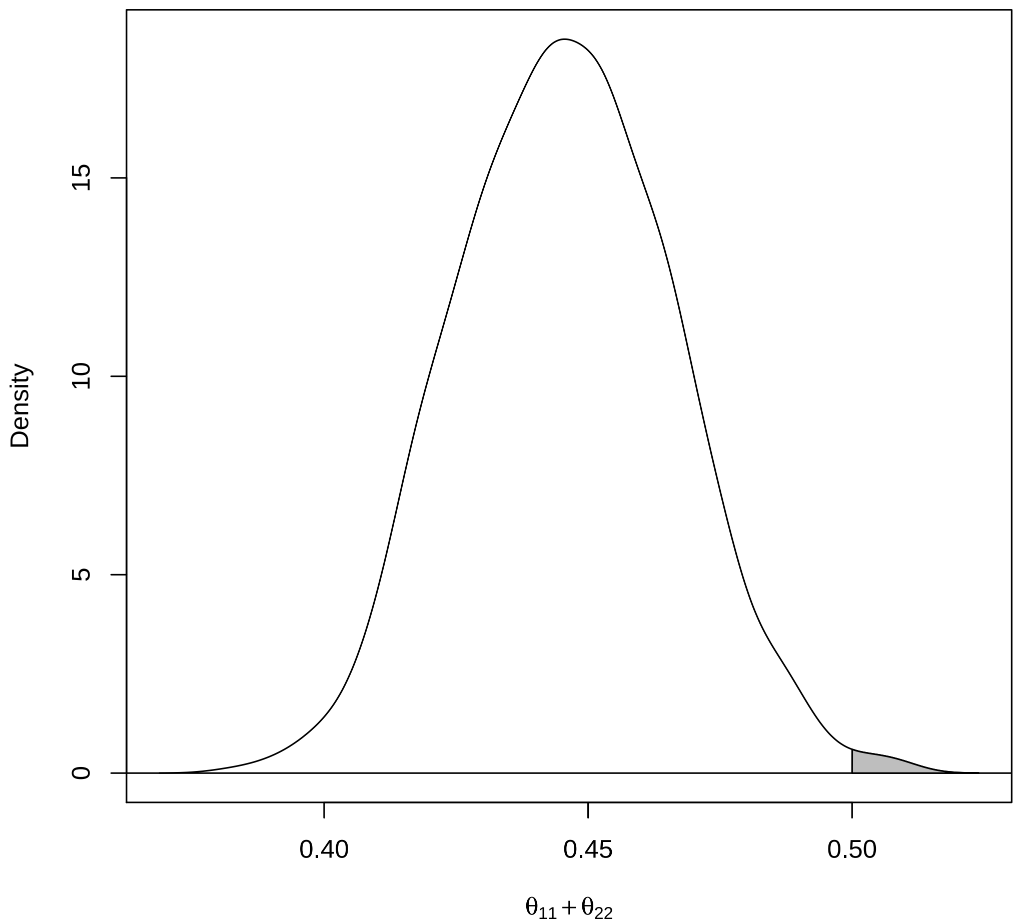

Figure 2: Posterior Distribution of $\theta_{11} + \theta_{22}$ for Beer Example

List of Tables

Table 1: Cola Replicated Forced-Choice Preference Test – Tests 3 and 4

	Test 4		
Test 3	Preferred A	Preferred B	Total
Preferred A	120	62	182
Preferred B	56	58	114
Total	176	120	296

Table 2: Priors for Cola Replicated Forced-Choice Preference Test

Prior	μ_{11}	μ_{12}	μ_{21}	μ_{22}	Explanation
A	1	1	1	1	noninformative
B	55	15	15	15	prior proportions similar to data
C	15	35	35	15	consumers more likely to switch preference than to repeat; switching preference from A to B equally likely as B to A; preferring A on both tests equally likely as preferring B on both; preferring A equally likely as preferring B on each test
D	55	5	35	15	consumers more likely to repeat product preference than switch; switching from B to A is more likely that A to B; preferring A on both tests more likely than preferring B on both; A is preferred more than B on each test
E	35	5	35	35	consumers more likely to repeat product preference than switch; switching from B to A is more likely that A to B; preferring A on both tests equally likely as preferring B on both; B is preferred more on time 3; A is preferred more on time 4
F	38	10	17	35	consumers more likely to repeat product preference than switch; switching from B to A is more likely that A to B; preferring A on both tests more likely than preferring B on both; B is preferred more on time 3; A is preferred more on time 4
G	30	15	35	20	consumers more likely to repeat product preference than switch; switching from B to A is more likely that A to B; preferring A on both tests more likely than preferring B on both; B is preferred more on time 3; A is preferred more on time 4
H	35	25	45	5	consumers more likely to switch preference than to repeat; switching from B to A is more likely that A to B; preferring A on both tests more likely than preferring B on both; A is preferred more than B on each test

Table 3: Posterior Probabilities for Cola Replicated Forced-Choice Preference Test

Prior	P_1	P_2	P_3	P_4
A	0.2912	0.9998	0.0000	0.9995
B	0.3104	0.9999	0.0000	1.0000
C	0.3305	0.8429	0.0003	0.9975
D	0.9725	0.9999	0.0001	0.9999
E	0.9725	0.9999	0.0295	0.9705
F	0.5332	1.0000	0.0003	0.9993
G	0.8607	0.9988	0.0015	0.9985
H	0.8470	0.9320	0.0001	0.9999

Table 4: Beer Double Preference Test with No-Preference Option

	Test 2			
Test 1	Preferred A	Preferred B	No Preference	Total
Preferred A	148	123	12	283
Preferred B	142	130	20	292
No Preference	12	12	20	44
Total	302	265	52	617