

Two-sample Bayesian Nonparametric Hypothesis Testing

Chris C. Holmes^{*}, François Caron[†], Jim E. Griffin[‡], and David A. Stephens[§]

Abstract. In this article we describe Bayesian nonparametric procedures for two-sample hypothesis testing. Namely, given two sets of samples $\mathbf{y}^{(1)} \stackrel{\text{iid}}{\sim} F^{(1)}$ and $\mathbf{y}^{(2)} \stackrel{\text{iid}}{\sim} F^{(2)}$, with $F^{(1)}, F^{(2)}$ unknown, we wish to evaluate the evidence for the null hypothesis $H_0 : F^{(1)} \equiv F^{(2)}$ versus the alternative $H_1 : F^{(1)} \neq F^{(2)}$. Our method is based upon a nonparametric Pólya tree prior centered either subjectively or using an empirical procedure. We show that the Pólya tree prior leads to an analytic expression for the marginal likelihood under the two hypotheses and hence an explicit measure of the probability of the null $\Pr(H_0|\{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}\})$.

Keywords: Bayesian nonparametrics, Pólya tree, hypothesis testing.

1 Introduction

Nonparametric hypothesis testing is an important branch of statistics with wide applicability. For example we often wish to evaluate the evidence for systematic differences between real-valued responses under two different treatments without specifying an underlying distribution for the data. That is, given two sets of samples $\mathbf{y}^{(1)} \stackrel{\text{iid}}{\sim} F^{(1)}$ and $\mathbf{y}^{(2)} \stackrel{\text{iid}}{\sim} F^{(2)}$, with $F^{(1)}, F^{(2)}$ unknown, we wish to evaluate the evidence for the competing hypotheses

$$H_0 : F^{(1)} \equiv F^{(2)} \text{ versus } H_1 : F^{(1)} \neq F^{(2)}.$$

In this article we describe a nonparametric Bayesian procedure for this scenario. Our Bayesian method quantifies the weight of evidence in favour of H_0 in terms of an explicit probability measure $\Pr(H_0|\mathbf{y}^{(1,2)})$, where $\mathbf{y}^{(1,2)}$ denotes the pooled data $\mathbf{y}^{(1,2)} = \{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}\}$. To perform the test we use a Pólya tree prior (Lavine, 1992; Mauldin et al., 1992; Lavine, 1994) centered on some distribution G where under H_0 we have $F^{(1,2)} = F^{(1)} = F^{(2)}$ and under H_1 , $F^{(1)} \neq F^{(2)}$ are modelled as independent draws from the Pólya tree prior. In this way we frame the test as a model comparison problem and evaluate the Bayes Factor for the two competing models. The Pólya tree is a well known nonparametric prior distribution for random probability measures F on Ω where Ω denotes the domain of Y (Ferguson, 1974).

Bayesian nonparametrics is a fast developing discipline, but while there has been considerable interest in nonparametric inference there has somewhat surprisingly been

^{*}Department of Statistics and Oxford-Man Institute, University of Oxford, England, cholmes@stats.ox.ac.uk

[†]Department of Statistics, University of Oxford, England francois.caron@stats.ox.ac.uk

[‡]School of Mathematics, Statistics and Actuarial Science, University of Kent, England, J.E.Griffin-28@kent.ac.uk

[§]Department of Mathematics and Statistics, McGill University, Canada, d.stephens@math.mcgill.ca

little written on nonparametric hypothesis testing. Bayesian parametric hypothesis testing where $F^{(1)}$ and $F^{(2)}$ are of known form is well developed in the Bayesian literature, see e.g. Bernardo and Smith (2000), and most nonparametric work has concentrated on testing a parametric model versus a nonparametric alternative (the Goodness of Fit problem). Initial work on the Goodness of Fit problem (Florens et al., 1996; Carota and Parmigiani, 1996) used a Dirichlet process prior for the alternative distribution and compared to a parametric model. In this case, the nonparametric distributions will be almost surely discrete, and the Bayes factor will include a penalty term for ties. The method can lead to misleading results if the data is absolutely continuous, and has motivated the development of methods using nonparametric priors that guarantee almost surely continuous distributions. Dirichlet process mixture models are one such class. The calculation of Bayes factors for Dirichlet process-based models is discussed by Basu and Chib (2003). Goodness of fit testing using mixtures of triangular distributions is considered by McVinish et al. (2009). An alternative form of prior, the Pólya tree, was considered by Berger and Guglielmi (2001). Simple conditions on the prior lead to absolutely continuous distributions. Berger and Guglielmi (2001) develop a default approach and consider its properties as a conditional frequentist method. Hanson (2006) discusses the use of Savage-Dickey density ratios to calculate Bayes factors in favour of the centering distribution (see also Branscum and Hanson (2008)). Consistency issues are discussed by Dass and Lee (2004), Rousseau (2007), Ghosal et al. (2008) and McVinish et al. (2009). There has been some work on testing the hypothesis that two distributions are the same; Dunson and Peddada (2008) consider hypothesis testing of stochastic ordering using restricted Dirichlet process mixtures, but their methods could be modified to allow two-sided hypotheses. They consider an interval null hypothesis and rely on Gibbs sampling for posterior computation. Pennell and Dunson (2008) develop a Mixture of Dependent Dirichlet Processes approach to testing changes in an ordered sequence of distributions using a tolerance measure. Bhattacharya and Dunson (2012) develop an approach for nonparametric Bayesian testing of differences between groups, with the data within each group constrained to lie on a compact metric space or Riemannian manifold.

Recently, following work presented here (originally posted on arXiv in Holmes et al. (2009)), Ma and Wong (2011) propose to allow the two random distributions under the alternative to randomly couple on different parts of the sample space, thereby achieving borrowing of information. Moreover, Chen and Hanson (2014) propose to use Lavine's (1992) partition for each $F^{(j)}$ centered at the normal distribution. Their approach enables generalization to more than two samples, but contrary to our approach requires a truncation level to be set. They also follow Berger and Guglielmi (2001) by choosing the parameter c that maximizes the Bayes factor in favor of the alternative.

The rest of the paper is as follows. In Section 2 we discuss the Pólya tree prior and derive the marginal probability distributions that result from such a prior. In Section 3 we describe our method and algorithm for calculating $\Pr(H_0|\mathbf{y}^{(1,2)})$ based on a subjective partition. In Section 4 we discuss an empirical Bayes procedure where the Pólya tree priors are centered on the empirical cdf of the joint data. Section 5 discusses the sensibility of the procedures to tuning parameters. Section 6 provides a discussion of related approaches and Section 7 concludes with a discussion of potential extensions.




Figure 1: Construction of a Pólya tree distribution. Each of the θ_{ϵ_m} is independently drawn from $\text{Beta}(\alpha_{\epsilon_m 0}, \alpha_{\epsilon_m 1})$. Adapted from Ferguson (1974).

2 Pólya tree priors

Pólya trees form a class of distributions for random probability measures F on some domain Ω (Lavine, 1992; Mauldin et al., 1992; Lavine, 1994). Consider a recursive dyadic (binary) partition of Ω into disjoint measurable sets. Denote the k th level of the partition $\{B_j^{(k)}, j = 0, \dots, 2^k - 1\}$, where $B_i^{(k)} \cap B_j^{(k)} = \emptyset$ for all $i \neq j$. The recursive partition is constructed such that $B_j^{(k)} \equiv B_{2j}^{(k+1)} \cup B_{2j+1}^{(k+1)}$ for $k = 1, 2, \dots, j = 0, \dots, 2^k - 1$. Figure 1 illustrates a bifurcating tree navigating the partition down to level three for $\Omega = [0, 1)$. It will be convenient to index the partition elements using base 2 subscript and drop the superscript so that, for example, B_{000} indicates the first set in level 3, B_{0011} the fourth set in level 4 and so on.

To define a random measure on Ω we construct random measures on the sets B_j . It is instructive to imagine a particle cascading down through the tree such that at the j th junction the probability of turning left or right is θ_j and $(1 - \theta_j)$ respectively. In addition we consider θ_j to be a random variable with some appropriate distribution $\theta_j \sim \pi_j$. The sample path of the particle down to level k will be recorded in a vector $\epsilon_k = \{\epsilon_{k1}, \epsilon_{k2}, \dots, \epsilon_{kk}\}$ with elements $\epsilon_{ki} \in \{0, 1\}$, such that $\epsilon_{ki} = 0$ if the particle went left at level i , $\epsilon_{ki} = 1$ if it went right. Hence B_{ϵ_k} denotes which partition the particle belongs to at level k . By convention, set $\epsilon_0 = \emptyset$. Given a set of θ_j s it is clear that the probability of the particle falling into the set B_{ϵ_k} is just

$$P(B_{\epsilon_k}) = \prod_{i=1}^k (\theta_{\epsilon_{i-1}})^{(1-\epsilon_{ii})} (1 - \theta_{\epsilon_{i-1}})^{\epsilon_{ii}},$$

which is just the product of the probabilities of falling left or right at each junction that the particle passes through. This defines a random measure on the partitioning sets.

Let Π denote the collection of sets $\{B_0, B_1, B_{00}, \dots\}$ and let \mathcal{A} denote the collection of parameters that determine the distribution at each junction, $\mathcal{A} = (\alpha_{00}, \alpha_{01}, \alpha_{000}, \dots)$.

Definition 2.1. Lavine (1992) A random probability measure F on Ω is said to have a Pólya tree distribution, or a Pólya tree prior, with parameters (Π, \mathcal{A}) , written $F \sim$

$PT(\Pi, \mathcal{A})$, if there exists nonnegative numbers $\mathcal{A} = (\alpha_0, \alpha_1, \alpha_{00}, \dots)$ and random variables $\Theta = (\theta, \theta_0, \theta_1, \theta_{00}, \dots)$ such that the following hold:

1. the random variables in Θ are mutually independent;
2. for every $k = 1, 2, \dots$ and every $\epsilon_k \in \{0, 1\}^k$,

$$\theta_{\epsilon_k} \sim \text{Beta}(\alpha_{\epsilon_k 0}, \alpha_{\epsilon_k 1});$$

3. for every $k = 1, 2, \dots$ and every $\epsilon_k \in \{0, 1\}^k$,

$$F(B_{\epsilon_k} | \Theta) = \prod_{i=1}^k (\theta_{\epsilon_{i-1}})^{(1-\epsilon_{ii})} (1 - \theta_{\epsilon_{i-1}})^{\epsilon_{ii}}. \quad (1)$$

A random probability measure $F \sim PT(\Pi, \mathcal{A})$ is realized by sampling the θ_j s from the Beta distributions. Θ is countably infinite as the tree extends indefinitely, and hence for most practical applications the tree is specified only to a depth m . Lavine (1994) refers to this as a “partially specified” Pólya tree. It is worth noting that we will not need to make this truncation in what follows: our test will be fully specified with analytic expressions for the marginal likelihood.¹

By defining Π and \mathcal{A} , the Pólya tree can be centered on some chosen distribution G so that $E[F] = G$ where $F \sim PT(\Pi, \mathcal{A})$. Perhaps the simplest way to achieve this is to place the partitions in Π at the quantiles of G and then set $\alpha_{\epsilon_k 0} = \alpha_{\epsilon_j 1}$ for all $k = 1, 2, \dots$ and all $\epsilon_k \in \{0, 1\}^k$. (Lavine, 1992). For $\Omega \equiv \mathbb{R}$ this leads to $B_0 = (-\infty, G^{-1}(0.5))$, $B_1 = [G^{-1}(0.5), \infty)$ and, at level k ,

$$B_{\epsilon_k} = [G^{-1}\{(k^* - 1)/2^k\}, G^{-1}(k^*/2^k)], \quad (2)$$

where k^* is the decimal representation of the binary number ϵ_k .

It is usual to set the α 's to be constant in a level $\alpha_{\epsilon_m 0} = \alpha_{\epsilon_m 1} = c_m$ for some constant c_m . The setting of c_m governs the underlying continuity of the resulting F 's. For example, setting $c_m = cm^2$, $c > 0$, implies that F is absolutely continuous with probability 1 while $c_m = c/2^m$ defines a Dirichlet process which makes F discrete with probability 1 (Lavine, 1992; Ferguson, 1974). We will follow the approach of Walker and Mallick (1999) and define $c_m = cm^2$. The choice of c is discussed in Section 5.

2.1 Conditioning and marginal likelihood

An attractive feature of the Pólya tree prior is the ease with which we can condition on data. Pólya trees exhibit conjugacy: given a Pólya tree prior $F \sim PT(\Pi, \mathcal{A})$ and data \mathbf{y} drawn independently from F , then *a posteriori* F also has a Pólya tree distribution, $F | \mathbf{y} \sim PT(\Pi, \mathcal{A}^*)$ where \mathcal{A}^* is the set of updated parameters, $\mathcal{A}^* = \{\alpha_{00}^*, \alpha_{01}^*, \alpha_{000}^*, \dots\}$

$$\alpha_{\epsilon_i}^* | \mathbf{y} = \alpha_{\epsilon_i} + n_{\epsilon_i}, \quad (3)$$

¹Note however that consistency results only hold for a truncated version of the proposed test.

where n_{ϵ_i} denotes the number of observations in \mathbf{y} that lie in the partition B_{ϵ_i} . The corresponding random variables θ_j^* are therefore distributed *a posteriori* as

$$\theta_j^* | \mathbf{y} = \text{Beta}(\alpha_{j0} + n_{j0}, \alpha_{j1} + n_{j1}) \tag{4}$$

where n_{j0} and n_{j1} are the numbers of observations falling left and right at the junction in the tree indicated by j . This conjugacy allows for a straightforward calculation of the marginal likelihood for any set of observations, as

$$\Pr(\mathbf{y} | \Theta, \Pi, \mathcal{A}) = \prod_j \theta_j^{n_{j0}} (1 - \theta_j)^{n_{j1}} \tag{5}$$

where $\theta_j | \mathcal{A} \sim \text{Be}(\alpha_{j0}, \alpha_{j1})$ and where the product in (5) is over the set of all partitions, $j \in \{0, 1, 00, \dots\}$, though clearly for many partitions we have $n_{j0} = n_{j1} = 0$. Equation (5) has the form of a product of independent Binomial-Beta trials hence the marginal likelihood is,

$$\Pr(\mathbf{y} | \Pi, \mathcal{A}) = \prod_j \left(\frac{\Gamma(\alpha_{j0} + \alpha_{j1})}{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})} \frac{\Gamma(\alpha_{j0} + n_{j0})\Gamma(\alpha_{j1} + n_{j1})}{\Gamma(\alpha_{j0} + n_{j0} + \alpha_{j1} + n_{j1})} \right) \tag{6}$$

where $j \in \{0, 1, 00, \dots\}$. This marginal probability will form the basis of our test for H_0 which we describe in the next section.

3 A procedure for Bayesian nonparametric hypothesis testing

We are interested in providing a weight of evidence in favour of H_0 given the observed data. From Bayes theorem,

$$\Pr(H_0 | \mathbf{y}^{(1,2)}) \propto \Pr(\mathbf{y}^{(1,2)} | H_0) \Pr(H_0). \tag{7}$$

Under the null hypothesis H_0 , $\mathbf{y}^{(1)}$ and $\mathbf{y}^{(2)}$ are samples from some common distribution $F^{(1,2)}$ with $F^{(1,2)}$ unknown. We specify our uncertainty in $F^{(1,2)}$ via a Pólya tree prior, $F^{(1,2)} \sim PT(\Pi, \mathcal{A})$. Under H_1 , we assume $\mathbf{y}^{(1)} \sim F^{(1)}$, $\mathbf{y}^{(2)} \sim F^{(2)}$ with $F^{(1)}, F^{(2)}$ unknown. Again we adopt a Pólya tree prior for $F^{(1)}$ and $F^{(2)}$ with the same prior parameterization as for $F^{(1,2)}$ so that

$$F^{(1)}, F^{(2)}, F^{(1,2)} \stackrel{\text{iid}}{\sim} PT(\Pi, \mathcal{A}) \tag{8}$$

The logic for adopting a common prior distribution is that we regard the F s as random draws from some universe of distributions that we describe probabilistically through the Pólya tree distribution. Π is constructed from the quantiles of some *a priori* centering distribution. Following the approach of Walker and Mallick (1999); Mallick and Walker (2003) we take common values for the α_j s at each level as $\alpha_{j0} = \alpha_{j1} = cm^2$ for an α parameter at level m .

The posterior odds on H_0 is

$$\frac{\Pr(H_0|\mathbf{y}^{(1,2)})}{\Pr(H_1|\mathbf{y}^{(1)},\mathbf{y}^{(2)})} = \frac{\Pr(\mathbf{y}^{(1,2)}|H_0)}{\Pr(\mathbf{y}^{(1)},\mathbf{y}^{(2)}|H_1)} \frac{\Pr(H_0)}{\Pr(H_1)} \quad (9)$$

where the first term is just the ratio of marginal likelihoods, the Bayes Factor, which from (6) and conditional on our specification of Π and \mathcal{A} , is

$$\frac{P(\mathbf{y}^{(1,2)}|H_0)}{P(\mathbf{y}^{(1)},\mathbf{y}^{(2)}|H_1)} = \prod_j b_j \quad (10)$$

where

$$b_j = \frac{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})}{\Gamma(\alpha_{j0} + \alpha_{j1})} \frac{\Gamma(\alpha_{j0} + n_{j0}^{(1)} + n_{j0}^{(2)})\Gamma(\alpha_{j1} + n_{j1}^{(1)} + n_{j1}^{(2)})}{\Gamma(\alpha_{j0} + n_{j0}^{(1)} + n_{j0}^{(2)} + \alpha_{j1} + n_{j1}^{(1)} + n_{j1}^{(2)})} \\ \times \frac{\Gamma(\alpha_{j0} + n_{j0}^{(1)} + \alpha_{j1} + n_{j1}^{(1)})}{\Gamma(\alpha_{j0} + n_{j0}^{(1)})\Gamma(\alpha_{j1} + n_{j1}^{(1)})} \frac{\Gamma(\alpha_{j0} + n_{j0}^{(2)} + \alpha_{j1} + n_{j1}^{(2)})}{\Gamma(\alpha_{j0} + n_{j0}^{(2)})\Gamma(\alpha_{j1} + n_{j1}^{(2)})} \quad (11)$$

and the product in (10) is over all partitions, $j \in \{\emptyset, 0, 1, 00, \dots\}$, $n_{j0}^{(1)}$ and $n_{j1}^{(1)}$ represent the numbers of observations in $\mathbf{y}^{(1)}$ falling right and left at each junction and $n_{j0}^{(2)}$ and $n_{j1}^{(2)}$ are the equivalent quantities for $\mathbf{y}^{(2)}$. We can see from (10) that the overall Bayes Factor has the form of a product of Beta-Binomial tests at each junction in the tree to be interpreted as “does the data support one θ_j or two, $\{\theta_j^{(1)}, \theta_j^{(2)}\}$, in order to model the distribution of the observations going left and right at each junction?”, where for each j , $\theta_j \sim \text{Beta}(\alpha_j, \alpha_j)$. The product in (10) is defined over the infinite set of partitions. However, for each branch j , $b_j = 1$ if $n_{j0}^{(1)} + n_{j1}^{(1)} = 0$ or $n_{j0}^{(2)} + n_{j1}^{(2)} = 0$; hence to calculate (10) for the infinite partition structure we just have to multiply terms from junctions which contain at least some data from the two sets of samples. Hence, we only need specify Π to the quantile level where partitions contain observations from both samples. Note also that in the complete absence of data (that is, when $n_{j0}^{(1)} + n_{j0}^{(2)} = n_{j1}^{(1)} + n_{j1}^{(2)} = 0$)

$$b_j = \frac{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})}{\Gamma(\alpha_{j0} + \alpha_{j1})} \frac{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})}{\Gamma(\alpha_{j0} + \alpha_{j1})} \frac{\Gamma(\alpha_{j0} + \alpha_{j1})}{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})} \frac{\Gamma(\alpha_{j0} + \alpha_{j1})}{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})} = 1$$

for all j , so the Bayes Factor is 1.

The test procedure is described in Algorithm 1.

3.1 Prior specification

The Bayesian procedure requires the specification of $\{\Pi, \mathcal{A}\}$ in the Pólya tree. While there are good guidelines for setting \mathcal{A} the setting of Π is more problem specific, and the results will be quite sensitive to this choice. Our current, default, guideline is to first standardise the joint data $\mathbf{y}^{(1,2)}$ with the median and interquantile range of $\mathbf{y}^{(1,2)}$ and then set the partition on the quantiles of a standard normal density, $\Pi = \Phi(\cdot)^{-1}$.

Algorithm 1 Bayesian nonparametric test

1. Fix the binary tree on the quantiles of some centering distribution G .
 2. For level $m = 1, 2, \dots$, for each j set $\alpha_j = cm^2$ for some c .
 3. Add the log of the contributions of terms in (10) for each junction in the tree that has non-zero numbers of observations in $\mathbf{y}^{(1,2)}$ going both right and left.
 4. Report $\Pr(H_0|\mathbf{y}^{(1,2)})$ as $\Pr(H_0|\mathbf{y}^{(1,2)}) = \frac{1}{1+\exp(-LOR)}$, where LOR denotes the log odd ratio calculated at step 3.
-

We have found this to work well as a default in most situations, though of course the reader is encouraged to set Π according to their subjective beliefs.

In our algorithm, parameter c is treated as a fixed hyperparameter. As we demonstrate in Section 3.2, a truncated version of our test with a subjective partition is consistent under null and alternative hypotheses irrespective of the choice of c . However, c does have an impact on finite sample properties, that is, the finite sample posterior probabilities. This is always the case for Bayesian model selection/hypothesis testing based on the marginal likelihood, which is effectively a measure of how well the prior predicts the observed data, and not a feature restricted to our nonparametric procedure. In Section 5 we provide some guidelines on the sensitivity of the testing procedure to the value of this parameter, and discuss empirical Bayes estimation of c .

3.2 Consistency

Conditions for the consistency of the procedure under the null hypothesis and alternative hypothesis are developed for a related test based on a truncation of the Bayes factor (9). Let $n = n_{\emptyset}^{(1)} + n_{\emptyset}^{(2)}$ be the total sample size for both samples. Let $l(\epsilon)$ be the length of the vector ϵ . This also indicates that B_ϵ forms part of the partition at level $l(\epsilon)$, and in our construction there are $2^{l(\epsilon)}$ partition elements at level $l(\epsilon)$. We consider the test statistics based on the truncated Bayes factor

$$BF_{\kappa_0} = \prod_{\{j|l(j) \leq \kappa_0\}} b_j \tag{12}$$

where $\kappa_0 \in \mathbb{N}$ defines the level of truncation and can be set arbitrarily large. We also consider a truncated version of the hypothesis test:

$$\begin{aligned} H_{0,\kappa_0} &: \forall \epsilon |l(\epsilon) \leq \kappa_0, \quad F^{(1)}(B_\epsilon) = F^{(2)}(B_\epsilon) \\ \text{versus} \\ H_{1,\kappa_0} &: \exists \epsilon |l(\epsilon) \leq \kappa_0 \text{ and } F^{(1)}(B_\epsilon) \neq F^{(2)}(B_\epsilon). \end{aligned}$$

First, assume H_{0,κ_0} is true and let F_0 denote the true distribution. To prove consistency under H_{0,κ_0} , it is sufficient to show that

$$\lim_{n \rightarrow \infty} \log BF_{\kappa_0} = \infty$$

as $n \rightarrow \infty$ if both samples are drawn from the same distribution.

Theorem 3.1. *Suppose that the limiting proportion of observations in the first sample exists and is β_\emptyset :*

$$\beta_\emptyset = \lim_{n \rightarrow \infty} \frac{n_\emptyset^{(1)}}{n_\emptyset^{(1)} + n_\emptyset^{(2)}}. \quad (13)$$

If $0 < \beta_\emptyset < 1$ then, under H_{0, κ_0} ,

$$\lim_{n \rightarrow \infty} \log BF_{\kappa_0} = \infty$$

and the test defined by Algorithm 1, truncated at level κ_0 , is consistent under the null.

Proof. See Appendix. ■

We now consider consistency under H_{1, κ_0} for the truncated version of the test.

Theorem 3.2. *Assume that $0 < \beta_\emptyset < 1$, and that exists B_ϵ , $l(\epsilon) \leq \kappa_0$, such that $F^{(1)}(B_\epsilon)F^{(2)}(B_\epsilon) > 0$ and $\frac{F^{(1)}(B_\epsilon)}{F^{(1)}(B_\epsilon)} \neq \frac{F^{(2)}(B_\epsilon)}{F^{(2)}(B_\epsilon)}$. Then*

$$\lim_{n \rightarrow \infty} BF_{\kappa_0} = 0$$

and the test defined by Algorithm 1, truncated at level κ_0 , is consistent under the alternative.

Proof. See Appendix. ■

The proofs for consistency for the non-truncated test are much more challenging, as one needs to bound terms at each level of the Pólya tree. In the next section, we provide numerical experiments on the evolution of the Bayes factor with respect to the sample size under both H_0 and H_1 , suggesting consistency for the non-truncated test.

3.3 Simulations

To examine the operating performance of the method we consider the following experiments designed to explore various canonical departures from the null.

- a) Mean shift: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim \mathcal{N}(\theta, 1)$, $\theta = 0, \dots, 3$.
- b) Variance shift: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim \mathcal{N}(0, \theta^2)$, $\theta = 1, \dots, 3$.
- c) Mixture: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim \frac{1}{2}\mathcal{N}(\theta, 1) + \frac{1}{2}\mathcal{N}(-\theta, 1)$, $\theta = 0, \dots, 3$.
- d) Tails: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim t(\theta^{-1})$, $\theta = 10^{-3}, \dots, 10$ where $t(\nu)$ denotes the standard Student t distribution with ν degrees of freedom.
- e) Lognormal mean shift: $\log Y^{(1)} \sim \mathcal{N}(0, 1)$, $\log Y^{(2)} \sim \mathcal{N}(\theta, 1)$, $\theta = 0, \dots, 3$.
- f) Lognormal variance shift: $\log Y^{(1)} \sim \mathcal{N}(0, 1)$, $\log Y^{(2)} \sim \mathcal{N}(0, \theta^2)$, $\theta = 1, \dots, 3$.

The default mean distribution $F_0^{(1,2)} = \mathcal{N}(0, 1)$ was used in the Pólya tree to construct the partition Π and $\alpha = m^2$. Data are standardized. Comparisons are performed with $n_0 = n_1 = 50$ against the two-sample Kolmogorov-Smirnov and Wilcoxon rank test. To compare the models we explore the “power to detect the alternative”. As a test statistic for the Bayesian model we simulate data under the null and then take the empirical 0.95 quantile of the distribution of Bayes Factors as a threshold to declare H_1 . This is known as “the Bayes, non-Bayes compromise” by Good (1992). Results, based on 1000 replications, are reported in Figure 2. As a general rule we can see that the KS test is more sensitive to changes in central location while the Bayes test is more sensitive to changes to tails or higher moments.

The dyadic partition structure of the Pólya Tree allows us to breakdown the contribution to the Bayes Factor by levels. That is, we can explore the contribution, in the log of equation (10), by level. This is shown in Figure 3 as boxplots of the distribution of log BF statistics across the levels for the simulations generated for Figure 2. This is a strength of the Pólya tree test in that it provides a qualitative and quantitative decomposition of the contribution to the evidence against the null from differing levels of the tree.

It is also of interest to investigate the behavior of the Bayes factor as a function of the sample size, both under the null and various alternatives. Under the alternative, we consider in particular the following cases:

- a) Mean shift: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim \mathcal{N}(1, 1)$.
- b) Variance shift: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim \mathcal{N}(0, 4)$.
- c) Tails: $Y^{(1)} \sim \mathcal{N}(0, 1)$, $Y^{(2)} \sim t(1)$.

The results are reported in Figure 4 for sample size $n = 10, 50, 100, 200$ with 500 replications, and seem to suggest that the non-truncated test is consistent under the null and alternative.

4 A conditional procedure

The Bayesian procedure above requires the subjective specification of the partition structure Π . This subjective setting may make some users uneasy regarding the sensitivity to specification. In this section we explore an empirical procedure whereby the partition Π is centered on the data via the empirical cdf of the joint data $\hat{\Pi} = [\hat{F}^{(1,2)}]^{-1}$.

Let $\hat{\Pi}$ be the partition constructed with the quantiles of the empirical distribution $\hat{F}^{(1,2)}$ of $\mathbf{y}^{(1,2)}$. Under H_0 , there are now no free parameters and only one degree of freedom in the random variables $\{n_{j0}^{(1)}, n_{j1}^{(1)}, n_{j0}^{(2)}, n_{j1}^{(2)}\}$ as conditional on the partition centered on the empirical cdf of the joint, once one of the variables has been specified the others are then known. We consider, arbitrarily, the marginal distribution of $\{n_{j0}^{(1)}\}$

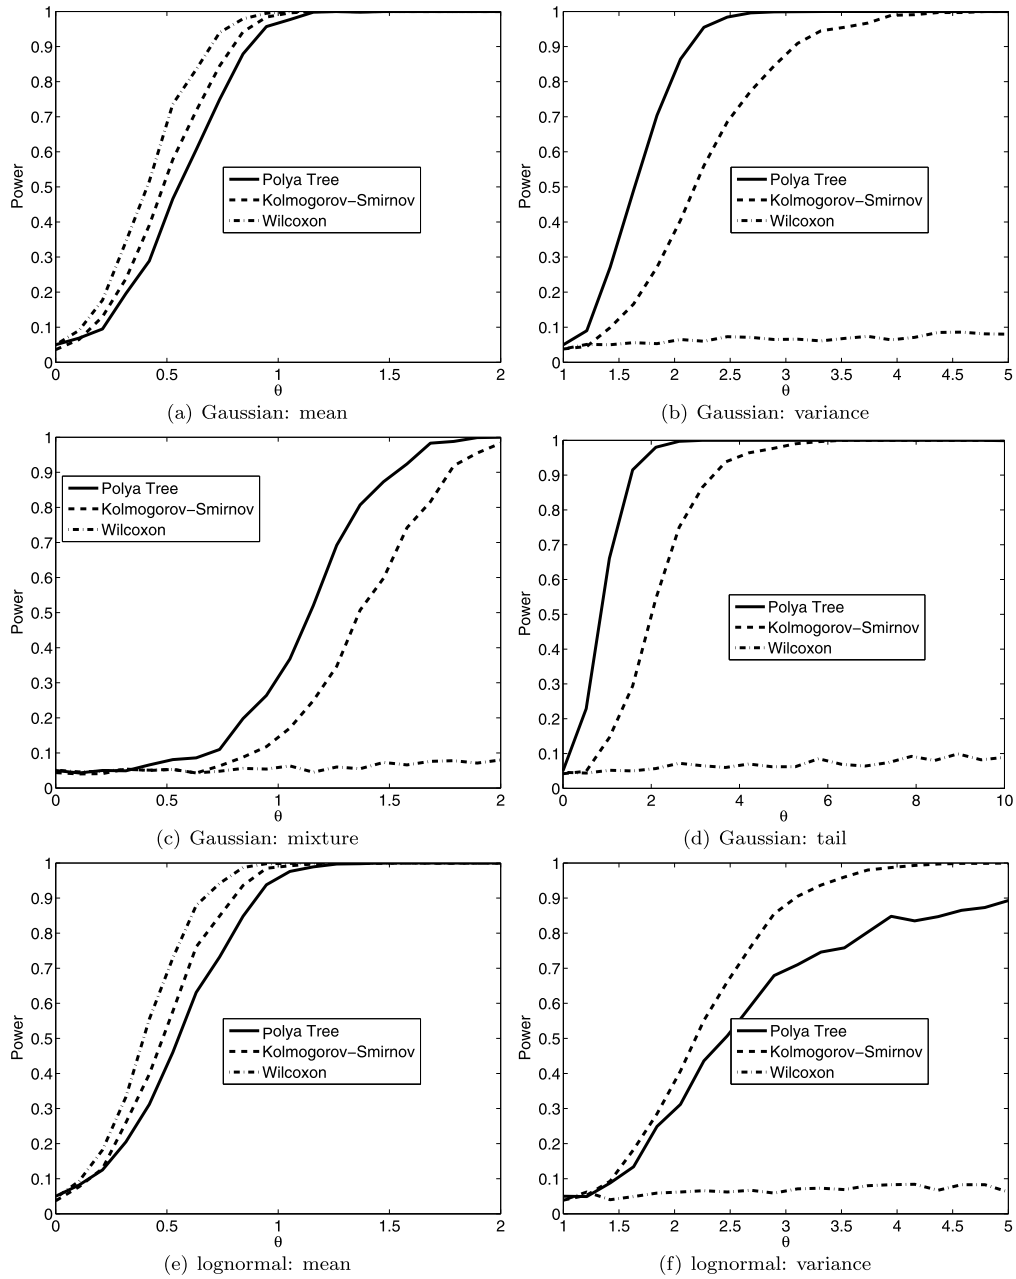


Figure 2: Power of Bayes test with $\alpha_j = m^2$ on simulations from Section 3.3., with x-axis measuring θ , the parameter in the alternative. Legend: K-S (dashed), Wilcoxon (dot-dashed), Bayesian test (solid).

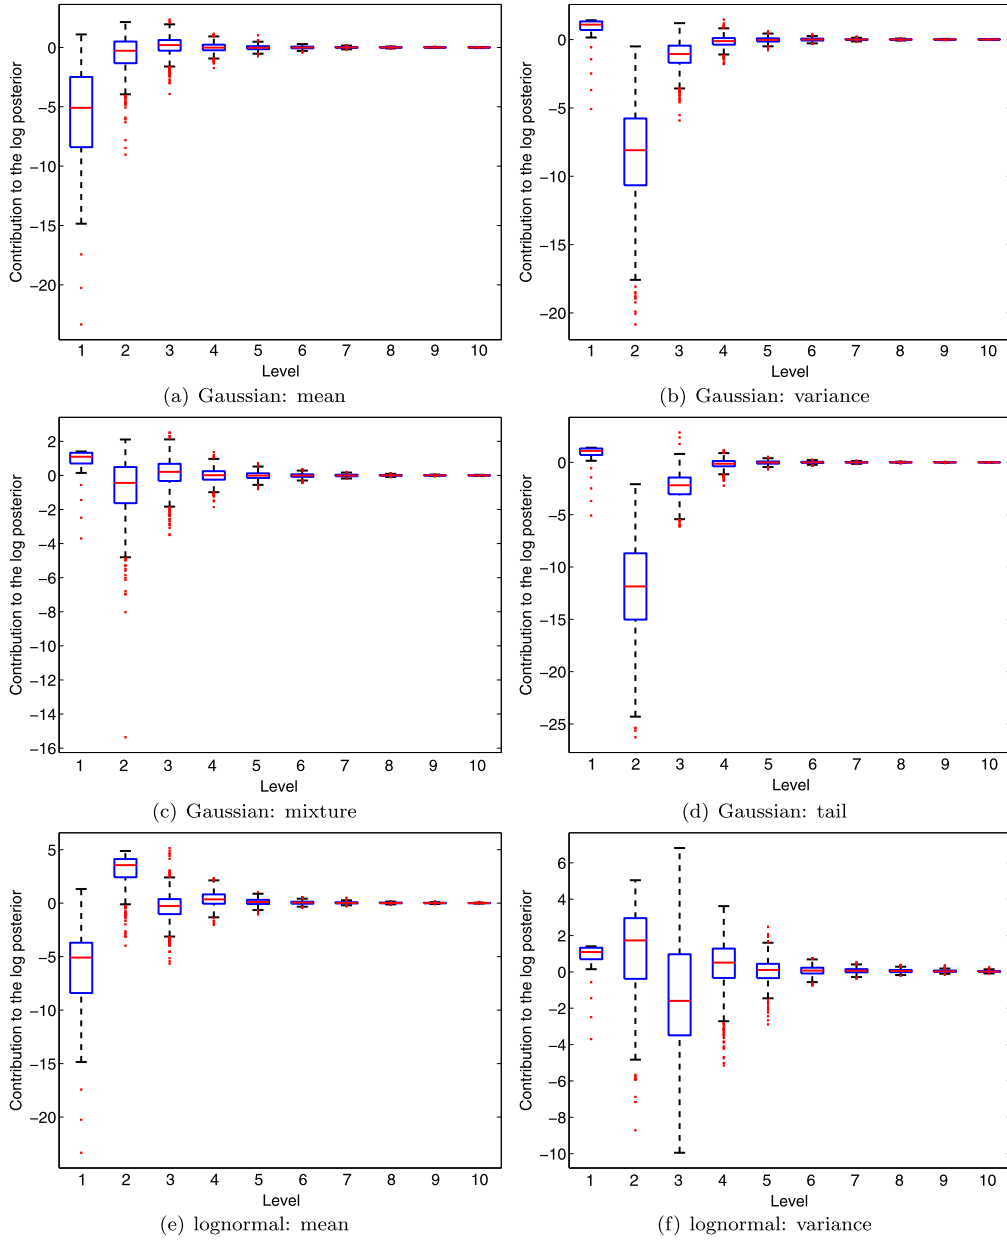


Figure 3: Contribution to Bayes Factors from different levels of the Pólya Tree under the alternative. Gaussian distribution with varying (a) mean (b) variance (c) mixture (d) tails; log-normal distribution with varying (e) mean (f) variance, from Section 3.2. Parameters of H_1 were set to the mid-points of the x-axis in Figure 2.

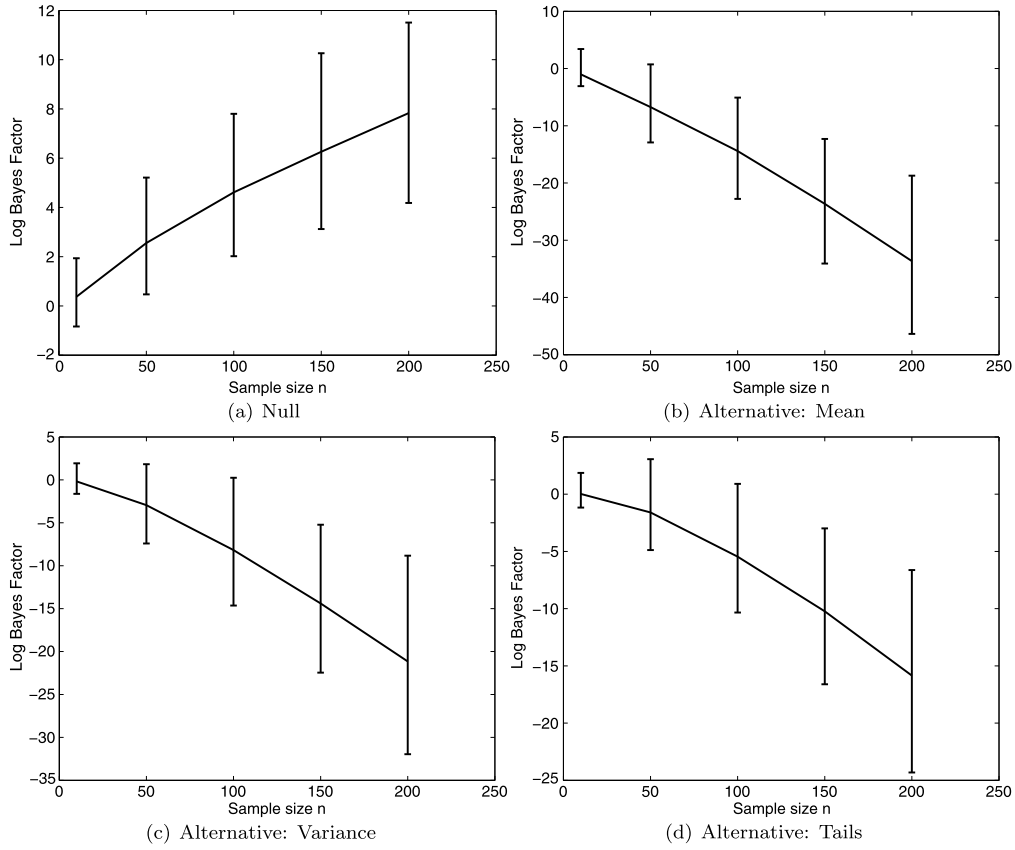


Figure 4: Mean Bayes factor and 90% confidence interval with respect to the sample size n under (a) the null and (b-c) two Gaussian distributions with different (b) means, (c) variances and (d) a Gaussian and a Student t .

which is now a product of hypergeometric distributions (we only consider levels where $n_j^{(1,2)} > 1$)

$$\Pr(\{n_{j0}^{(1)}\} | H_0, \hat{\Pi}, \mathcal{A}) \propto \prod_j \binom{n_j^{(1)}}{n_{j0}^{(1)}} \binom{n_j^{(1,2)} - n_{j0}^{(1)}}{n_{j0}^{(1,2)} - n_{j0}^{(1)}} / \binom{n_j^{(1,2)}}{n_{j0}^{(1,2)}} \quad (14)$$

$$= \prod_j \text{HG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}) \quad (15)$$

if $\max(0, n_{j0}^{(1,2)} + n_j^{(1)} - n_j^{(1,2)}) \leq n_{j0}^{(1)} \leq \min(n_j^{(1)}, n_{j0}^{(1,2)})$, and zero otherwise.

Under H_1 , the marginal distribution of $\{n_{j0}^{(1)}\}$ is a product of the conditional distri-

bution of independent binomial variates, conditional on their sum,

$$\Pr \left(\{n_{j0}^{(1)}\} | H_1, \hat{\Pi}, \mathcal{A} \right) \propto \prod_j \frac{g(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \theta_j^{(1)}, \theta_j^{(2)})}{\sum_x g(x; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \theta_j^{(1)}, \theta_j^{(2)})} \tag{16}$$

if $\max(0, n_{j0}^{(1,2)} + n_j^{(1)} - n_j^{(1,2)}) \leq n_j^{(1)} \leq \min(n_j^{(1)}, n_{j0}^{(1,2)})$, zero otherwise, and where

$$g(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \theta_j^{(1)}, \theta_j^{(2)}) = \text{Binomial}(n_{j0}^{(1)}; n_j^{(1)}, \theta_j^{(1)}) \times \dots \\ \text{Binomial}(n_{j0}^{(1,2)} - n_{j0}^{(1)}; n_j^{(1,2)} - n_j^{(1)}, \theta_j^{(2)})$$

and

$$\theta_j^{(1)} | \mathcal{A} \sim \text{Beta}(\alpha_{j0}, \alpha_{j1}) \quad \theta_j^{(2)} | \mathcal{A} \sim \text{Beta}(\alpha_{j0}, \alpha_{j1}).$$

Now, consider the odds ratio $\omega_j = \frac{\theta_j^{(1)}(1 - \theta_j^{(2)})}{\theta_j^{(2)}(1 - \theta_j^{(1)})}$ and let

$$\text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \omega_j) = \frac{g(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \theta_j^{(1)}, \theta_j^{(2)})}{\sum_x g(x; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \theta_j^{(1)}, \theta_j^{(2)})}.$$

Then it can be seen that $\text{EHG}(x; N, m, n, \omega)$ is the extended hypergeometric distribution (Harkness, 1965) whose pdf is proportional to

$$\text{HG}(x; N, m, n)\omega^x, \quad a \leq x \leq b,$$

where $a = \max(0, n + m - N)$, $b = \min(m, n)$. Note there are C++ and R routines to evaluate the pdf. The extended hypergeometric distribution models a biased urn sampling scheme whereby there is a different likelihood of drawing one type of ball over another at each draw. The Bayes factor is now given by

$$BF = \prod_j \frac{\text{HG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)})}{\int_0^\infty \text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \omega_j) p(\omega_j) d\omega_j} \tag{17}$$

where the marginal likelihood in the denominator can be evaluated using importance sampling or one-dimensional quadrature.

The conditional Bayes two-sample test can then be given in a similar way to Algorithm 1 but now using (17) for the contribution at each junction. Conditions for the consistency of the procedure under the null hypothesis are developed for a related test based on a truncation of the Bayes factor (17) although we have been unable to show consistency under the alternative.

Theorem 4.1. Consider the Bayes factor (17) truncated at level κ_0

$$BF_{\kappa_0} = \prod_{j: l(j) \leq \kappa_0} \frac{\text{HG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)})}{\int_0^\infty \text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \omega_j) p(\omega_j) d\omega_j}. \tag{18}$$

Suppose that β_\emptyset is as defined in Equation (13). If $0 < \beta_\emptyset < 1$ then, under H_{0,κ_0} ,

$$\lim_{n \rightarrow \infty} \log BF_{\kappa_0} = \infty$$

and the test is consistent under the null.

Proof. See Appendix. ■

We repeated the simulations from Section 3.3 with $\alpha = m^2$. The results are shown in Figures 5 and 6. We observe similar behaviour to the test with subjective partition but importantly we see that the problem in detecting the difference between normal and t-distribution is corrected. Note that no standardisation of the data is required for this test.

5 Sensitivity to the parameter c

The parameter c acts as a precision parameter in the Pólya tree and consequently can have an effect on the hypothesis testing procedures previously described. In principle, the parameter can be chosen subjectively as with precision parameters in other models (such as the linear model). Its effect is perhaps most easily understood through the prior variance of $P(B_{\epsilon_k})$ which has the form (Hanson, 2006)

$$\text{Var}[P(B_{\epsilon_k})] = 4^{-k} \left[\prod_{j=1}^k \frac{2cj^2 + 2}{2cj^2 + 1} - 1 \right].$$

The prior variance tends to zero as $c \rightarrow \infty$ and so the nonparametric prior places mass on distributions which more closely resemble the centering distribution as c increases. Another consequence of this is that, under H_1 , c determines the *a priori* expected squared Euclidean distance between $F^{(1)}(B_{\epsilon_k})$ and $F^{(2)}(B_{\epsilon_k})$, where $F^{(1)}$ and $F^{(2)}$ are presumed independently drawn from $PT(\Pi, \mathcal{A})$; this distance diminishes as c increases:

$$\begin{aligned} \mathbb{E}[(F^{(1)}(B_{\epsilon_k}) - F^{(2)}(B_{\epsilon_k}))^2] &= \text{Var}[F^{(1)}(B_{\epsilon_k})] + \text{Var}[F^{(2)}(B_{\epsilon_k})] \\ &= 4^{-k+\frac{1}{2}} \left[\prod_{j=1}^k \frac{2cj^2 + 2}{2cj^2 + 1} - 1 \right]. \end{aligned}$$

The value of c can be chosen to control the rate at which the variances decreases. We have found values of c between 1 and 10 work well in practice. Figures 7 and 8 show results for different values of c . As with any Bayesian testing procedure, we recommend checking the sensitivity of their results to the chosen value of the hyperparameter c .

An alternative approach to the choice of c in hypothesis testing is given by Berger and Guglielmi (2001) in the context of testing a parametric model against a nonparametric alternative. They argue that the minimum of the Bayes factor in favour of the parametric model is useful since the parametric model can be considered satisfactory

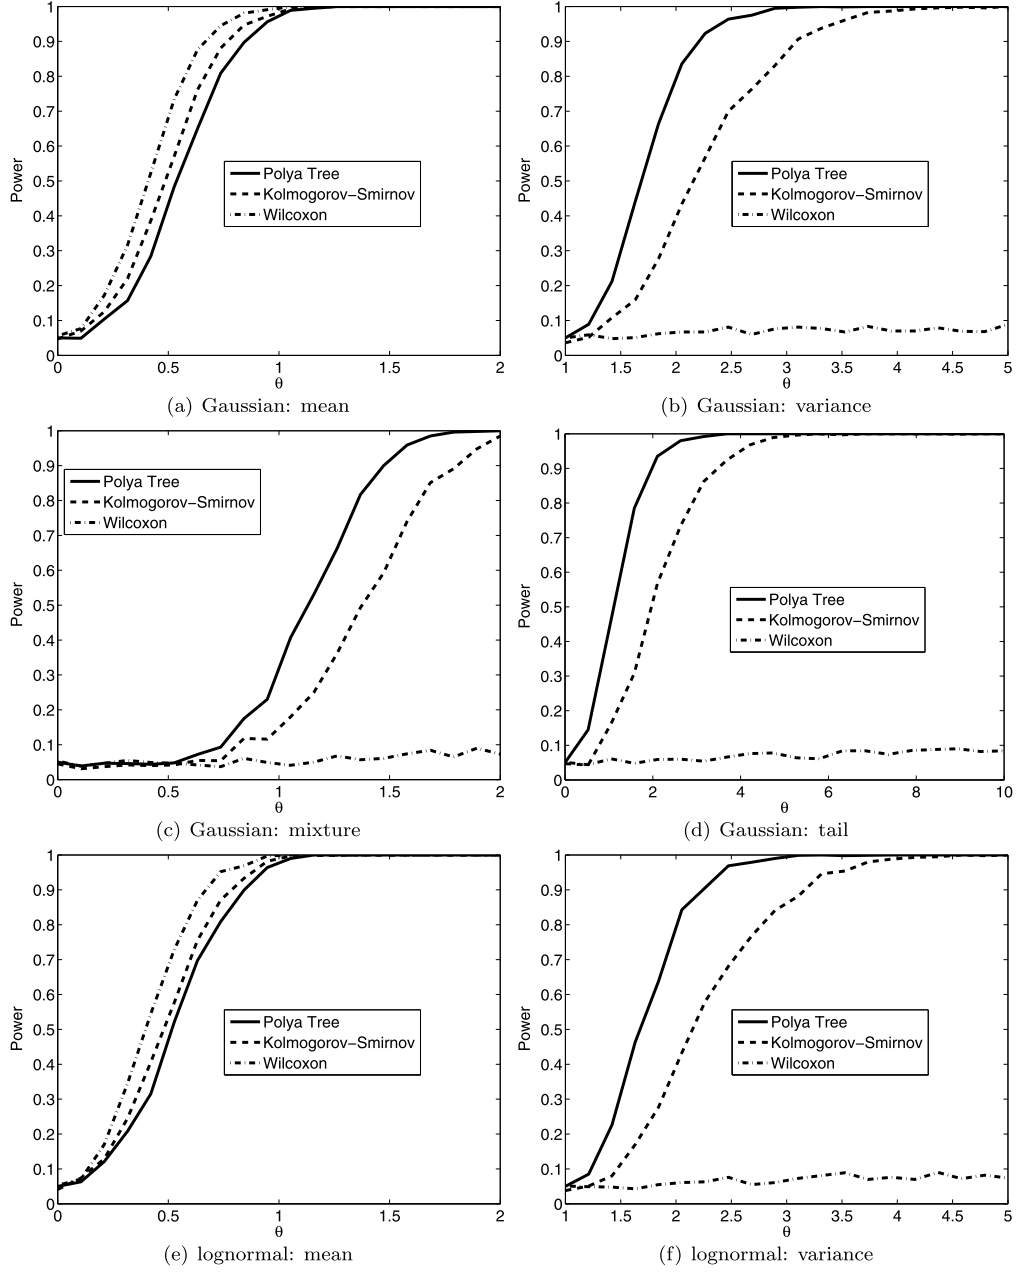


Figure 5: As in Figure 2 but now using conditional Bayes Test with $\alpha_j = m^2$.

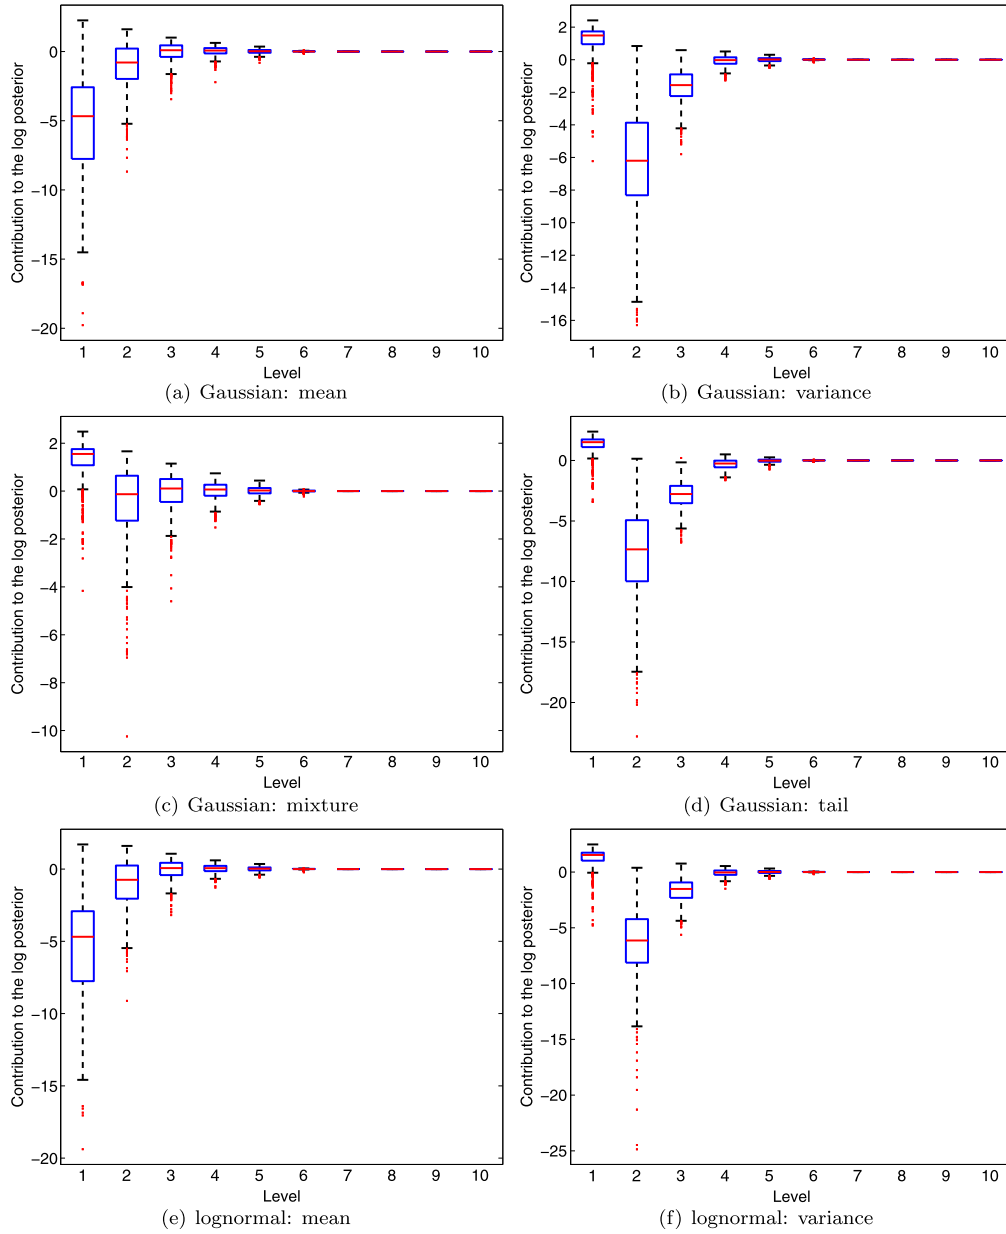


Figure 6: Contribution to the Bayes Factor for different levels of the conditional Pólya tree prior for Gaussian distribution with varying (a) mean (b) variance (c) mixture (d) tails; log-normal distribution with varying (e) mean (f) variance.

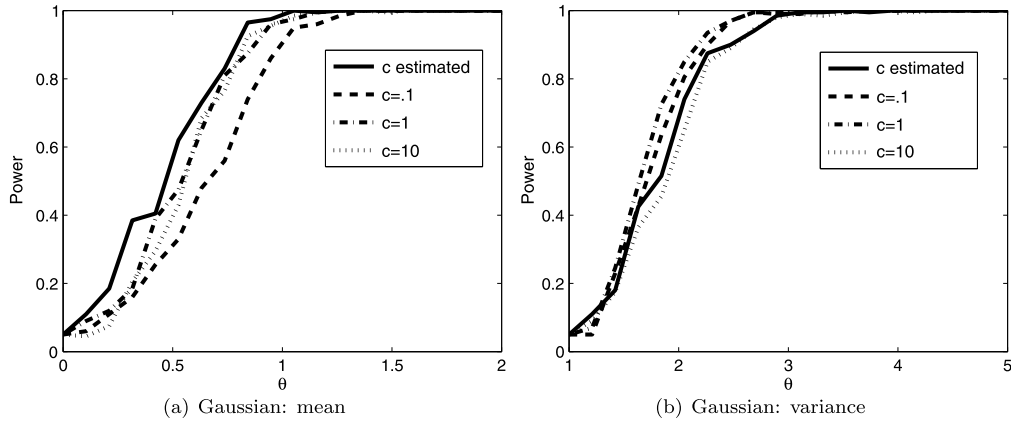


Figure 7: Subjective test with empirical Bayes estimation of the parameter c for (a) mean shift (b) variance shift. Point estimates of c are obtained by maximizing both the marginal likelihood under the null and alternative over the grid of values 10^i for $i = -2, -1, \dots, 3$.

if the “minimum is not too small”. It is the Bayes factor calculated using the empirical Bayes (Type II maximum likelihood) estimate of c . We suggest taking a similar approach if c cannot be subjectively chosen. In the test with subjective partition, the empirical Bayes estimates \hat{c} are calculated under H_0 and under H_1 . Using these values, the Bayes factor can be interpreted as a likelihood ratio statistic for the comparison of the two hypotheses. In the conditional test, the empirical Bayes estimate is calculated only under H_1 (since the marginal likelihood under H_0 does not depend on c). Figures 7 and 8 provide results for mean and variance shifts with c estimated over a fixed grid using this procedure.

We also performed experiments to test the sensitivity of the procedure to the partition. Experiments (not reported here) with a partition centered on a standard Student t distribution showed little difference compared to a partition centered on a standard Gaussian distribution.

6 Discussion and related work

There have been several other approaches to testing the difference between two distributions using Pólya tree based approaches. Ma and Wong (2011) propose the Coupling Optional Pólya tree (co-opt) prior which extends their previous work on Optional Pólya tree priors (Wong and Ma, 2010). The Optional Pólya tree defines a prior for a single distribution. A prior is defined on the sequence of partitions used to construct the Pólya tree. This allows the partition to be concentrated on areas of the sample space which have the largest difference from the base measure (which is the uniform in their case). The co-opt prior is suitable for two distributions with a partition defined for each distribution. The prior allows coupling of two partitions so that if a set A (which is member

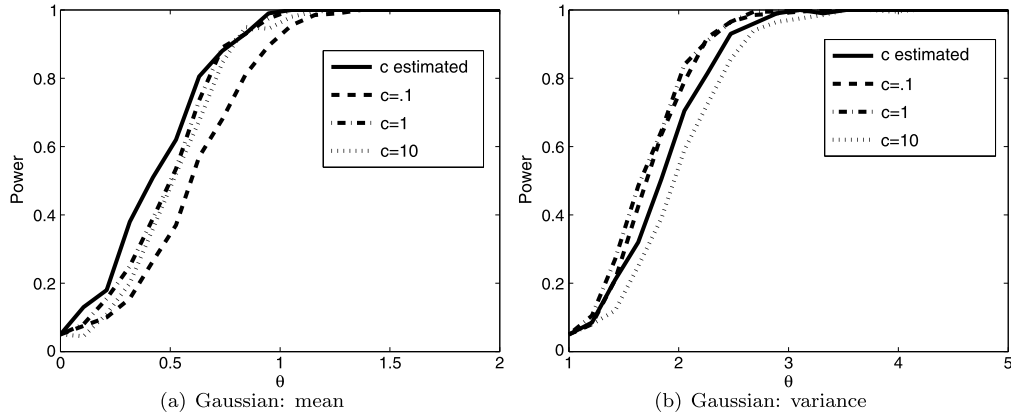


Figure 8: Conditional test with empirical Bayes estimation of the parameter c for (a) mean shift (b) variance shift. Point estimate of c is obtained by maximizing the marginal likelihood under the alternative over the grid of values 10^i for $i = -2, -1, \dots, 3$.

of the partition at level m) is coupled then all subsequent partitions of A will be the same for the two distributions. This allows the posterior to concentrate these couplings on parts of the support of the two distributions where they are similar and so allows the borrowing of strength between the two distributions. The prior is conjugate and can be used to both test for differences between two distributions and to infer where these differences occur in the posterior. The prior is particularly suited to multivariate problems due to its ability to efficiently learn the partition of the data. The posterior is available in closed form but can be computationally expensive to calculate in practice with computational time scaling exponentially with sample size.

Chen and Hanson (2014) propose a method for comparing k -samples of data which may be censored. They test the null hypothesis that the distribution of each sample is the same against the alternative that the samples arise from at least two different distributions. Under the null hypothesis, the common distribution is given a Pólya tree prior whereas, under the alternative hypothesis, each distribution is given an independent Pólya tree prior. All Pólya tree priors are centered over a normal distribution whose parameters are estimated using maximum likelihood to define an empirical Bayes procedure. Different partitions are used for the Pólya tree distributions under the null and the alternative hypotheses and so the partition must be truncated in order to compute the Bayes factor, contrary to our simpler approach which involves no truncation.

7 Conclusions

We have described a Bayesian nonparametric hypothesis test for real valued data which provides an explicit measure of $\Pr(H_0|\mathbf{y}^{(1,2)})$. The test is based on a fully specified Pólya tree prior for which we are able to derive an explicit form for the Bayes Factor.

Conditioning on a particular partition can lead to a predictive distribution that exhibits jumps at the partition boundary points. This is a well-known phenomenon of Pólya tree priors and some interesting directions to mitigate its effects can be found in Hanson and Johnson (2002); Paddock et al. (2003); Hanson (2006). We do not consider these approaches here as mixing over partitions would lose the analytic tractability of our approach, but it is an interesting area for future study and is considered by Chen and Hanson (2014).

Acknowledgements

This research was partly supported by the Oxford-Man Institute (OMI) through a visit by FC to the OMI. FC thanks Pierre Del Moral for helpful discussions on the proofs of consistency and acknowledges the support of the European Commission under the Marie Curie Intra-European Fellowship Programme. DS acknowledges the support of a Discovery Grant from the Natural Sciences and Engineering Council of Canada. The authors thank reviewers of earlier versions of this article for helpful comments.

References

- Basu, S. and Chib, S. (2003). “Marginal likelihood and Bayes factors for Dirichlet process mixture models.” *Journal of the American Statistical Association*, 98: 224–235. [298](#)
- Berger, J. O. and Guglielmi, A. (2001). “Testing of a parametric model versus nonparametric alternatives.” *Journal of the American Statistical Association*, 96: 174–184. [298](#), [310](#)
- Bernardo, J. M. and Smith, A. F. M. (2000). *Bayesian theory*. Chichester: John Wiley. [298](#)
- Bhattacharya, A. and Dunson, D. B. (2012). “Nonparametric Bayes Classification and Hypothesis Testing on Manifolds.” *Journal of multivariate analysis*, 111: 1–19. [298](#)
- Branscum, A. J. and Hanson, T. J. (2008). “Bayesian nonparametric meta-analysis using Polya tree mixture models.” *Biometrics*, 64: 825–833. [298](#)
- Carota, C. and Parmigiani, G. (1996). “On Bayes Factors for Nonparametric Alternatives.” In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M. (eds.), *Bayesian Statistics 5*, 508–511. London: Oxford University Press. [298](#)
- Chen, Y. and Hanson, T. E. (2014). “Bayesian nonparametric k-sample tests for censored and uncensored data.” *Computational Statistics & Data Analysis*, 71: 335–340. [298](#), [314](#), [315](#)
- Dass, S. C. and Lee, J. (2004). “A note on the consistency of Bayes factors for testing point null versus non-parametric alternatives.” *Journal of Statistical Planning and Inference*, 119: 143–152. [298](#)
- Dunson, D. B. and Peddada, S. D. (2008). “Bayesian nonparametric inference on stochastic ordering.” *Biometrika*, 95: 859–874. [298](#)

- Ferguson, T. S. (1974). "Prior distributions on spaces of probability measures." *The Annals of Statistics*, 2: 615–629. [297](#), [299](#), [300](#)
- Florens, J. P., Richard, J. F., and Rolin, J. M. (1996). "Bayesian Encompassing Specification Tests of a Parametric Model Against a Nonparametric Alternative." Technical Report 96.08, Institut de Statistique, Université Catholique de Louvain. [298](#)
- Ghosal, S., Lember, J., and Van der Vaart, A. (2008). "Nonparametric Bayesian model selection and averaging." *Electronic Journal of Statistics*, 2: 63–89. [298](#)
- Good, I. J. (1992). "The Bayes/non-Bayes compromise: a brief review." *Journal of the American Statistical Association*, 87: 597–606. [305](#)
- Hanson, T. E. (2006). "Inference for mixtures of finite Polya tree models." *Journal of the American Statistical Association*, 101: 1548–1565. [298](#), [310](#), [315](#)
- Hanson, T. E. and Johnson, W. O. (2002). "Modeling regression error with a mixture of Polya trees." *Journal of the American Statistical Association*, 97: 1020–1033. [315](#)
- Harkness, W. L. (1965). "Properties of the extended hypergeometric distribution." *The Annals of Mathematical Statistics*, 36: 938–945. [309](#), [320](#)
- Holmes, C. C., Caron, F., Griffin, J. E., and Stephens, D. A. (2009). "Two-sample Bayesian nonparametric hypothesis testing." Technical report, arXiv:0910.5060 [stat.ME]. [298](#)
- Kass, R. and Raftery, A. (1995). "Bayes factors." *Journal of the American Statistical Association*, 90: 773–795. [320](#)
- Lavine, M. (1992). "Some aspects of Polya tree distributions for statistical modelling." *The Annals of Statistics*, 20: 1222–1235. [297](#), [299](#), [300](#)
- (1994). "More aspects of Polya tree distributions for statistical modelling." *The Annals of Statistics*, 22: 1161–1176. [297](#), [299](#), [300](#)
- Ma, L. and Wong, W. H. (2011). "Coupling optional Pólya trees and the two sample problem." *Journal of the American Statistical Association*, 106(496): 1553–1565. [298](#), [313](#)
- Mallick, B. K. and Walker, S. G. (2003). "A Bayesian semiparametric transformation model incorporating frailties." *Journal of Statistical Planning and Inference*, 112: 159–174. [301](#)
- Mauldin, R. D., Sudderth, W. D., and Williams, S. C. (1992). "Polya trees and random distributions." *The Annals of Statistics*, 20: 1203–1221. [297](#), [299](#)
- McVinish, R., Rousseau, J., and Mengersen, K. (2009). "Bayesian Goodness of Fit Testing with Mixtures of Triangular Distributions." *Scandinavian Journal of Statistics*, 36: 337–354. [298](#)
- Paddock, S. M., Ruggeri, F., Lavine, M., and West, M. (2003). "Randomized Polya tree models for nonparametric Bayesian inference." *Statistica Sinica*, 13: 443–460. [315](#)
- Pennell, M. L. and Dunson, D. B. (2008). "Nonparametric Bayes Testing of Changes in a Response Distribution with an Ordinal Predictor." *Biometrics*, 64: 413–423. [298](#)

- Rousseau, J. (2007). “Approximating Interval hypothesis : p-values and Bayes factors.” In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M. (eds.), *Bayesian Statistics 8*. Oxford University Press. 298
- Walker, S. G. and Mallick, B. K. (1999). “A Bayesian semiparametric accelerated failure time model.” *Biometrics*, 55: 477–483. 300, 301
- Wilks, S. (1938). “The large-sample distribution of the likelihood ratio for testing composite hypotheses.” *The Annals of Mathematical Statistics*, 9: 60–62. 318, 320
- Wong, W. H. and Ma, L. (2010). “Optional Polya tree and Bayesian inference.” *The Annals of Statistics*, 38(3): 1433–1459. 313

Appendix: Proofs

Proof of Theorem 1

Clearly the log Bayes factor is

$$\log BF_{\kappa_0} = \sum_{\{j|l(j) \leq \kappa_0\}} \log b_j^{(n)}.$$

Stirling’s approximation of the Gamma function allows us to write

$$\begin{aligned} b_j^{(n)} &\simeq \frac{\Gamma(\alpha_{j0})\Gamma(\alpha_{j1})}{\Gamma(\alpha_{j0} + \alpha_{j1})} \frac{1}{\sqrt{2\pi}} \frac{(\hat{p}_{j0}^{(1,2)})^{\alpha_{j0}-1/2} (\hat{p}_{j1}^{(1,2)})^{\alpha_{j1}-1/2}}{(\hat{p}_{j0}^{(1)})^{\alpha_{j0}-1/2} (\hat{p}_{j1}^{(1)})^{\alpha_{j1}-1/2} (\hat{p}_{j0}^{(2)})^{\alpha_{j0}-1/2} (\hat{p}_{j1}^{(2)})^{\alpha_{j1}-1/2}} \\ &\times \sqrt{\frac{n_j^{(1)} n_j^{(2)}}{n_j^{(1,2)}}} \times \frac{(\hat{p}_{j0}^{(1,2)})^{n_{j0}^{(1,2)}} (\hat{p}_{j1}^{(1,2)})^{n_{j1}^{(1,2)}}}{(\hat{p}_{j0}^{(1)})^{n_{j0}^{(1)}} (\hat{p}_{j1}^{(1)})^{n_{j1}^{(1)}} (\hat{p}_{j0}^{(2)})^{n_{j0}^{(2)}} (\hat{p}_{j1}^{(2)})^{n_{j1}^{(2)}}} \end{aligned} \quad (19)$$

where

$$\hat{p}_{j0}^{(k)} = \frac{n_{j0}^{(k)}}{n_{j0}^{(k)} + n_{j1}^{(k)}} \quad \hat{p}_{j1}^{(k)} = 1 - \hat{p}_{j0}^{(k)}.$$

We have, under the null,

$$\sqrt{\frac{n_j^{(1)} n_j^{(2)}}{n_j^{(1,2)}}} \simeq \sqrt{n} \sqrt{F_0(B_j) \beta_\emptyset (1 - \beta_\emptyset)}. \quad (20)$$

The term

$$L_j = \frac{(\hat{p}_{j0}^{(1,2)})^{n_{j0}^{(1,2)}} (\hat{p}_{j1}^{(1,2)})^{n_{j1}^{(1,2)}}}{(\hat{p}_{j0}^{(1)})^{n_{j0}^{(1)}} (\hat{p}_{j1}^{(1)})^{n_{j1}^{(1)}} (\hat{p}_{j0}^{(2)})^{n_{j0}^{(2)}} (\hat{p}_{j1}^{(2)})^{n_{j1}^{(2)}}} \quad (21)$$

is a likelihood ratio for testing composite hypotheses

$$H_{j0} : p_{j0}^{(1)} = p_{j0}^{(2)} = p_{j0}^{(1,2)} \text{ vs } H_{j1} : (p_{j0}^{(1)}, p_{j0}^{(2)}) \in [0, 1]^2$$

with $n_{j0}^{(1)} \sim \text{Binomial}(n_j^{(1)}, p_{j0}^{(1)})$ and $n_{j0}^{(2)} \sim \text{Binomial}(n_j^{(2)}, p_{j0}^{(2)})$. Clearly $\widehat{p}_{j0}^{(1,2)}$ and $(\widehat{p}_{j0}^{(1)}, \widehat{p}_{j0}^{(2)})$ are the maximum likelihood estimators under H_{j0} and H_{j1} respectively. It follows that, under H_{j0} , $-2 \log L_j$ asymptotically follows a χ^2 distribution (Wilks, 1938).

Finally, if $\beta_\emptyset(1 - \beta_\emptyset) > 0$ and using Equation (20), then Theorem 1 follows.

Proof of Theorem 2

If $F^{(1)}(B_j) = 0$ or $F^{(2)}(B_j) = 0$, then we have trivially $\log(b_j) = 0$. We assume that

$$F^{(1)}(B_j)F^{(2)}(B_j) > 0 \quad (22)$$

$$0 < \beta_\emptyset < 1. \quad (23)$$

If $p_{j0}^{(1)} = p_{j0}^{(2)}$ then, from the previous section, $\log(b_j)$ goes to ∞ in $o(\log(n))$. We consider now the case $p_{j0}^{(1)} \neq p_{j0}^{(2)}$.

$$\text{Let } \widehat{\beta}_j^{(n)} = n_j^{(1)} / n_j^{(1,2)},$$

$$\beta_j = \lim_{n \rightarrow \infty} \widehat{\beta}_j^{(n)} = \frac{\beta_\emptyset F^{(1)}(B_j)}{F^{(1,2)}(B_j)}$$

with $F^{(1,2)}(B_j) = \beta_\emptyset F^{(1)}(B_j) + (1 - \beta_\emptyset)F^{(2)}(B_j)$. Under assumptions (22) and (23), we have $0 < \beta_j < 1$. Let L_j be defined as in Equation (21). We have

$$\log L_j = \eta_j - n_j^{(1,2)} \zeta_j$$

where

$$\eta_j = n_j^{(1,2)} H_{p_{j0}^{(1,2)}}(\widehat{p}_{j0}^{(1,2)}) - n_j^{(1)} H_{p_{j0}^{(1)}}(\widehat{p}_{j0}^{(1)}) - n_j^{(2)} H_{p_{j0}^{(2)}}(\widehat{p}_{j0}^{(2)}),$$

$$\zeta_j = \widehat{\beta}_j^{(n)} H_1(p_{j0}^{(1)}) + (1 - \widehat{\beta}_j^{(n)}) H_1(p_{j0}^{(2)}) - H_1(p_{j0}^{(1,2)}),$$

$$p_{j0}^{(k)} = \lim_{n \rightarrow \infty} \widehat{p}_{j0}^{(k)}, \quad k = 1, 2, \{1, 2\}$$

and the function $H_p : x \in (0, 1) \rightarrow \mathbb{R}_+$ is defined for $p \in (0, 1)$ by

$$H_p(x) = x \log\left(\frac{x}{p}\right) + (1 - x) \log\left(\frac{1 - x}{1 - p}\right).$$

Consider first the term ζ_j . We have

$$p_{j0}^{(1,2)} = \beta_j p_{j0}^{(1)} + (1 - \beta_j) p_{j0}^{(2)}$$

and $n \rightarrow \infty$, therefore

$$\zeta_j \rightarrow \beta_j H_1(p_{j0}^{(1)}) + (1 - \beta_j) H_1(p_{j0}^{(2)}) - H_1(p_{j0}^{(1,2)}).$$

As the function H_p is convex, it follows that ζ_j tends to a positive constant if $p_{j0}^{(1)} \neq p_{j0}^{(2)}$. Hence $n_j^{(1,2)} \zeta_j$ tends to ∞ in $o(n)$.

Consider now η_j which is approximately equal to

$$Y_j = \frac{n_j^{(1,2)}}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} (\widehat{p}_{j0}^{(1,2)} - p_{j0}^{(1,2)})^2 - \frac{n_j^{(1)}}{2p_{j0}^{(1)}(1-p_{j0}^{(1)})} (\widehat{p}_{j0}^{(1)} - p_{j0}^{(1)})^2 - \frac{n_j^{(2)}}{2p_{j0}^{(2)}(1-p_{j0}^{(2)})} (\widehat{p}_{j0}^{(2)} - p_{j0}^{(2)})^2.$$

Then

$$Y_j = -\frac{n_j^{(1,2)}\widehat{\beta}_j^{(n)}(1-\widehat{\beta}_j^{(n)})}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} [\widehat{p}_{j0}^{(1)} - \widehat{p}_{j0}^{(2)} - (p_{j0}^{(1)} - p_{j0}^{(2)})]^2 + \frac{n_j^{(1,2)}\widehat{\beta}_j^{(n)}(1-\rho_j^{(1)})}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} (\widehat{p}_{j0}^{(1)} - p_{j0}^{(1)})^2 + \frac{n_j^{(1,2)}(1-\widehat{\beta}_j^{(n)})(1-\rho_j^{(2)})}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} (\widehat{p}_{j0}^{(2)} - p_{j0}^{(2)})^2$$

where

$$\rho_j^{(k)} = \frac{p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})}{p_{j0}^{(k)}(1-p_{j0}^{(k)})}.$$

We have asymptotically

$$Y_j \simeq nF^{(1,2)}(B_j) \left(\frac{-\beta_j(1-\beta_j)}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} [\widehat{p}_{j0}^{(1)} - \widehat{p}_{j0}^{(2)} - (p_{j0}^{(1)} - p_{j0}^{(2)})]^2 + \frac{\beta_j(1-\rho_j^{(1)})}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} (\widehat{p}_{j0}^{(1)} - p_{j0}^{(1)})^2 + \frac{(1-\beta_j)(1-\rho_j^{(2)})}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} (\widehat{p}_{j0}^{(2)} - p_{j0}^{(2)})^2 \right).$$

We have a quadratic form

$$Y_j \simeq -nF^{(1,2)}(B_j) \frac{\beta_j(1-\beta_j)}{2p_{j0}^{(1,2)}(1-p_{j0}^{(1,2)})} \begin{bmatrix} (\widehat{p}_{j0}^{(1)} - p_{j0}^{(1)}) & (\widehat{p}_{j0}^{(2)} - p_{j0}^{(2)}) \end{bmatrix} A \begin{bmatrix} (\widehat{p}_{j0}^{(1)} - p_{j0}^{(1)}) \\ (\widehat{p}_{j0}^{(2)} - p_{j0}^{(2)}) \end{bmatrix}$$

with square matrix

$$A = \begin{bmatrix} \frac{\rho_j^{(1)} - \beta_j}{1 - \beta_j} & -1 \\ -1 & \frac{\rho_j^{(2)} - 1 + \beta_j}{\beta_j} \end{bmatrix}$$

and so asymptotically

$$\sqrt{n} \begin{bmatrix} (\widehat{p}_{j0}^{(1)} - p_{j0}^{(1)}) \\ (\widehat{p}_{j0}^{(2)} - p_{j0}^{(2)}) \end{bmatrix} \sim \mathcal{N} \left(0_2, \begin{pmatrix} \frac{p_{j0}^{(1)}(1-p_{j0}^{(1)})}{\beta_j F^{(1,2)}(B_j)} & 0 \\ 0 & \frac{p_{j0}^{(2)}(1-p_{j0}^{(2)})}{(1-\beta_j)F^{(1,2)}(B_j)} \end{pmatrix} \right)$$

which is independent of the sample size, and it follows that Y_j asymptotically follows a scaled χ^2 distribution and $\log L_j$ goes to $-\infty$ in $o(n)$. To conclude, we have that in Equation (19)

$$\sqrt{\frac{n_j^{(1)} n_j^{(2)}}{n_j^{(1,2)}}} \simeq \sqrt{n} \sqrt{F^{(1,2)}(B_j) \beta_j (1 - \beta_j)}.$$

It follows that

- If $F^{(1)}(B_j)F^{(2)}(B_j) = 0$, $\log(b_j) = 0$.
- If $F^{(1)}(B_j)F^{(2)}(B_j) > 0$
 - If $p_{j0}^{(1)} \neq p_{j0}^{(2)}$, the log-contribution $\log(b_j)$ goes to $-\infty$ at a rate of $o(n)$.
 - If $p_{j0}^{(1)} = p_{j0}^{(2)}$, the log-contribution $\log(b_j)$ goes to $+\infty$ at a rate of $o(\log n)$.

Proof of Theorem 3

The condition $\beta_\emptyset(1 - \beta_\emptyset)$ implies that $\beta_j(1 - \beta_j) > 0$ for all j . Let

$$b_j^{(n)} = \frac{\text{HG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)})}{\int_0^\infty \exp\{u(\omega_j)\} d\omega_j} \quad (24)$$

where $u(\omega_j) = \log \text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \omega_j) + \log p(\omega_j)$. Then

$$\log BF_{\kappa_0} = \sum_{\{j | l(j) \leq \kappa_0\}} \log b_j^{(n)}.$$

Under the conditions $\beta_j(1 - \beta_j) > 0$, the maximum likelihood estimate $\hat{\omega}_j$ of the parameter ω_j in the extended hypergeometric distribution converges in probability to the true parameter (Harkness, 1965, p. 944). We can therefore use a Laplace approximation (Kass and Raftery, 1995) of the denominator in (24), we obtain for $n_{j0}^{(1)}$ large

$$b_j^{(n)} \simeq \frac{\text{HG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)})}{\sqrt{2\pi} |\hat{\Sigma}_j|^{1/2} \text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \hat{\omega}_j) p(\hat{\omega}_j)}$$

where $\hat{\omega}_j = \text{argmax}_{\omega_j} u(\omega_j)$ and $\hat{\Sigma}_j^{-1} = -\mathbf{D}^2 u_j(\hat{\omega}_j)$, where $\mathbf{D}^2 u_j(\hat{\omega}_j)$ is the Hessian matrix of second derivatives. The ratio

$$r_j = \frac{\text{HG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)})}{\text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \hat{\omega}_j)} = \frac{\text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, 1)}{\text{EHG}(n_{j0}^{(1)}; n_j^{(1,2)}, n_j^{(1)}, n_{j0}^{(1,2)}, \hat{\omega}_j)}$$

is a likelihood ratio for testing the composite hypotheses

$$H_{j0} : \omega_j = 1 \text{ vs } H_{j1} : \omega_j > 0,$$

hence $-2 \log r_j$ is asymptotically χ^2 -distributed (Wilks, 1938). And as $|\hat{\Sigma}_j| \rightarrow 0$ as $n \rightarrow \infty$, then $b_j^{(n)} \rightarrow \infty$ for all j as $n \rightarrow \infty$.