

# Flexible Bayesian Multiple Comparison Adjustment Using Dirichlet Process and Beta-Binomial Model Priors

Don van den Bergh\* and Fabian Dablander\*

*Department of Psychological Methods, University of Amsterdam*

July 19, 2022

## Abstract

Researchers frequently wish to assess the equality or inequality of groups, but this comes with the challenge of adequately adjusting for multiple comparisons. Statistically, all possible configurations of equality and inequality constraints can be uniquely represented as partitions of the groups, where any number of groups are equal if they are in the same partition. In a Bayesian framework, one can adjust for multiple comparisons by constructing a suitable prior distribution over all possible partitions. Inspired by work on variable selection in regression, we propose a class of flexible beta-binomial priors for Bayesian multiple comparison adjustment. We compare this prior setup to the Dirichlet process prior suggested by Gopalan and Berry (1998) and multiple comparison adjustment methods that do not specify a prior over partitions directly. Our approach to multiple comparison adjustment not only allows researchers to assess all pairwise (in)equalities, but in fact all possible (in)equalities among all groups. As a consequence, the space of possible partitions grows quickly — for ten groups, there are already 115,975 possible partitions — and we set up a stochastic search algorithm to efficiently explore the space. Our method is implemented in the Julia package *EqualitySelection*, and we illustrate it on examples related to the comparison of means, variances, and proportions.

## 1 Introduction

Assessing the equality or inequality of groups is a key problem in science and applied settings. If a confirmatory hypothesis is lacking, a standard approach is to first test whether all groups are equal and, if they are not, engage in multiple post-hoc comparisons. A large swathe of multiple comparisons techniques to guard against inflated false-positive errors exist in classical statistics, dating back to the work of John Tukey and others (e.g., Benjamini & Braun, 2002; Rao, 2009). From a Bayesian perspective, the problem of multiple comparisons can be addressed by changing the model prior (e.g., Berry & Hochberg, 1999; de Jong, 2019; Jeffreys, 1961; Westfall et al., 1997), an approach that has found prominent application in variable selection for regression (e.g., Scott & Berger, 2006, 2010). Here, we focus on a Bayesian multiplicity adjustment for testing the (in)equality between groups. Statistically, all possible configurations of equality and inequality constraints can be uniquely represented as partitions of the groups, where two groups are equal if they are in the same partition. In a Bayesian framework, one can adjust for multiple comparisons by constructing a suitable prior distribution over all possible partitions. This allows the researcher to explore the set of all possible equality and inequality relations among the groups while penalizing for multiple comparisons.

---

\*These authors share first authorship.

The first to propose a prior over all partitions to adjust for multiple hypotheses testing were, to our knowledge, Gopalan and Berry (1998), who suggested the Dirichlet process prior. Here, we propose a class of flexible beta-binomial priors for Bayesian multiple comparison adjustment, inspired by work on variable selection in regression (Scott & Berger, 2006, 2010) and explore its properties vis-à-vis previous work on multiple comparisons. More specifically, the current paper is structured as follows. In Section 2, we set up the problem and describe the Pólya urn scheme from which a number of priors can be derived. We characterize three such priors — the Dirichlet process, the beta-binomial, and the uniform prior — and outline our methodology in Section 3. In Section 4 we contrast the three priors, illustrate our method on a simulated example, and present a simulation study assessing the multiplicity adjustment of each prior. We also assess the method proposed by Westfall et al. (1997) and an uncorrected testing procedure based only on pairwise Bayes factors. As the space of possible partitions grows quickly — for ten groups, there are already 115,975 possible partitions — we set up a stochastic search algorithm to efficiently explore the space. Our method is implemented in Julia and available in the *EqualitySelection* package from <https://github.com/vandenman/EqualitySelection>. In Section 5, we apply our method to examples related to the comparison of proportions and variances. We conclude in Section 6.

## 2 Preliminary Remarks

In this section, we set up the hypothesis testing problem, discuss the relation between partitions and models, and describe Pólya’s urn scheme that will unify the presentation of the priors in the following section.

### 2.1 Problem Setup

Our goal is to adjust for multiple comparisons in a flexible manner. Multiple comparisons are not a problem if we wish to compare only two hypotheses, denoted as  $\mathcal{H}_0$  and  $\mathcal{H}_1$ . The Bayes factor quantifies how strongly we should update our prior beliefs about  $\mathcal{H}_0$  relative to  $\mathcal{H}_1$  after observing the data (Kass & Raftery, 1995; Ly et al., 2016). Let group  $j$  consist of  $n_j$  observations  $\vec{y}_j = \{y_{j1}, \dots, y_{jn_j}\}$  for  $j \in \{1, \dots, K\}$  and  $i \in \{1, \dots, n_j\}$ , and let  $\vec{y} = \{\vec{y}_1, \dots, \vec{y}_K\}$ . The Bayes factor is given by:

$$\underbrace{\frac{p(\mathcal{H}_0 | \vec{y})}{p(\mathcal{H}_1 | \vec{y})}}_{\text{Posterior odds}} = \underbrace{\frac{p(\vec{y} | \mathcal{H}_0)}{p(\vec{y} | \mathcal{H}_1)}}_{\text{Bayes factor}} \times \underbrace{\frac{p(\mathcal{H}_0)}{p(\mathcal{H}_1)}}_{\text{Prior odds}} , \quad (1)$$

which does not depend on the number of hypotheses a researcher wishes to test.

A principled way to account for multiplicity is by adjusting the prior probability of the hypotheses (e.g., Jeffreys, 1961; Westfall et al., 1997). Suppose a researcher is interested in comparing  $K$  groups, parameterized by  $\vec{\theta} = (\theta_1, \dots, \theta_K)$ . She is not only interested in whether all parameters are equal ( $\mathcal{H}_0$ ) or whether they are unequal ( $\mathcal{H}_1$ ), but also which pairs of parameters are equal or not. In the language of classical statistics, she is interested in post-hoc comparisons. We focus on a Bayesian solution to this problem in the current paper. More specifically, going beyond classical testing, we consider the problem of assessing all possible equalities and inequalities between the groups. In general terms, the inference problem is:

$$\begin{aligned} \rho &\sim \pi_\rho(\cdot) \\ \vec{\theta} | \rho &\sim \pi_{\vec{\theta}}(\cdot) \\ f(\vec{y}; \vec{\theta}, \rho) &= \prod_{j=1}^K g(\vec{y}_j; \theta_j, \phi) , \end{aligned}$$

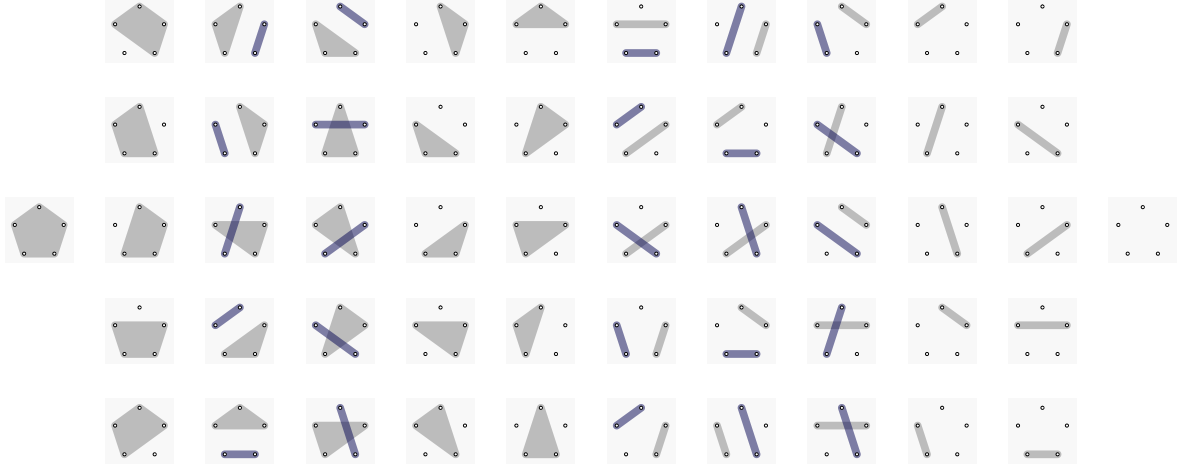


Figure 1: All 52 possible models given  $K = 5$ , represented as partitions. Circles represent individual parameters and shaded regions indicate which parameters are equal.

where  $\rho$  is a partition,  $\phi$  is a nuisance parameter (in case it exists), and  $f$  and  $g$  are the likelihood functions. Using the posterior distribution of  $\vec{\theta}$ , we have that:

$$\begin{aligned} p(\mathcal{H}_0 \mid \vec{y}) &= p(\theta_1 = \theta_2 = \dots = \theta_K \mid \vec{y}) \\ p(\mathcal{H}_1 \mid \vec{y}) &= p(\theta_1 \neq \theta_2 \neq \dots \neq \theta_K \mid \vec{y}) . \end{aligned}$$

There are many more possible hypotheses, however, depending on the combination of equalities and inequalities. We can represent those as partitions, as we detail in the next section.

## 2.2 Partitions

The space of possible equality constraints for some parameter vector  $\vec{\theta} = (\theta_1, \dots, \theta_K)$  of size  $K$  is equivalent to the partitions of that vector. For example, for  $K = 3$  the model that states  $\theta_1 = \theta_2 \neq \theta_3$  is equivalent to the partition  $\{\{\theta_1, \theta_2\}, \{\theta_3\}\}$ . The space of possible models for  $K = 5$  is shown in Figure 1. The correspondence between (in)equality constraints and partitions is useful as partitions have been studied extensively in combinatorics. Given  $K$  parameters, the number of partitions of size  $j$  is given by the Stirling numbers of the second kind, denoted  $\left\{ \begin{smallmatrix} K \\ j \end{smallmatrix} \right\}$ . The total number of partitions is given by the  $K^{\text{th}}$ -Bell number, which is defined as a sum over the Stirling numbers:

$$B_K = \sum_{j=0}^K \left\{ \begin{smallmatrix} K \\ j \end{smallmatrix} \right\} . \quad (2)$$

The Bell numbers grow very quickly, with the number of partitions for a vector  $\vec{\theta}$  of size 10 being  $B_{10} = 115,975$ .

The Stirling numbers and Bell numbers can be generalized to the  $r$ -Stirling (Broder, 1984) and  $r$ -Bell numbers (Mezo, 2011), respectively. These generalizations help to construct conditional distributions, as we will see later. The  $r$ -Stirling numbers  $\left\{ \begin{smallmatrix} K \\ j \end{smallmatrix} \right\}_r$  give the number of partitions of size  $j$  given  $K+r$  groups such that the first  $r$  parameters are all in distinct subsets. The  $r$ -Bell numbers give the total number of partitions given  $K$  parameters where the first  $r$

parameters are in distinct subsets. Specifically, we have:

$$\left\{ \begin{matrix} K \\ j \end{matrix} \right\}_r = \sum_{i=0}^K \binom{K}{i} \left\{ \begin{matrix} i \\ j \end{matrix} \right\}_r r^{K-i} \quad (3)$$

$$B_{K,r} = \sum_{i=0}^K \left\{ \begin{matrix} K+r \\ i+r \end{matrix} \right\}_r. \quad (4)$$

Note that  $\left\{ \begin{matrix} K \\ j \end{matrix} \right\}_1 = \left\{ \begin{matrix} K \\ j \end{matrix} \right\}$  and that  $B_{K,0} = B_K$ . Both the  $r$ -Stirling and  $r$ -Bell numbers are defined through recurrence relations, although explicit expressions exist which are easier to compute for large values; see Broder (1984) and Mezo (2011) for details.

### 2.3 Urn Schemes

We can represent the different partitions using an urn with  $K$  different balls labeled 1 through  $K$ . For each parameter  $\theta_j$ , a ball  $b_j$  is drawn from the urn with  $b_j \in \{1, \dots, K\}$ . If two drawn balls are equal,  $b_i = b_j$ , then the two parameters are assigned to the same subset of the partition, that is, the two parameters  $\theta_i$  and  $\theta_j$  are equal if  $b_i = b_j$ . Note that different draws from an urn can represent the same partition. For example, the draws  $(1, 1, 2)$  and  $(3, 3, 1)$  both represent the partition  $\{\{\theta_1, \theta_2\}, \{\theta_3\}\}$ . The prior distributions introduced in the next sections assign probabilities to the unique partitions. Note that the prior probability of a particular draw can be obtained by dividing the probability of the corresponding partition by the total number of draws that correspond to that partition. The total number of draws that represent the same partition is given by  $d! \binom{K}{d}$  where  $d$  is the number of non-empty subsets of a particular draw.

Although the urn consists of  $K$  different balls, the event of interest is whether the next ball drawn equals one of the balls already drawn — in other words, whether an equality or inequality is introduced. This event reduces the urn to a Pólya urn. All prior distributions discussed below are related to the Pólya urn. Specifically, the joint prior distribution on  $(\theta_1, \dots, \theta_K)$  is characterized by a (generalized) Pólya urn such that:

$$\theta_K \mid \theta_1, \dots, \theta_{K-1} \sim \begin{cases} \zeta_j & \text{with probability } P_\pi \\ \theta_j^* & \text{with probability } 1 - P_\pi \end{cases}, \quad (5)$$

where  $\zeta_j$  denotes a new value for  $\theta_K$  (with  $\theta_1 = \zeta_1$ ) and  $\theta_j^*$  denotes a value equal to any previously observed value. We characterize the priors we discuss in the next section in terms of (5), which is known as a *prediction rule* (e.g., Ishwaran & James, 2001); in terms of the induced prior over partitions; and in terms of their penalty for multiplicity.

## 3 Methodology

Let  $\vec{\theta}^* = (\theta_1^*, \dots, \theta_r^*)$  denote the vector of unique population parameters out of  $\vec{\theta} = (\theta_1, \dots, \theta_K)$ ,  $\vec{\theta}_{-j}$  the vector of parameters without parameter  $\theta_j$ , and the number of repeats of  $\theta_j^*$  as  $n_j^*$ . Let  $\rho$  denote a partition and  $|\rho|$  its size. For example, if  $\rho = \{\{\theta_1, \theta_2\}, \{\theta_3\}\}$ , then  $|\rho| = 2$ . Similarly, for this example  $\vec{\theta}^* = (\theta_1^*, \theta_2^*)$  and  $n^* = (2, 1)$ . In the next sections, we discuss and contrast a number of priors.

### 3.1 Dirichlet Process Prior

The Dirichlet process (DP) is a distribution over distributions (Ferguson, 1973). We say that  $\mathcal{G} \sim \text{DP}(\alpha, \mathcal{K})$  is distributed according to a DP if its marginal distributions are Dirichlet distributed, where  $\alpha$  is a concentration parameter and  $\mathcal{K}$  is the base distribution, which will depend on the application; for details, see for example Teh (2010). The DP can be understood

as the infinite-dimensional generalization of the Dirichlet distribution, which makes it popular for mixture modeling (e.g., Rasmussen et al., 1999). Our modeling approach is similar to mixture modeling, except that we do not cluster data but parameters — a cluster corresponds to a partition. The prediction rule of the DP is given by (e.g., Blackwell & MacQueen, 1973; Ishwaran & James, 2001):

$$\theta_{j+1} \mid \theta_1, \dots, \theta_j \sim \begin{cases} \mathcal{K} & \text{with probability } \frac{\alpha}{\alpha+j-1} \\ \text{Categorical}(\theta_1^*, \dots, \theta_r^* \mid n_1^*, \dots, n_r^*) & \text{else ,} \end{cases} \quad (6)$$

where  $\alpha$  is the concentration parameter and the base distribution of the DP depends on the application (see Section 5). In other words, we draw a new value for  $\theta_j$  from  $\mathcal{K}$  with probability  $\alpha/(\alpha+j-1)$ , or else set it to a previously observed value. The particular value  $\theta_j^*$  the parameter  $\theta_j$  is set to is proportional to the number of times  $\theta_j^*$  was observed previously, given by  $n_j^*$ , resulting in the well-known “rich-get-richer” property (e.g., Teh, 2010).

The Dirichlet process implies a prior distribution over partitions. The prior on the partitions  $\rho$  is:

$$\pi(\rho \mid \alpha) = \frac{\alpha^{|\rho|} \Gamma(\alpha)}{\Gamma(n + \alpha)} \prod_{c \in \rho} \Gamma(|c|) , \quad (7)$$

where  $c$  is an element of  $\rho$ , and  $|c|$  is its size. While the Dirichlet process features the infinite-dimensional object  $\mathcal{K}$ , the prior over partitions results from integrating it out. Hence the nonparametric model (in which the number of parameters is not fixed) implies a parametric model (in which the number of parameters is fixed) for the partitions (Quintana, 2006). This makes it usable for our purposes, where we have a fixed number of parameters.

The leftmost column in Figure 2 shows the DP prior over partitions (top) and number of inequalities (bottom) for different values of  $\alpha$ . Intuitively, one reasonable requirement for a prior in the context of penalizing multiplicity is to be monotonically decreasing in the number of partitions, which further implies a monotonically decreasing prior probability over the number of inequalities. This is the case for  $\alpha = 0.50$  (beige diamonds) as shown in the top and bottom panels, and indeed for any value  $\alpha < 1$ . The value suggested by Gopalan and Berry (1998) creates a symmetric prior over the partitions (yellow suns), implying that the model with no inequalities is a priori as likely as the model with all inequalities (in the  $K = 5$  case, this yields  $\alpha = 2.213$ ). The prior with  $\alpha = 1$  (pink stars) results in a nonincreasing prior over the number of partitions, but in an increasing prior over the number of inequalities: the model with one inequality is more likely than the model with no inequalities.

As  $\alpha \rightarrow 0$ , the prior of the model with all  $K - 1$  equalities  $\mathcal{M}_0$  (i.e., the null model) converges to one, while as  $\alpha \rightarrow \infty$ , the prior of the model with  $K - 1$  inequalities  $\mathcal{M}_{B_K}$  (i.e., the full model) converges to one. For prior elicitation, Gopalan and Berry (1998) note that  $\alpha$  is determined by specifying two of either  $P(\mathcal{M}_0)$ ,  $P(\mathcal{M}_{B_K})$ , or their ratio, since  $P(\mathcal{M}_0) = \alpha^{(K-1)!} / \prod_{j=1}^K (\alpha + j - 1)$  and  $P(\mathcal{M}_{B_K}) = \alpha^K / \prod_{j=1}^K (\alpha + j - 1)$ ; see also Table 1.

### 3.2 Beta-binomial Prior

The beta-binomial model prior is a popular choice for stochastic search variable selection in linear regression (George & McCulloch, 1993) and Bayesian model averaging (e.g., Hinne et al., 2020; Hoeting et al., 1999). It states that the prior probability of including  $j$  predictors out of a total of  $K$  predictors is given by:

$$\text{BB}(j \mid K, \alpha, \beta) = \binom{K}{j} \frac{B(j + \alpha, K - j + \beta)}{B(\alpha, \beta)} , \quad (8)$$

where  $\alpha$  and  $\beta$  are hyperparameters. The prior probability of a particular regression model is obtained by dividing by the number of ways  $j$  out of  $K$  predictors can be included:  $\text{BB}(j \mid K, \alpha, \beta) / \binom{K}{j}$ .

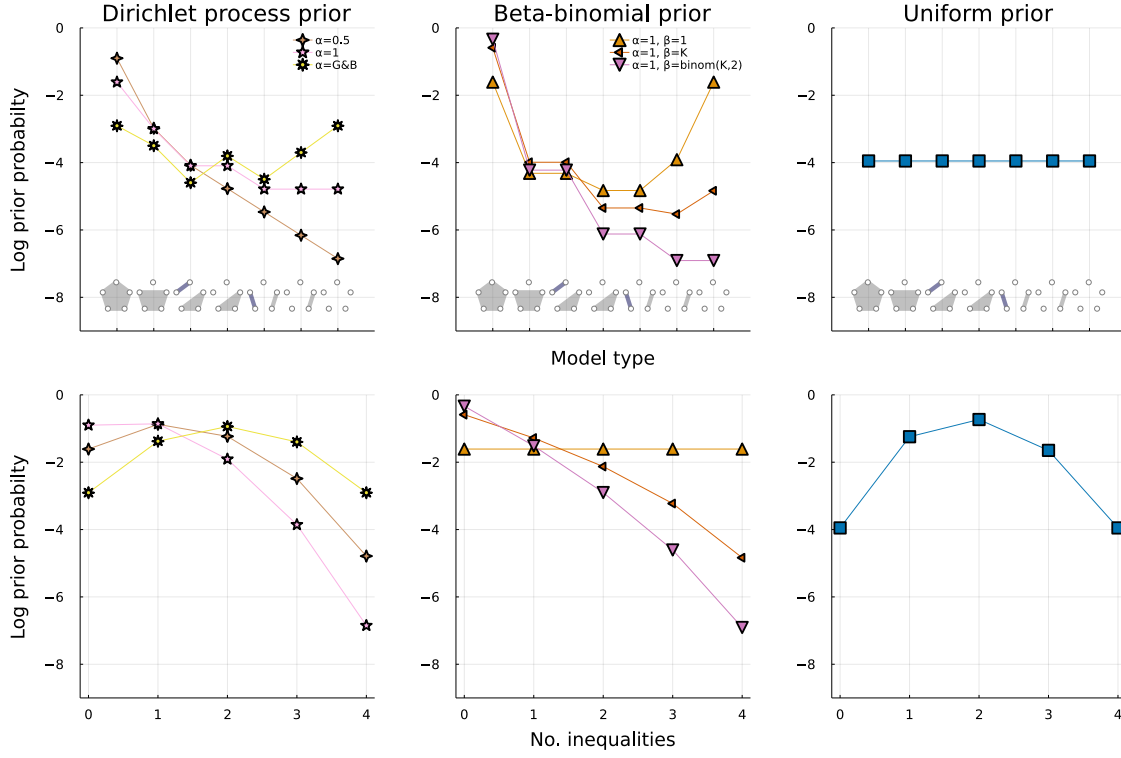


Figure 2: Top: Dirichlet process (left), beta-binomial (middle), and uniform prior (right) across distinct model types for  $K = 5$  groups and different prior parameters. Bottom: Same but for the number of inequalities across models.

	Dirichlet process prior	Beta-binomial prior	Uniform prior
Parameters	$\alpha$	$(\alpha = 1, \beta)$	$\mathbf{X}$
Prior over partitions	$\frac{\alpha^{ \rho } \Gamma(\alpha)}{\Gamma(n+\alpha)} \prod_{c \in \rho} \Gamma( c )$	$\binom{K-1}{ \rho -1} \frac{B( \rho -1+\alpha, K- \rho +\beta)}{B(\alpha, \beta) \binom{K}{ \rho }}$	$(B_K)^{-1}$
Prior monotonically decreasing	$\alpha \leq 1$	$\beta \geq K, \beta \geq \binom{K}{2}$	$\mathbf{X}$
Prior probability of null model	$\alpha(K-1)! / \prod_{j=1}^K (\alpha+j-1)$	$B(\alpha, K-1+\beta) / B(\alpha, \beta)$	$(B_K)^{-1}$
Prior probability of full model	$\alpha^K / \prod_{j=1}^K (\alpha+j-1)$	$B(K-1+\alpha, \beta) / B(\alpha, \beta)$	$(B_K)^{-1}$
Prior probability of ratio (null / full)	$\alpha(K-1)! / \alpha^K$	$B(\alpha, K-1+\beta) / B(K-1+\alpha, \beta)$	1

Table 1: Characterizations of the different priors studied in this paper. Note:  $\beta \geq K$  implies a prior decreasing in terms of the number of inequalities, but not in terms of the partitions.  $\beta \geq \binom{K}{2}$  implies both.

The beta-binomial distribution introduces a penalty for including additional predictors and in that way introduces a correction for multiplicity (Scott & Berger, 2006, 2010).

For the multiple comparison problem discussed in this paper, we consider the number of inequality constraints and use the beta-binomial prior to introduce a penalty for each additional inequality among the groups considered. For  $K$  groups, there can be a maximum of  $K - 1$  inequalities, resulting in a  $\text{BB}(i \mid K - 1, \alpha, \beta)$  prior distribution over the number of included inequalities  $i$  out of  $K$  groups. To see how this translates to a prior over the partitions  $\rho$ , note that there is a one-to-many correspondence between the number of inequalities  $i$  out of  $K$  groups and the resulting partitions  $\rho$ . For example, having  $i = 1$  inequalities with  $K = 3$  groups is consistent with the partitions  $\{\{\theta_1, \theta_2\}, \{\theta_3\}\}$ ,  $\{\{\theta_1, \theta_3\}, \{\theta_2\}\}$ , and  $\{\{\theta_2, \theta_3\}, \{\theta_1\}\}$ , all of which are of size  $|\rho| = i + 1$ . The number of partitions of size  $|\rho|$  is given, as discussed above, by the Stirling number  $\left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}$ . For the assignment of the prior probability, it is only the size of the partition (the number of inequalities) that counts. With these observations in hand, we arrive at the following (adjusted) beta-binomial prior distribution over partitions  $\rho$ :

$$\pi(\rho \mid K, \alpha, \beta) = \binom{K-1}{|\rho|-1} \frac{\text{B}(|\rho| - 1 + \alpha, K - |\rho| + \beta)}{\text{B}(\alpha, \beta) \left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}} . \quad (9)$$

The prediction rule of the beta-binomial prior is given by:

$$\theta_{j+1} \mid \theta_1, \dots, \theta_j \sim \begin{cases} \mathcal{K} & \text{with probability } P_\pi \\ \text{Categorical}(\theta_1^*, \dots, \theta_r^* \mid 1, \dots, 1) & \text{else} \end{cases} , \quad (10)$$

where

$$P_\pi = \frac{\sum_{\substack{\rho \in \mathcal{P} \\ \theta_j \notin \vec{\theta}_{-j} \subseteq \rho}} \text{BB}(\rho \mid K, \alpha, \beta)}{\sum_{\substack{\rho \in \mathcal{P} \\ \theta_j \notin \vec{\theta}_{-j} \subseteq \rho}} \text{BB}(\rho \mid K, \alpha, \beta) + \sum_{\substack{\rho \in \mathcal{P} \\ \theta_j \in \vec{\theta}_{-j} \subseteq \rho}} \text{BB}(\rho \mid K, \alpha, \beta)} , \quad (11)$$

and where  $\mathcal{P}$  denotes the set of all possible partitions. In essence, Equation (11) takes the probability of all possible partitions where  $\theta_j$  is distinct from  $\vec{\theta}_{-j}$ , conditional on  $\vec{\theta}_{-j}$  being a subset of the considered partition. The sum over all possible partitions can be simplified using the  $r$ -Stirling numbers:

$$P_\pi = \frac{\sum_{i=1}^K \text{BB}(i \mid K, \alpha, \beta) \left\{ \begin{smallmatrix} K-j+r+1 \\ i \end{smallmatrix} \right\}_{r+1}}{r \sum_{j=i}^K \text{BB}(i \mid K, \alpha, \beta) \left\{ \begin{smallmatrix} K-j+r \\ i \end{smallmatrix} \right\}_r + \sum_{i=1}^K \text{BB}(i \mid K, \alpha, \beta) \left\{ \begin{smallmatrix} K-j+r+1 \\ i \end{smallmatrix} \right\}_{r+1}} , \quad (12)$$

where  $r$  is number of unique parameters in  $\vec{\theta}$ , that is, the size of the partition.

The beta-binomial prior on the partitions and the induced prior on the number of inequalities are shown for different parameterizations in the middle column in Figure 2. For  $\alpha = \beta = 1$ , the beta-binomial distribution over the partitions has a characteristic U-shape (orange triangles). This prior specification in turn implies a uniform prior on the number of inequalities. We follow Wilson et al. (2010) who, in the context of regression, suggested to set  $\alpha = 1$  as a default so that the distribution over model size (here the number of inequalities) is nonincreasing, and to scale  $\beta = \lambda K$  with the number of groups to force the prior to be monotonically decreasing, with a default of  $\lambda = 1$  (Wilson et al., 2010). This is illustrated as the red line (leftward pointing triangles) in Figure 2 using  $\beta = 5$ . In the multiple comparison case, we additionally investigate  $\beta = \binom{K}{2}$ , which implies that the prior on the number of inequalities of individual models is nonincreasing, see Appendix A. The purple line (upside-down triangles) in Figure 2 shows a decreasing prior for  $\beta = \binom{5}{2} = 10$ . This prior assigns the least mass to models with an increasing number of inequalities compared to all others beta-binomial priors.

Figure 2 shows that the DP prior makes a distinction that the beta-binomial is, by design, not making: while the beta-binomial prior assigns the same prior mass to partitions with the same number of (in)equalities, the DP prior assigns more mass to the partition with the larger cluster. For example, the beta-binomial does not distinguish between  $\{\{\theta_1, \theta_2, \theta_3\}, \{\theta_4\}, \{\theta_5\}\}$  and  $\{\{\theta_1, \theta_2\}, \{\theta_3, \theta_4\}, \{\theta_5\}\}$ , while the DP assigns more mass to the former (see Figure 2). We return to this distinction in the discussion.

Lastly, note that for the beta-binomial prior we have that  $P(\mathcal{M}_0) = B(\alpha, K-1+\beta)/B(\alpha, \beta)$  and  $P(\mathcal{M}_{B_K}) = B(K-1+\alpha, \beta)/B(\alpha, \beta)$ . Fixing  $\alpha = 1$ , we have that as  $\beta \rightarrow \infty$ , the prior of the model with all  $K - 1$  equalities  $\mathcal{M}_0$  converges to one, while as  $\beta \rightarrow 0$ , the prior of the model with  $K - 1$  inequalities  $\mathcal{M}_{B_K}$  converges to one; see also Table 1. As with the Dirichlet process prior discussed above, one can use these relations in prior elicitation.

### 3.3 Uniform Prior

For completeness, we give a prior that is uniform over the space of partitions. The probability mass function is straightforward. All valid configurations of size  $K$  have probability  $1/B_K$ . The prediction rule of the uniform prior is given by:

$$\theta_{j+1} \mid \theta_1, \dots, \theta_j \sim \begin{cases} \mathcal{K} & \text{with probability } P_{\pi_U} \\ \text{Categorical}(\theta_1^*, \dots, \theta_r^* \mid 1, \dots, 1) & \text{else ,} \end{cases} \quad (13)$$

where

$$P_{\pi_U} = \frac{B_{K-j+1, r+1}}{B_{K-j+1, r+1} + r B_{K-j+1, r}} \quad (14)$$

Here,  $B_{K-j+1, r+1}$  counts the number of models where  $\theta_{j+1} \notin (\theta_1^*, \dots, \theta_r^*)$  conditional on  $\theta_1, \dots, \theta_j$  being assigned to  $r$  distinct subsets. Complementarily,  $B_{K-j+1, r}$  counts the number of models where  $\theta_{j+1} \in (\theta_1^*, \dots, \theta_r^*)$  conditional on  $\theta_1, \dots, \theta_j$  being assigned to  $r$  distinct subsets, which is multiplied by  $r$  as there are  $r$  subsets that  $\theta_{j+1}$  could be assigned to. Under this uniform prior, all partitions  $\rho$  are equally likely, as can be seen in the top right panel in Figure 2. Note that this uniform prior induces a non-uniform prior on the number of inequalities, as shown in the bottom right panel.

### 3.4 Posterior Model Consistency

Model selection consistency is a key desiderata that a good Bayes factor should fulfill (e.g., Bayarri et al., 2012; Consonni et al., 2018; Ly et al., 2016). In the situation of multiple models, the notion of pairwise model selection consistency needs to be extended. This extension is referred to as posterior model selection consistency. Posterior model consistency in a model class  $\mathfrak{M}$  is the convergence to one, in probability, of the posterior probabilities to the true model (e.g., Casella et al., 2009; Moreno et al., 2015). Let  $\mathcal{M}_j \in \mathfrak{M}$  be the model that instantiates the hypothesis  $\mathcal{H}_j$  that specifies the (in)equalities among  $K$  groups. The posterior probability of  $\mathcal{M}_j$  is given by:

$$p(\mathcal{M}_j \mid \mathcal{D}) = \frac{p(\mathcal{D} \mid \mathcal{M}_j)\pi(\mathcal{M}_j)}{\sum_{i=0}^{B_K} p(\mathcal{D} \mid \mathcal{M}_i)\pi(\mathcal{M}_i)} = \frac{\text{BF}_{j0}\pi(\mathcal{M}_j)}{\sum_{i=0}^{B_K} \text{BF}_{i0}\pi(\mathcal{M}_i)} . \quad (15)$$

It follows that if the Bayes factor is model selection consistent, posterior model consistency holds (see also Moreno et al., 2015, Theorem 1) — unless the prior assigns zero mass to the true model. This is not the case for any of the priors discussed above, and hence whether posterior model consistency holds depends solely on the priors on the parameters within models.

### 3.5 Stochastic Search Method

When the number of groups is small and the computation of Bayes factors is swift, one can directly compute the Bayes factors for all hypotheses. Using the priors we outlined above, one can then obtain posterior distributions over hypotheses that incorporate the desired multiplicity adjustment. The number of (in)equalities grows extremely quickly with the number of groups, however, and for larger number of groups one must rely on stochastic search methods. Moreover, while directly computing the Bayes factors results in posterior distributions over hypotheses, it does not yield posterior distributions over parameters. We therefore set up a stochastic search method that yields both, allowing researchers to incorporate uncertainty across hypotheses through model averaging (e.g., Hinne et al., 2020; Hoeting et al., 1999).

Our method is implemented in the programming language Julia (Bezanson et al., 2017). First, we implemented the prior distributions in Julia. Next, we used the library *Turing.jl*, which is designed for general-purpose probabilistic programming (Ge et al., 2018). Turing enabled us to directly reuse the distributions defined in Julia code and also provided a multitude of options for composing different MCMC samplers. We set up a Gibbs sampler that explored the posterior space in two steps. The first step used Turing’s built-in Hamiltonian Monte Carlo methods for sampling from the posterior distributions of the continuous parameters. In all models discussed here, all parameters are continuous except for the partitions. The second step used a custom Gibbs algorithm for sampling from the posterior distribution over partitions. The partitions were represented as a vector of integers denoted  $\vec{\gamma}$  that indicate partition membership. By partition membership, we mean that two parameters  $\theta_i$  and  $\theta_j$  are in the same partition if and only if  $\gamma_i = \gamma_j$ . For example,  $\{\{\theta_1\}, \{\theta_2, \theta_3\}\}$  could be represented by  $(1, 2, 2)$  but also by  $(3, 1, 1)$ . We first explain the remainder of the sampling scheme and motivate the duplicate representations in the next paragraph. The number of possible duplicate representations in  $\vec{\gamma}$  for one partition is straightforward to compute, and the prior over  $\vec{\gamma}$  is obtained by taking the prior over the partitions and dividing uniformly over duplicate representations. Next, we sample each element of  $\vec{\gamma}$  conditional on the other elements. Since the partition membership is discrete, we enumerate all possible values and draw from the resulting categorical distribution. Sampling individual elements of  $\vec{\gamma}$  from the conditional distributions rather than the joint distribution reduces the complexity from  $\mathcal{O}(B_K)$  to  $\mathcal{O}(K^2)$ .

Although the duplicate representations of  $\vec{\gamma}$  for one partition introduce some additional computational cost, they facilitate exploration of the posterior space. For example, if we had used a one-to-one mapping from partitions to  $\vec{\gamma}$ , then updating the first membership in  $(1, 2, 2)$  to  $(2, 2, 2)$  would not be a valid configuration, as this should be represented by  $(1, 1, 1)$ . However, a transition from  $(1, 2, 2)$  to  $(1, 1, 1)$  requires updating two parameters and is therefore less likely to occur. Nevertheless, on the level of partitions, it makes sense to propose a move from  $\{\{\theta_1\}, \{\theta_2, \theta_3\}\}$  to  $\{\{\theta_1, \theta_2, \theta_3\}\}$ .

## 4 Investigating Multiplicity Adjustment

In this section, we investigate the differences between the above priors in more detail and compare them to the method proposed by Westfall et al. (1997) and an uncorrected approach using pairwise Bayes factors. In Section 4.1, we use a small simulation study to illustrate the implications of multiplicity adjustment. In Section 4.2, we present the results of a more extensive simulation study.

### 4.1 Illustrating Multiplicity Adjustment

Here we illustrate the different multiplicity penalties that the different priors impose using a small simulation study. We simulate data from a one-way ANOVA model and analyze it using the specification by Rouder et al. (2012). The ANOVA model extended with a prior over

partitions is given by:

$$\begin{aligned}
Y_{ij} &\sim \mathcal{N}(\mu + \sigma\theta_j, 1) \\
\mu &\propto 1 \\
\sigma^2 &\propto 1/\sigma^2 \\
g &\sim \mathcal{IG}(1/2, 1/2) \\
\vec{\theta}^u &\sim \mathcal{N}_{K-1}(0, g) \\
\vec{\theta}^c &\leftarrow \mathbf{Q}\vec{\theta}^u \\
\theta_j &\leftarrow \text{mean of elements of } \vec{\theta}^c \text{ in the same partition} \\
\rho &\sim \pi_\rho(\cdot) \quad .
\end{aligned} \tag{16}$$

The data follow a Gaussian distribution with a grand mean  $\mu$  and a group-specific offset  $\theta_j$ . The offsets sum to zero to avoid identification constraints. This is achieved by projecting  $\vec{\theta}^u$  from a  $K - 1$  dimensional space onto a  $K$  dimensional space using the matrix  $\mathbf{Q}$ , which consists of the first  $K - 1$  columns of an eigendecomposition of a degenerate covariance matrix as defined in Rouder et al. (2012).<sup>1</sup> Next, the elements of  $\vec{\theta}^c$  within the same partition are averaged to obtain  $\theta_j$ . The unconstrained offsets  $\vec{\theta}^u$  are assigned a  $g$  prior where  $g$  itself is assigned an inverse gamma prior with shape and scale equal to  $1/2$  (Liang et al., 2008). Note that the model reduces to the approach of Rouder et al. (2012) whenever the partition indicates that all elements are distinct.

We simulated from the null model, which assumes that all the groups are equal, and from the full model, which assumes that all groups are unequal, drawing 100 observations per group and varying the number of groups  $K \in [2, 3, \dots, 10]$ , repeating each combination 100 times. In the full model, the means were of increasing size with successive differences of 0.20. For the analysis we considered six priors: the Dirichlet process prior with  $\alpha \in \{0.50, 1\}$  and  $\alpha$  set adaptively to have equal prior mass assigned to the model with all equalities and the model with all inequalities (i.e.,  $p(\mathcal{H}_0) = p(\mathcal{H}_1)$ ), as done by Gopalan and Berry (1998); the beta-binomial prior with  $\alpha = 1$  and  $\beta \in \{1, K, \binom{K}{2}\}$ ; and the uniform prior. We also included the prior adjustment method proposed by Westfall et al. (1997) and an uncorrected method using pairwise Bayes factors. We used our methodology as described in Section 3.5, drawing 12,000 MCMC samples and discarding the first 2,000 as a burn-in.

To assess how well the respective priors adjust for multiplicity, we calculated how frequently the posterior probability that any two groups differ is larger than 0.50, using the null model as data-generating model. Similarly, to assess how well the respective priors are capable of detecting true differences, we calculated how frequently the posterior probability that any two groups *do not* differ is larger than 0.50, using the full model as data-generating model.

The left panel in Figure 3 shows that using a uniform prior (blue squares) very quickly leads to false positives as the number of groups increases. This is not surprising: the uniform prior assigns each model the same prior mass, hence diminishing the plausibility assigned to  $\mathcal{H}_0$  dramatically as  $K$  increases, thus increasing the probability of an error. The Dirichlet process prior which assigns equal mass to the full and the null model (yellow suns), as suggested by Gopalan and Berry (1998), performs better than the uniform prior but still does not provide adequate error control. It performs roughly as poorly as the method which simply computes pairwise Bayesian  $t$ -tests (green circles). The correction proposed by Westfall et al. (1997) performs much better (light blue circles) but still leads to a relatively high probability of making at least one error as the number of groups increases. The DP prior with  $\alpha = 1$  (pink stars) performs better, with the DP prior with  $\alpha = 0.50$  (beige diamonds) and the set of beta-binomial priors providing good error control.

---

<sup>1</sup>Note that this projection is not unique. It can also be achieved with, for example, a QR decomposition, as recommended by the Stan Development Team (2022).

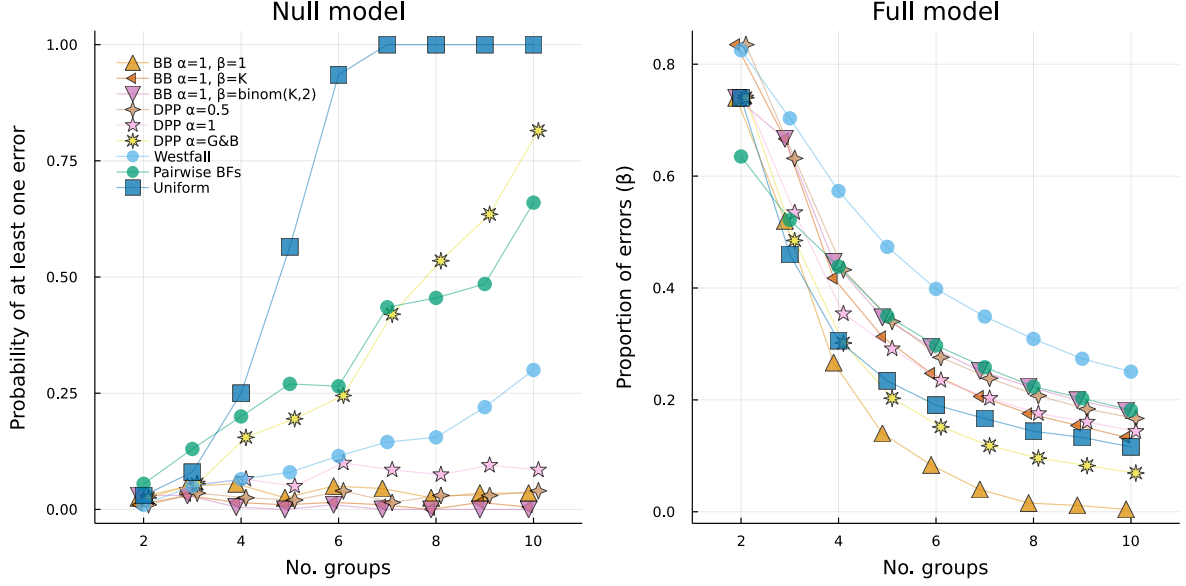


Figure 3: Left: Probability of making at least one false claim about a difference between two groups when there is none. Right: Proportion of falsely claiming no difference between two groups when there is one.

The right panel in Figure 3 shows that the beta-binomial prior with  $\alpha = \beta = 1$  leads to the lowest proportion of falsely claiming no difference between two groups, followed by the Dirichlet process prior for which  $p(\mathcal{H}_0) = p(\mathcal{H}_1)$  and the uniform prior. The method proposed by Westfall et al. (1997) performs worst, followed by the beta-binomial prior with  $\alpha = 1$  and  $\beta = \binom{K}{2}$  and the DP prior with  $\alpha = 0.50$ . The performance of the uncorrected pairwise Bayes factor approach is somewhere in the middle. Note that all approaches perform better as the group size increases, but this is due to our simulation design: each additional group exhibits a mean larger than the previous one by 0.20 and adds  $n$  more observations, which makes falsely claiming no difference less likely with an increasing number of groups. Instead of looking at absolute error, we therefore focus on the relative ordering of the priors. Overall, we conclude that not adjusting for multiple comparisons — either by using a uniform prior or by using pairwise Bayes factors — naturally leads to the worst performance and that the method by Westfall et al. (1997) is overly conservative and does not provide adequate error control with an increasing number of groups. In the next section, we report on a more extensive simulation study to further disentangle the differences between the multiple comparison methods.

## 4.2 Simulation Study

In the previous section, we illustrated the importance of adjusting the prior model probabilities in reducing the familywise error rate when all groups are equal. Here we explore the multiplicity adjustment of the different methods in a more exhaustive simulation study. We used the same ANOVA model as in the previous section and varied the total number of groups  $K \in \{5, 9\}$  and the sample size per group  $n \in \{50, 100, 250, 500\}$ . In addition, we varied the true number of equalities to be  $\{0\%, 25\%, 50\%, 75\%, 100\%\}$ . For  $K = 5$ , there are 4 possible (in)equalities which resulted in models that have either 0, 1, 2, 3, or 4 equalities. For  $K = 9$ , there are 8 possible (in)equalities, resulting in 0, 2, 4, 6, or 8 equalities in the true model. Given the number of equalities, we sampled a particular partition uniformly from all possible partitions with that amount of equalities and used this model to simulate data. Each unique combination was repeated 100 times and each generated data set was analyzed with the same prior specifications as above. We assessed the familywise error control as well as statistical power. The results for

$K = 5$  and  $K = 9$  were similar. Therefore, we focus on the  $K = 5$  in the main text and discuss the  $K = 9$  case in Appendix B.

Note that the hierarchical approach has an additional source of  $\alpha$  error in contrast to pairwise comparisons when there are more than 0 inequalities because it imposes transitivity. For example, imagine that the true model postulates that  $\theta_1 = \theta_2 = \theta_3 \neq \theta_4$ . However, the sample means are (by random sampling)  $\bar{x}_1 = 0.1, \bar{x}_2 = 0.2, \bar{x}_3 = 0.3, \bar{x}_4 = 0.35$ . The hierarchical approach would find that  $\theta_3 = \theta_4$ , but not that  $\theta_1 = \theta_3$  since that also implies  $\theta_1 = \theta_4$ . Therefore, the model  $\theta_1 = \theta_2 = \theta_3 \neq \theta_4$  and even the equality  $\theta_1 = \theta_2$  are not retrieved. In contrast, the pairwise methods violate transitivity as they only look at two pairs at the time and will happily suggest that  $\theta_1 = \theta_2$ ,  $\theta_2 = \theta_3$ , and  $\theta_3 = \theta_4$  while simultaneously suggesting that  $\theta_1 \neq \theta_4$ .

#### 4.2.1 Familywise Error Rate

Figure 4 shows the probability of at least one error for different methods across the number of *inequalities* in the true model and sample sizes. The top left panel shows that the uniform prior (blue squares), the pairwise Bayes factors (green circles), the Dirichlet process prior with  $(p(\mathcal{H}_0) = p(\mathcal{H}_1))$  (yellow stars), and the method proposed by Westfall et al. (1997) (light blue circles) perform worst and that the other Dirichlet process and beta-binomial priors provide adequate error control. This mirrors the results above, which is natural since this part of the simulation is a special case for  $K = 5$ . Increasing the number of inequalities to 1 (top right) and 2 (bottom left), we find that the pairwise Bayes factors, the method by Westfall et al. (1997), and the uniform improve in performance. This is likely due to the fact that, with more inequalities, there are simply less opportunities to incorrectly claim that two population means are different. In contrast, the performance of the other methods decreases when there is at least one inequality; it is difficult to disentangle a trend with increasing inequalities.

The rightmost panel in Figure 4 shows the results averaged over the number of inequalities in the true model. We find that the method by Westfall et al. (1997) shows the strongest familywise error control, closely followed by the the beta-binomial priors with  $\beta = K$  and  $\beta = \binom{K}{2}$  and the DP prior with  $\alpha = 0.50$ . The pairwise Bayes factors perform similar to the Dirichlet process prior with  $\alpha = 1$ , with the beta-binomial prior with  $\beta = 1$ , the symmetric DP prior, and the uniform prior performing worst. The differences between the methods become less pronounced with increasing sample size since the data starts to dominate the prior.

#### 4.2.2 Statistical Power

Figure 5 shows the proportion of falsely claiming a difference between two groups when there is none for different methods across the number of *equalities* in the true model and sample sizes. The top left panel shows that the beta-binomial prior with  $\beta = 1$  performs best and the method proposed by Westfall et al. (1997) performs worst, again mirroring the results of the small simulation study above. Increasing the number of equalities in the true model, we find that the performance of virtually all methods decreases except for the uniform prior, which shows a slight increase, especially for large sample sizes. This overall decrease in performance is likely due to the fact that the average pairwise difference between groups *decreases* with the number of equalities. To illustrate, note that the model with no equalities for  $K = 4$  groups has population means  $\vec{\mu} = \{-0.30, -0.10, 0.10, 0.30\}$ , which yields pairwise differences  $[0.20, 0.20, 0.20, 0.40, 0.40, 0.60]$  with an average of 0.33. In contrast, including one equality results in  $\vec{\mu} = \{-0.25, -0.05, 0.15, 0.15\}$ <sup>2</sup>, yielding pairwise differences  $[0.20, 0.20, 0.20, 0.40, 0.40]$  with an average of 0.28.

The rightmost panel in Figure 5 shows the results averaged over the number of equalities in the true model. We find that the method by Westfall et al. (1997) is highly conservative, trading

<sup>2</sup>This is due to the sum-to-zero constraint and the constraint that all successive unequal groups have a difference of 0.20.

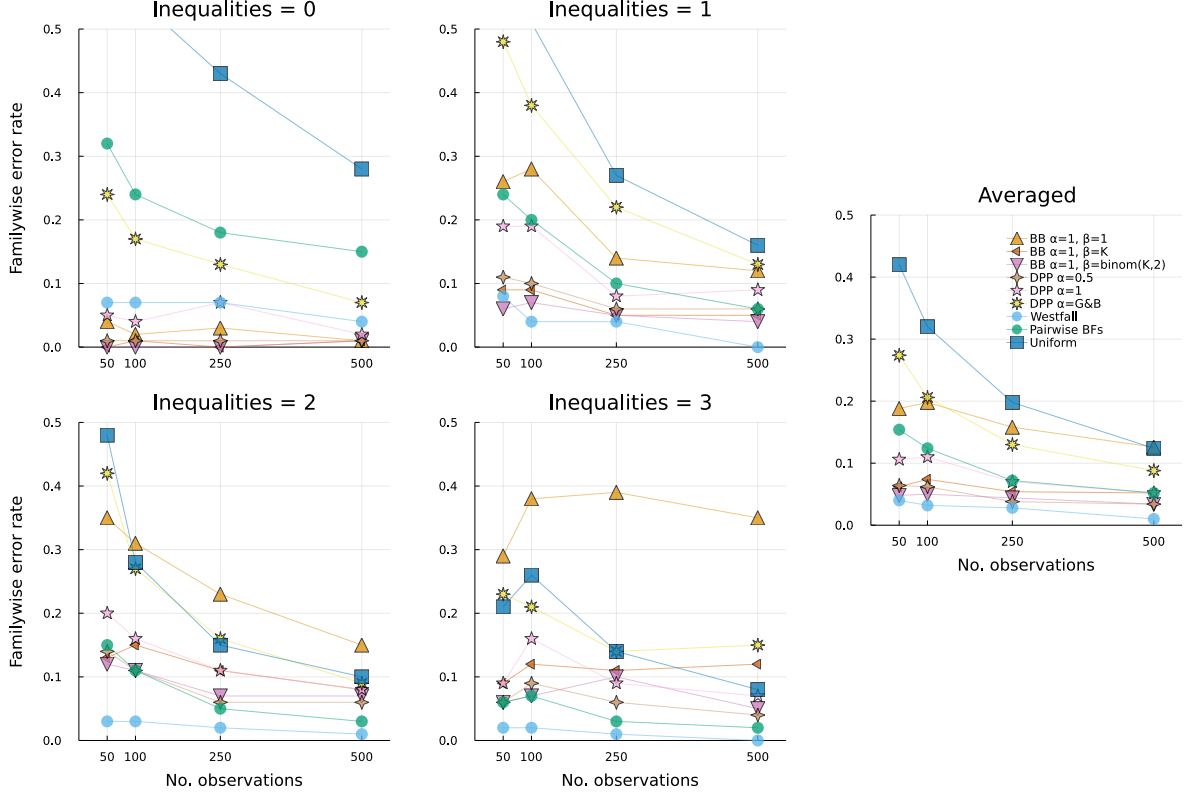


Figure 4: Familywise error rate across priors and sample sizes under a model with 0 (top left), 1 (top right), 2 (bottom left), and 3 (bottom right) true inequalities for  $K = 5$  groups. The rightmost panel shows the average familywise error rate across inequalities.

off the strong familywise error control with an increase in the proportion of false negatives. Similarly, the priors that performed worst with respect to familywise error control — the uniform, symmetric DP, and beta-binomial prior with  $\beta = 1$  — perform best here. The other DP and beta-binomial priors as well as the pairwise Bayes factors are somewhere in between those two extremes. Note that again the differences between the methods become less pronounced with increasing sample size.

#### 4.2.3 Simulation Discussion

Our results show that no single method dominates all others. While the beta-binomial prior with  $\beta = 1$  performed best in our initial simulation study described in Section 4.1, including models beyond the null and full model showed that this prior performed considerably worse in those settings. The beta-binomial prior with  $\beta = K$ ,  $\beta = \binom{K}{2}$ , and the DP prior with  $\alpha = 0.50$  perform very similarly overall. Importantly, both the method proposed by Westfall et al. (1997) and the pairwise Bayes factors can yield transitivity violations, while explicitly specifying a prior over partitions cannot. For example, we might find that  $\mu_1 = \mu_2$  and  $\mu_2 = \mu_3$  using pairwise Bayes factors with some threshold, but at the same time conclude that  $\mu_1 \neq \mu_3$ . This is one key reason why explicitly specifying the prior over partitions is preferable. In the next section, we focus on the beta-binomial prior with  $\beta = K$  and apply our method to two examples.

## 5 Applications

In this section, we apply the beta-binomial setup to two examples: testing the (in)equality of proportions and variances, respectively. We have developed a generic Julia package called

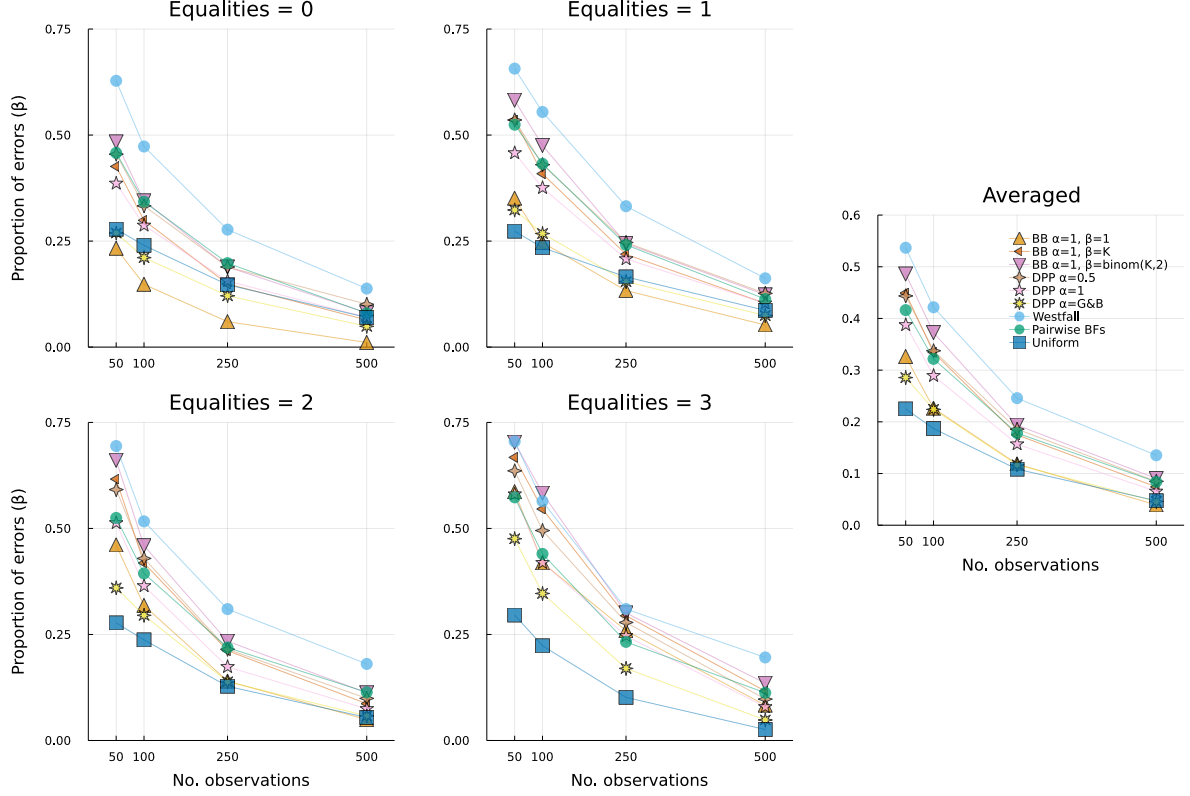


Figure 5: Proportion of falsely claiming a difference between two groups when there is none across priors and sample sizes under a model with 0 (top left), 1 (top right), 2 (bottom left), and 3 (bottom right) true inequalities for  $K = 5$  groups. The rightmost panel shows the average error rate across inequalities.

*EqualitySelection* that utilizes the probabilistic programming framework *Turing* to allow the user to adjust for multiplicity as proposed in this paper.

## 5.1 Testing Proportions

Nuijten et al. (2016) investigated a sample of 30,717 articles published between 1985 and 2013 in eight major psychology journals for statistical reporting errors. Our question here is: Which journals make the same amount of errors, and which make more errors? We answer the question using the following model specification. For journal  $j$ , denote the number of statistical errors found as  $e_j$  and the number of statistical tests analyzed as  $n_j$ . We assume that underlying each proportion there is a latent true chance of making an error,  $\theta_j$ . Thus, we modeled the data as independent binomials, that is,  $e_j \sim \text{Binomial}(\theta_j, n_j)$ . Next, we specify a hierarchical level over the partitions to assess for which journals the chances of making an error are equal. This leads to the following model specification:

$$\begin{aligned}
 e_j &\sim \text{Binomial}(\theta_j, n_j) \\
 \theta_j^u &\sim \text{Beta}(1, 1) \\
 \theta_j &\leftarrow \text{mean of elements of } \theta_j^u \text{ in the same partition} \\
 \rho &\sim \text{beta-binomial}(1, 8) .
 \end{aligned} \tag{17}$$

The unconstrained chances  $\theta_j^u$  are assigned beta priors from which — together with the partitions — the possibly constrained chances are created. Two chances  $\theta_i$  and  $\theta_j$  are equal if and only if their indices appear in the same partition  $\{i, j\} \subseteq \rho_k$  for some  $k$ . Note that the

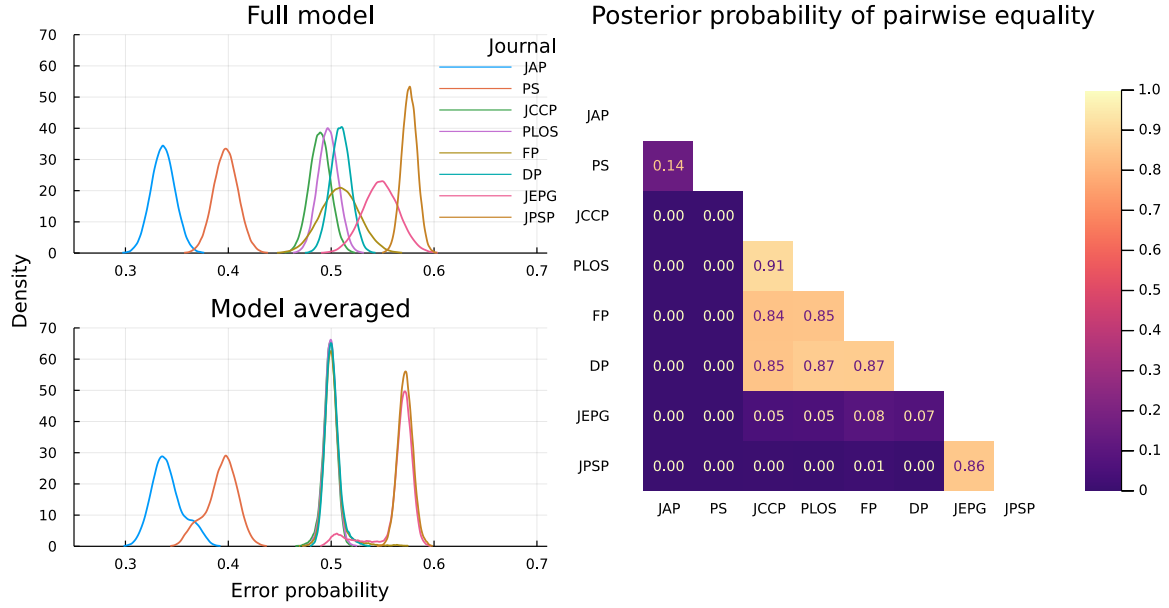


Figure 6: Left: Posterior means of the full model where all proportions are assumed to be different (top) and posterior means when averaging over all models using a beta-binomial( $\alpha = 1$ ,  $\beta = 8$ ) prior (bottom). Right: Posterior probabilities for pairwise equality across all journals. The abbreviations stand for: *Journal of Applied Psychology* (JAP), *Psychological Science* (PS), *Journal of Consulting and Clinical Psychology* (JCCP), *Public Library of Science* (PLOS), *Developmental Psychology* (DP), *Journal of Experimental Psychology: General* (JEPG), and *Journal of Personality and Social Psychology* (JPSP).

model reduces to the full model of independent binomials whenever the partitions state that all elements in  $\vec{\theta}$  are distinct. We use a beta-binomial prior with  $\alpha = 1$  and  $\beta = 8$ . The top left panel in Figure 6 shows the posterior distributions for the underlying error chance for each journal under a model that assumes that they are all different.

We can see that the posterior distributions for JCCP (green), PLOS (purple), DP (turquoise), and FP (beige) are very close to each other, with FP showing more pronounced uncertainty. The panel below shows the model-averaged posterior distributions, clearly demonstrating a shrinkage effect. The error chances for JAP and PS are pulled toward each other, with JCCP, PLOS, DP, and FP being shrunk towards each other almost completely, similarly to JEPG and JPSP. The right panel in Figure 6 gives the posterior distributions for pairwise equality across all journals, reflecting the two main clusters in the model-averaged density plot on the left.

## 5.2 Testing Standard Deviations

Borkenau et al. (2013) studied whether men and women differ in the variability of personality traits. Here we focus on five personality traits (agreeableness, extraversion, openness, conscientiousness, neuroticism) rated by participants' peers in an Estonian sample consisting of  $n_1 = 969$  women and  $n_2 = 716$  men. Our goal is to assess which personality traits across the sexes can be assumed equal in terms of their variability. This example shows how our methodology can be used to test group differences while taking the multivariate dependency of the outcome measure into account. We build on the parameterization proposed by Dablander et al. (2020), who developed a default Bayes factor test for testing the (in)equality of variances. Let  $\vec{y}_1$  and  $\vec{y}_2$  denote the five-element vectors of observed data for men and women, respectively, and  $K = 10$

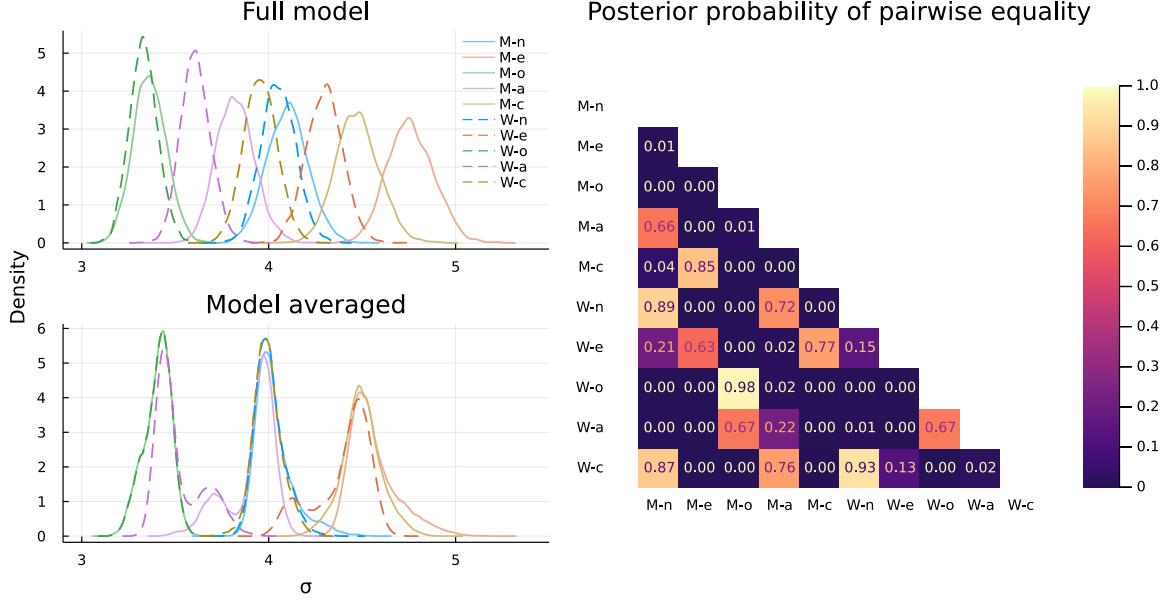


Figure 7: Left: Posterior means of the full model where all standard deviations are assumed to be different (top) and posterior means when averaging across all models using a beta-binomial( $\alpha = 1$ ,  $\beta = 10$ ) prior (bottom). Right: Posterior probabilities for pairwise equality across all personality traits. In the abbreviations the first letter stands for *men* (m) or *women* (w). The second letter stands for *neuroticism* (n), *extraversion* (e), *openness* (o), *agreeableness* (a), and *conscientiousness* (c).

be the total number of variables. For each sex  $k \in \{1, 2\}$ , we have:

$$\begin{aligned}
\vec{Y}_k &\sim \mathcal{N}(\vec{\mu}_k, \Sigma_k) \\
\vec{\mu}_k &\propto \vec{1} \\
\Sigma_k &= \text{diag}(\vec{\sigma}_k) \Omega_k \text{diag}(\vec{\sigma}_k) \\
\Omega_k &\sim \text{LKJ}(1) ,
\end{aligned}$$

where LKJ refers to the Lewandowski-Kurowicka-Joe prior (Lewandowski et al., 2009). To test the equality of variances both between and across groups, we define the ten-variable standard deviation vector  $\vec{\sigma} = [\vec{\sigma}_1, \vec{\sigma}_2]$  with  $\bar{\sigma}$  denoting the average standard deviation. Following Dablander et al. (2020), we write  $\sigma_j = (K\vartheta_j\bar{\sigma})^{-1}$ , where  $\vartheta_j = \sigma_j/\sum_{j=1}^K \sigma_j$  is the relative standard deviation and  $\vartheta_K = 1 - \sum_{j=1}^{K-1} \vartheta_j$ . To complete the model specification, we write:

$$\begin{aligned}
\sigma_j &= (K\vartheta_j\bar{\sigma}_j)^{-1} \\
\bar{\sigma}_j &\propto \bar{\sigma}_j^{-1} \\
\vartheta_j &\leftarrow \text{mean of elements of } \vartheta^u \text{ in the same partition} \\
\vec{\vartheta}^u &\sim \text{Dirichlet}(1, \dots, 1) \\
\rho &\sim \text{beta-binomial}(1, 10) .
\end{aligned} \tag{18}$$

Two standard deviations  $\sigma_i$  and  $\sigma_j$  are equal if and only if their indices appear in the same partition  $\{i, j\} \subseteq \rho_k$  for some  $k$ . When the partition states that all standard deviations are distinct we recover the full model. The top left panel of Figure 7 shows the posterior distributions under the full model that assumes all standard deviations are different.

While all posterior distributions lie close to each other, the standard deviations of openness for men and women overlap particularly much. The bottom panel shows the model-averaged

posterior distributions, which again demonstrate a shrinkage effect. The right panel of Figure 7 shows the posterior probability of pairwise equality across all personality traits for men and women. It appears that there are three clusters: (1) men–openness, women–openness, and women–agreeableness; (2) men–neuroticism, women–neuroticism, women–conscientiousness, and men–agreeableness; (3) men–conscientiousness, men–extraversion, and women–extraversion. However, for the personality traits women–agreeableness, men–agreeableness, and women–extraversion, the evidence is not overwhelming, as indicated by the bimodality in the model-averaged posterior distributions.

## 6 Discussion

Testing the (in)equality between groups while adjusting for multiple comparisons is a core challenge in many applied settings. In this paper, we have proposed a flexible class of beta-binomial priors to penalize multiplicity and make inferences over all possible (in)equalities in relatively general settings. We compared the beta-binomial priors to a Dirichlet process prior suggested by Gopalan and Berry (1998), to a uniform prior, to the method proposed by Westfall et al. (1997), and to an uncorrected method based on pairwise Bayes factors. We also illustrated our method, which is freely available in the Julia package *EqualitySelection*, on two examples.

We found that a beta-binomial prior with  $\alpha = 1$  and  $\beta \in \{K, \binom{K}{2}\}$  as well as a Dirichlet process prior with  $\alpha < 1$  adequately control the familywise error rate, while a uniform prior and using only pairwise Bayes factors, unsurprisingly, do not. We also found that the method proposed by Westfall et al. (1997) compares favorably in terms of error control but not in terms of power. While we have focused on a posterior probability threshold of 0.50 (i.e., a Bayes factor of 1), other thresholds will naturally impact the trade-off between the two types of errors. Importantly, and in contrast to conventional adjustments for multiple comparisons (e.g., Jeffreys, 1961; Westfall et al., 1997), specifying a prior over the partitions allows inferences over all possible (in)equalities. This means that researchers can use the methods we provide to assess not only the probability of pairwise (in)equalities — as is common in standard post-hoc tests for, say, ANOVA — but in fact can make probabilistic statements over any set of (in)equalities they wish to assess. Similarly, the outlined approach also allows for model-averaging, which as we have seen in the applications yields shrinkage of the groups towards each other. Using a prior over partitions further avoids violations of transitivity, i.e. claiming for example that  $\mu_1 \neq \mu_3$  while both  $\mu_1 = \mu_2$  and  $\mu_2 = \mu_3$ .

As with any statistical method, there are a number of points to keep in mind. First, while we suggest default values of  $\alpha = 1$  and  $\beta = K$  for the beta-binomial prior and  $\alpha \leq 1$  for the DP prior, researchers may wish to use a more informed prior specification. Values for the prior parameters can be elicited by specifying model priors for two out of the following: the prior on the null model, on the full model, or their ratio. Second, the beta-binomial prior differs from the DP prior in that it assigns models with the same number of partitions the same prior probability, while the DP prior assigns more mass to the model with the larger cluster. It is not obvious which of the two behaviors is more desirable, and it may well depend on the problem under study. Researchers using the methods we have made available should keep this difference in mind, although the extent to which it matters in practice remains to be seen.

There are some practical limitations of our implementation that we leave for future work. We currently do not allow for factorial designs, for example, for which dummy or contrast coding is more natural. The key challenge there is to specify the prior in such a way that it reflects the structure of the experimental design. For the present, we believe that the Bayesian approach outlined in this paper can help applied researchers who wish to compare multiple groups.

**Author Contributions.** DvdB and FD proposed the study. DvdB implemented the method and created the Julia package. DvdB conducted the simulation study and analyzed the data

with the help of FD. FD and DvdB wrote the manuscript. All authors read and approved the submitted version of the paper. They also declare that there were no conflicts of interest.

## References

- Bayarri, M. J., Berger, J. O., Forte, A., & García-Donato, G. (2012). Criteria for Bayesian model choice with application to variable selection. *The Annals of Statistics*, 40(3), 1550–1577.
- Benjamini, Y., & Braun, H. (2002). John W. Tukey’s contributions to multiple comparisons. *The Annals of Statistics*, 30(6), 1576–1594.
- Berg, S. (1975). Some properties and applications of a ratio of Stirling numbers of the second kind. *Scandinavian Journal of Statistics*, 91–94.
- Berry, D. A., & Hochberg, Y. (1999). Bayesian perspectives on multiple comparisons. *Journal of Statistical Planning and Inference*, 82(1-2), 215–227.
- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1), 65–98.
- Blackwell, D., & MacQueen, J. B. (1973). Ferguson distributions via Pólya urn schemes. *The Annals of Statistics*, 1(2), 353–355.
- Borkenau, P., Hřebíčková, M., Kuppens, P., Realo, A., & Allik, J. (2013). Sex differences in variability in personality: A study in four samples. *Journal of Personality*, 81(1), 49–60.
- Broder, A. Z. (1984). The  $r$ -Stirling numbers. *Discrete Mathematics*, 49(3), 241–259.
- Casella, G., Girón, F. J., Martínez, M. L., Moreno, E., et al. (2009). Consistency of Bayesian procedures for variable selection. *The Annals of Statistics*, 37(3), 1207–1228.
- Consonni, G., Fouskakis, D., Liseo, B., Ntzoufras, I., et al. (2018). Prior distributions for objective Bayesian analysis. *Bayesian Analysis*, 13(2), 627–679.
- Dablander, F., van den Bergh, D., Ly, A., & Wagenmakers, E.-J. (2020). Default Bayes Factors for Testing the (In) equality of Several Population Variances. *arXiv preprint arXiv:2003.06278*.
- de Jong, T. (2019). A Bayesian approach to the correction for multiplicity. <https://doi.org/10.31234/osf.io/s56mk>
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, 209–230.
- Ge, H., Xu, K., & Ghahramani, Z. (2018). Turing: A language for flexible probabilistic inference. *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, 1682–1690.
- George, E. I., & McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423), 881–889.
- Gopalan, R., & Berry, D. A. (1998). Bayesian multiple comparisons using Dirichlet process priors. *Journal of the American Statistical Association*, 93(443), 1130–1139.
- Hinne, M., Gronau, Q. F., van den Bergh, D., & Wagenmakers, E.-J. (2020). A conceptual introduction to Bayesian model averaging. *Advances in Methods and Practices in Psychological Science*, 3(2), 200–215.
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model averaging: A tutorial. *Statistical Science*, 382–401.
- Ishwaran, H., & James, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453), 161–173.
- Jeffreys, H. (1961). *Theory of Probability (3rd Ed.)* Oxford, UK: Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9), 1989–2001.

- Liang, F., Paulo, R., Molina, G., Clyde, M. A., & Berger, J. O. (2008). Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, *103*(481), 410–423.
- Ly, A., Verhagen, J., & Wagenmakers, E.-J. (2016). Harold Jeffreys’s default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, *72*, 19–32.
- Mezo, I. (2011). The  $r$ -Bell numbers. *Journal of Integer Sequences*, *14*(1), 1–14.
- Moreno, E., Girón, J., Casella, G., et al. (2015). Posterior model consistency in variable selection as the model dimension grows. *Statistical Science*, *30*(2), 228–241.
- Nuijten, M. B., Hartgerink, C. H., Van Assen, M. A., Epskamp, S., & Wicherts, J. M. (2016). The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior Research Methods*, *48*(4), 1205–1226.
- Quintana, F. A. (2006). A predictive view of Bayesian clustering. *Journal of Statistical Planning and Inference*, *136*(8), 2407–2429.
- Rao, C. (2009). Multiple comparison procedures-a note and a bibliography. *Journal of Statistics*, *16*(1), 66–109.
- Rasmussen, C. E., et al. (1999). The infinite Gaussian mixture model. *NIPS*, *12*, 554–560.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, *56*(5), 356–374.
- Scott, J. G., & Berger, J. O. (2006). An exploration of aspects of Bayesian multiple testing. *Journal of Statistical Planning and Inference*, *136*(7), 2144–2162.
- Scott, J. G., & Berger, J. O. (2010). Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*, 2587–2619.
- Stan Development Team. (2022). *Stan modeling language users guide and reference manual*, 2–30. <https://mc-stan.org>
- Teh, Y. W. (2010). Dirichlet Process.
- Westfall, P. H., Johnson, W. O., & Utts, J. M. (1997). A Bayesian perspective on the Bonferroni adjustment. *Biometrika*, *84*(2), 419–427.
- Wilson, M. A., Iversen, E. S., Clyde, M. A., Schmidler, S. C., & Schildkraut, J. M. (2010). Bayesian model search and multilevel inference for SNP association studies. *The Annals of Applied Statistics*, *4*(3), 1342–1364.

## A Beta-binomial Prior with Decreasing Prior Model Odds

**Proposition 1.** The prior density of the beta-binomial distribution over partitions is decreasing for  $\alpha = 1$  and  $\beta \geq \binom{K}{2}$ , and strictly decreasing for  $\alpha = 1$  and  $\beta > \binom{K}{2}$ .

*Proof.* The prior density of the Beta-binomial over partitions is given by:

$$\pi(\rho \mid K, \alpha, \beta) = \binom{K-1}{|\rho|-1} \frac{B(|\rho|-1+\alpha, K-|\rho|+\beta)}{B(\alpha, \beta) \left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}}.$$

To examine the ratio of two consecutive model sizes we evaluate the ratio of the prior for partitions  $\rho$  and  $q$  with  $|q| = |\rho| + 1$ :

$$\frac{\pi(\rho \mid K, \alpha, \beta)}{\pi(q \mid K, \alpha, \beta)} = \frac{\binom{K-1}{|\rho|-1}}{\binom{K-1}{|\rho|}} \frac{B(|\rho|-1+\alpha, K-|\rho|+\beta)}{B(|\rho|+\alpha, K-|\rho|-1+\beta)} \frac{\left\{ \begin{smallmatrix} K \\ |\rho|+1 \end{smallmatrix} \right\}}{\left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}}, \quad (19)$$

$$= \frac{|\rho|}{K-|\rho|} \frac{\beta+K-|\rho|-1}{\alpha+|\rho|-1} \frac{\left\{ \begin{smallmatrix} K \\ |\rho|+1 \end{smallmatrix} \right\}}{\left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}}. \quad (20)$$

Using the recurrence relation of the Stirling numbers  $\left\{ \begin{smallmatrix} n+1 \\ k \end{smallmatrix} \right\} = k \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} + \left\{ \begin{smallmatrix} n \\ k-1 \end{smallmatrix} \right\}$ , the ratio  $\left\{ \begin{smallmatrix} K \\ |\rho|+1 \end{smallmatrix} \right\} / \left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}$  is equivalent to  $\left\{ \begin{smallmatrix} K+1 \\ |\rho|+1 \end{smallmatrix} \right\} / \left\{ \begin{smallmatrix} K \\ |\rho|+1 \end{smallmatrix} \right\} - (|\rho| + 1)$ . This ratio of Stirling numbers was studied by Berg (1975) and their property 2 provides this inequality

$$\frac{\left\{ \begin{smallmatrix} K+1 \\ |\rho|+1 \end{smallmatrix} \right\}}{\left\{ \begin{smallmatrix} K \\ |\rho|+1 \end{smallmatrix} \right\}} - |\rho| - 1 \geq \frac{\left\{ \begin{smallmatrix} K+1 \\ |\rho| \end{smallmatrix} \right\}}{\left\{ \begin{smallmatrix} K \\ |\rho| \end{smallmatrix} \right\}} - |\rho|.$$

It follows that the ratio in Equation (19) is maximal at  $|\rho| = K - 1$  and has value  $\binom{K}{2}$ . Next, we fix  $\alpha = 1$  and solve  $\pi(K \mid K-1, 1, \beta) / \pi(K \mid K, 1, \beta) = 1$  for  $\beta$  which yields  $\beta = \binom{K}{2}$ . Thus  $\beta \geq \binom{K}{2}$  implies  $\pi(j+1 \mid K, 1, \beta) \geq \pi(j \mid K, 1, \beta)$  (resp.  $\beta > \binom{K}{2}$  implies  $\pi(j+1 \mid K, 1, \beta) > \pi(j \mid K, 1, \beta)$ ).  $\square$

## B Simulation Results for $K = 9$

Here we present the extended simulation results for the  $K = 9$  group case. Figure 8 mirrors the results for the  $K = 9$  case, namely that the pairwise Bayes factors, the method proposed by Westfall et al. (1997), and the uniform prior generally increase in performance as the number of inequalities increase, while the other priors generally decrease in performance. Averaging over the settings, we again find that the beta-binomial prior with  $\beta = 1$ , the uniform prior, and the symmetric DP prior exhibit the worst error control, with the method proposed by Westfall et al. (1997) performing best, closely followed by the beta-binomial prior with  $\beta = \binom{K}{2}$  and the DP prior with  $\alpha = 0.50$ .

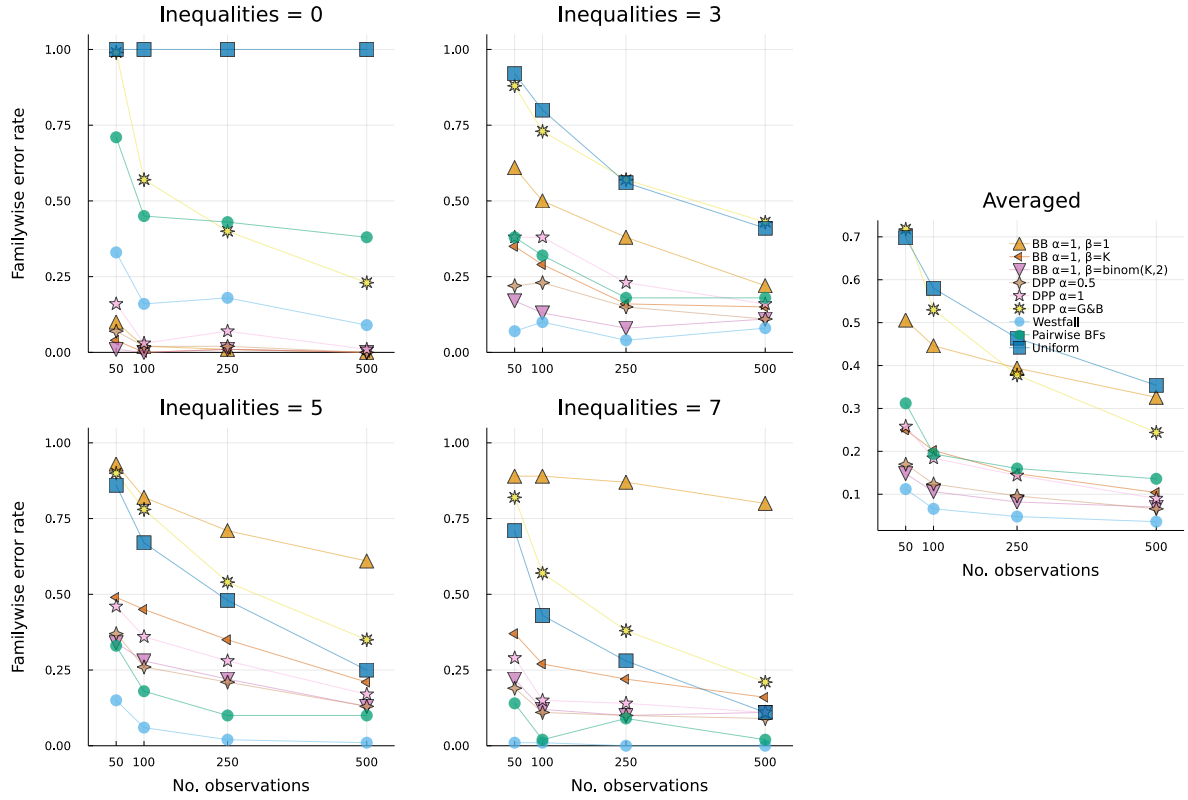


Figure 8: Familywise error rate across priors and sample sizes under a model with 0 (top left), 3 (top right), 5 (bottom left), and 7 (bottom right) true inequalities for  $K = 9$  groups. The rightmost panel shows the average familywise error rate across inequalities.

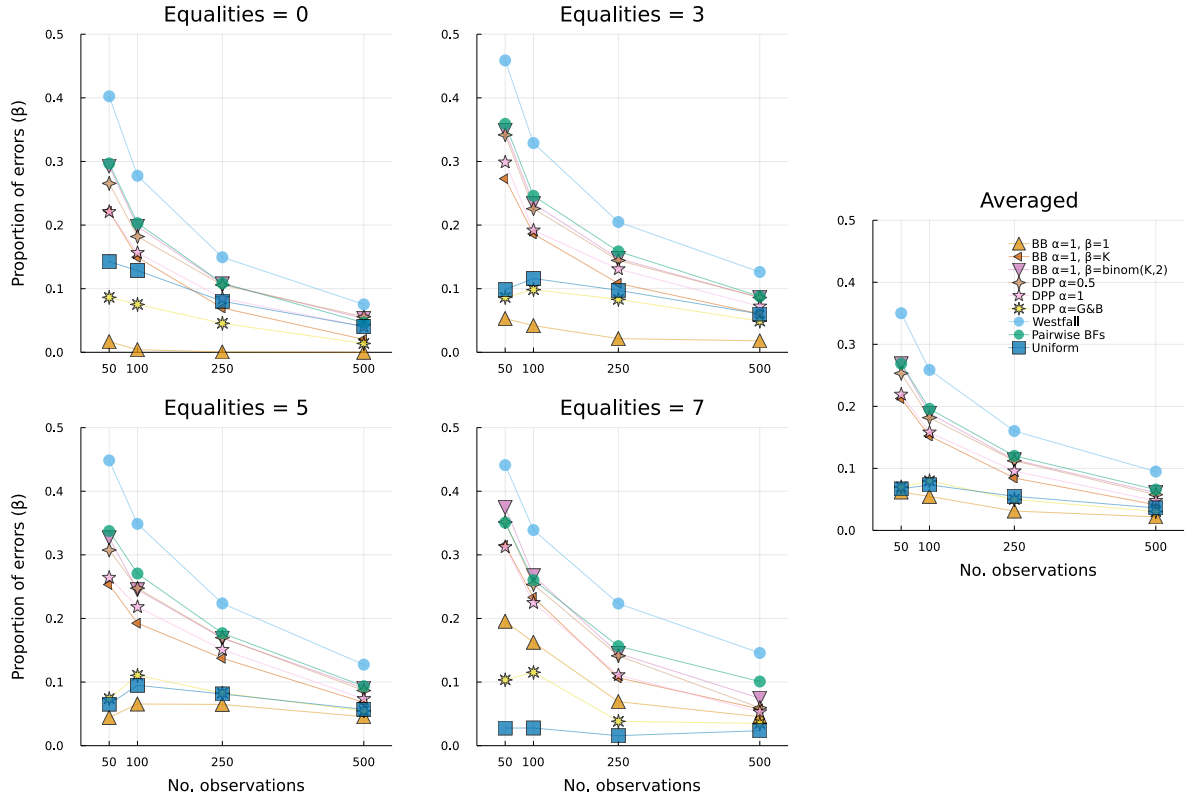


Figure 9: Proportion of falsely claiming a difference between two groups when there is none across priors and sample sizes under a model with 0 (top left), 3 (top right), 5 (bottom left), and 7 (bottom right) true inequalities for  $K = 9$  groups. The rightmost panel shows the average error rate across inequalities.