

THE DEPENDENCE OF ROUTINE BAYESIAN MODEL SELECTION METHODS ON IRRELEVANT ALTERNATIVES

PIOTR ZWIERNIK AND JIM Q. SMITH

ABSTRACT. Bayesian methods - either based on Bayes Factors or BIC - are now widely used for model selection. One property that might reasonably be demanded of any model selection method is that if a model M_1 is preferred to a model M_0 , when these two models are expressed as members of one model class \mathbb{M} , this preference is preserved when they are embedded in a different class \mathbb{M}' . However, we illustrate in this paper that with the usual implementation of these common Bayesian procedures this property does not hold true even approximately. We therefore contend that to use these methods it is first necessary for there to exist a "natural" embedding class. We argue that in any context like the one illustrated in our running example of Bayesian model selection of binary phylogenetic trees there is no such embedding.

1. INTRODUCTION

Bayesian method such as Bayes Factor (BF) (for example [Denison et al. \(2002\)](#)), or ones based on the Bayesian Information Criterion (BIC) [Schwarz \(1978\)](#) are now widespread in statistical analysis. In this paper we show that a disadvantage of these approaches, in the way these are routinely used, is that they can give rise to an awkward inferential ambiguity, which we later argue for some problems can never be resolved.

Selecting a model can be seen as a special case of providing a preference order over a number of options, which we assume to be finite. A long time ago [Nash \(1950\)](#) and [Thomas \(1984\)](#) argue that a particular property - restated below in terms that apply for model selection - is essential for preference orderings used to identify an optimal choice. This states that if M_1 is preferred to M_0 (written $M_1 \succ M_0$) in the model class \mathbb{M} then $M_1 \succ M_0$ also in $\mathbb{M}' := \mathbb{M} \cup \mathbb{M}^+$ whenever $M_1 \succ M^+$ for all $M^+ \in \mathbb{M}^+$. So in particular if M^* is a best model in \mathbb{M} and we extend our selection to contain other models all of which are worse than M^* , then M^* remains a best model in the larger set of models. This natural property is called *independence of irrelevant alternatives* (IIA).

If IIA does not hold, then, whether or not we include poor fitting models in the selection set will influence what model we label as "best". Why should the choice of this model be influenced by inclusion or exclusion of other candidate models later discovered to be poor representatives of the underlying data generating process? Surely *any* routine method of Bayesian Model selection: for example such as those reviewed by [Bernardo & Smith \(1994\)](#)

The first author was supported by Jan Draisma's Vidi grant from the Netherlands Organisation for Scientific Research (NWO).

and [Key et al. \(1999\)](#) - even ones assuming that all models were wrong - should be expected not to violate IIA.

Happily most applications of a naive Bayes Factor (BF) model selection satisfy the property of IIA in the following sense. Let \mathbf{X} be the random vector of observations over which models are selected, taking values \mathbf{x} in the sample space $\mathcal{X} \subseteq \mathbb{R}^m$. For simplicity assume that all predictive densities associated with these models are bounded and strictly positive over their shared support. Let M_0 be a reference model with predictive density $p_0(\mathbf{x})$. Let M_ϕ be another candidate model with predictive density $p_\phi(\mathbf{x})$ on $\mathbf{x} \in \mathcal{X}$ and a prior odds relative to M_0 of λ_ϕ , where $(p_\phi(\mathbf{x}), \lambda_\phi)$ are functions of ϕ alone and in particular not the predictive density and prior model probability $(p_{\phi'}(\mathbf{x}), \lambda_{\phi'})$ for any other candidate model $M_{\phi'}$. Then setting $\lambda_0 = 1$, we note immediately that $M_{\phi_2} \succeq M_{\phi_1}$ if and only if for the value of \mathbf{x} we observe

$$\log p_{\phi_2}(\mathbf{x}) + \log \lambda_{\phi_2} \geq \log p_{\phi_1}(\mathbf{x}) + \log \lambda_{\phi_1}.$$

Clearly this preference is therefore unaffected by the values of $(p_{\phi'}(\mathbf{x}), \lambda_{\phi'})$ for $M_{\phi'}$ which may or may not be contained in the selection class.

Our problems begin when the class of models \mathbb{M} over which selection takes place is extremely large. Then the necessary task of carefully and individually choosing separate prior distributions over the hyperparameters of each candidate model is clearly infeasible. We are therefore forced to reference our choice of prior density over the parameters of each candidate so that in some sense these are consistent with each other. In this situation $(p_\phi(\mathbf{x}), \lambda_\phi)$ is highly related with $(p_{\phi'}(\mathbf{x}), \lambda_{\phi'})$ for other candidate models and the BF may break the IIA property.

For example this is exactly what happens when selecting over the class of Bayesian networks (BNs). In this case [Heckerman et al. \(1995\)](#) introduced an additional demand that the prior densities satisfied *parameter modularity*. Here when parts of the structure of two of these multivariate models coincide, then the priors over this shared structure in these two models are assumed to be the same. This assumption, and others like it, not only makes the assignment of priors across a large model class feasible but also makes it possible to use greedy search algorithms to efficiently search the space for the best candidate model.

However, these methods come at a price. Because prior densities are chosen to match those given within the model class, the choice of class itself can affect the inference and in particular IIA can be violated. Of course occasionally, in simple applications, there are compelling reasons why a particular model class should be used. Then the violation of IIA is not a problem. However, the choice of embedding class is often chosen for convenience or convention rather than for some phenomenological reason associated with the modelled process. It is in these circumstances that the violation of IIA gives rise to poor inference.

One such setting occurs when the modeler must decide whether or not to include variables which usefully explain the process but cannot be observed. Within the Bayesian paradigm the fact that these variables are not observed does not matter in any formal sense: the score will simply be the log-marginal density of the observed variables where we marginalize over the missing ones. Thus this missingness causes no methodological problems. Indeed, if these integrations cannot be executed in closed form then their score can

be calculated in a straightforward manner using now standard numerical techniques. The problem is rather that two statistically equivalent models will be treated differently depending on whether or not the variables representing these underlying causes are included.

To be more concrete we consider the class of *phylogenetic tree models*. The evolution of a collection of different extant species is typically represented by a Bayesian network on a directed tree, possibly with some additional constraints on the parameter vector (see for example [Semple & Steel \(2003\)](#), [Yang \(2006\)](#)). The extant species are represented by the leaves of a tree, whose interior vertices label extinct ancestors. The model is usually depicted as a tree with edge lengths like for example in [Figure 1](#). The topology of the tree represents the underlying graphical model. The lengths are functions of the conditional probabilities parametrising the model and in the phylogenetic context they give a measure of the phylogenetic distance between two species.




FIGURE 1. A simple directed phylogenetic tree.

In its simplest form each variable on the tree is binary and represents the presence or absence of a characteristic in a large class of genetic locations, believed to occur at random with a particular probability determined by the vertex. The usual evolutionary hypotheses tell us that, if the tree is valid, then collections of variables separated from each other by any internal vertex are independent of each other given the value of that vertex. Our problem is to select an evolutionary tree that gives the most plausible evolutionary explanation of the data we have observed on the extant species. For any tree there is a formal Bayesian selection method to do this. We simply assign a prior density to each parameter, using methods such as those described by [Heckerman \(1999\)](#) respecting parameter modularity, calculate the corresponding log marginal likelihood score marginalising over the hidden variables and choose the tree scoring the highest.

The problem occurs because it can be proved (see for example [Settimi & Smith \(2000\)](#)) that a tree is statistically the same collection of sample densities over observed vertices of the graph as a simpler tree if it contains an interior vertex with only two neighbours. A common procedure is therefore to restrict the class of tree models to include only trees whose interior vertices have 3 or more neighbours. However, if the ensuing inference were to depend on the associated hypothesis that a hereditary ancestor existed *only* if it had

at least two non-extinct associated species is surely not satisfactory for two reasons. First, it seems quite conceivable that, in fact, there is only one direct descendant still surviving from a particular species. Second, even if we accepted the above, from an inferential point of view we should note that the property of having no degree two vertices (i.e. with two neighbours) is not closed under taking margins over a subset of the set of leaves of a tree. By this we mean that a subtree of the original tree will typically have degree two vertices.

The fact that a marginal tree contains degree two vertices has important practical implications whenever we want to include an additional species in our data set. This happens for example in the procedure of *outgroup rooting* when an outgroup is added in phylogenetic analysis in order to find the root of an unrooted phylogenetic tree as described by (Yang, 2006, Section 3.1.1.2). For a simple example imagine that the undirected tree on the left of Figure 2 has been chosen for four extant species. To find the root of this tree an outgroup 5 has been added. Let say the tree on the right side of Figure 2 has been found for the augmented data set. We have now introduced an additional hidden vertex c between 3 and 4. In this illustrated case the marginal model over $\{1, 2, 3, 4\}$ in the tree on the right-hand side coincides with the original model on the left. However, as we show in this paper, if we use automatic model selection methods, it is not clear that the best marginal model for $\{1, 2, 3, 4\}$ in the augmented data set will be the original model; and this is the minimal requirement for the outgroup rooting method to work in a consistent way.




FIGURE 2. Including an outgroup may distort the analysis.

In examples like the one given by phylogenetic tree models both standard BF and also BIC model selection methods do not in general respect the IIA property. In this paper we examine this phenomenon in much more detail with reference to the simplest possible manifestation of this ambiguity in Section 2. A natural question to ask is then whether these model selection methods are at least *approximately* invariant to this choice of embedding class (as appears to be the case of Consonni & La Rocca (2011) when addressing a rather different issue). Sadly the answer to this question is no! The embedding class can have a critical impact on the model selection even in the simplest cases. In Section 3 we analyse our basic example in a full Bayes factor model selection with conjugate priors, which is a default method in numerous R packages and Hugin software. In Section 4 we also show that the parametrization ambiguity also applies to BIC model selection and provide an in-depth geometrical explanation of this phenomenon. We show again that the embedding

class can have a critical impact on the model selection. We end the paper with a short discussion of the more general implications of this phenomenon.

2. THE RUNNING EXAMPLE

In our discussion we first want to make a clear distinction between different notions of statistical models used in this paper. A *model* for a random variable \mathbf{X} is a family of distributions of \mathbf{X} . A *parametric model* is a parametric family M_ψ of probability distributions of \mathbf{X} together with the defining parametrization ψ , which is a map from the given finite dimensional parameter space Θ to the space of all probability distributions. We say that a model M has a parametric formulation if there exists a parametrization defining this model. Let M_ϕ, M_ψ be two parametric models with parametrisations ϕ, ψ and parameter sets Θ, Ω respectively. We say that M_ϕ and M_ψ are parametric formulations of the same model if $\phi(\Theta) = \psi(\Omega)$ as models. Finally, a *Bayesian model* is a parametric model together with an associated family of prior distributions on the parameter space Θ . Hence the same model can have many parametric formulations and each, when combined with the associated prior distribution over parameters, can lead to a different Bayesian model.

To demonstrate the problem described in the introduction we will focus on a comparison of two simple models for a vector of two binary random variables X and Z under two different parametric formulations. The first model is the saturated model and the second is the model of independence $X \perp\!\!\!\perp Z$. The motivation is as follows. Suppose that we have two hypotheses: H_0 that X and Z are unrelated ($X \perp\!\!\!\perp Z$) and H_1 that there are evolutionary related ($X \rightarrow Z$). What we will demonstrate is that - with routine model settings - the second hypothesis must be distinguished from the one where we have an evolutionary relationship of the form $X \rightarrow Y \rightarrow Z$, where the ancestor of X and the predecessor of Z has not been observed. This is so even though both lead to exactly the same model for (X, Z) . The problem is that when we refer to evolutionary relation we certainly mean both $X \rightarrow Z$, $X \rightarrow Y \rightarrow Z$ and even $X \rightarrow Y_1 \rightarrow Y_2 \rightarrow Z$ simultaneously. The issue we have here is that therefore from a Bayesian point of view, for model selection we need to add all those intermediate vertices that might have occurred between the two observed vertices. But how do we determine this number and why should it impinge on our choice?

First, consider the “natural” parametric formulation where the saturated model is parametrised by the joint distributions and the independence model by the corresponding marginal distributions. Denote these parametric models by $M_1^{(0)}$ and $M_0^{(0)}$ respectively and the parameter spaces by $\Theta_0^{(0)} = [0, 1]^2$ and $\Theta_1^{(0)} = \{\theta_{xz} \in \mathbb{R}^4 : \sum_{i,j=0}^1 \theta_{xz}(i, j) = 1, \theta_{xz}(i, j) \geq 0\}$. Thus $M_1^{(0)}$ is given by $p_{xz}(i, j) = \theta_{xz}(i, j)$ and $M_0^{(0)}$ by $p_{xz}(i, j) = \theta_x(i)\theta_z(j)$ for $i, j = 0, 1$. The directed acyclic graphs (see [Lauritzen \(1996\)](#)) representing these models are given in [Figure 3](#).

Alternatively, consider two other parametric models for (X, Z) : model $M_1^{(1)}$ of conditional independence $X \perp\!\!\!\perp Z|Y$, implying the saturated model on the (X, Z) margin and represented by a graph $X \rightarrow Y \rightarrow Z$; and its submodel $M_0^{(1)}$ of marginal independence of




FIGURE 3. The directed acyclic graphs of the saturated model and the independence model under the first parametric formulation.

X and Z . Here we assume that Y is binary and not observed and the model is parametrized by the marginal distribution of X and conditional distributions of Y given X and of Z given Y .

More generally by $M_1^{(k)}$ for $k \geq 1$ denote the parametric model for (X, Z) given by the graph $X \rightarrow Y_1 \rightarrow \dots \rightarrow Y_k \rightarrow Z$, where all the Y_i are assumed to be binary and hidden. The corresponding marginal independence model is denoted by $M_0^{(k)}$ and it is a submodel of $M_1^{(k)}$. The parameter spaces are denoted by $\Theta_0^{(k)}$ and $\Theta_1^{(k)}$. The parameters of $M_1^{(k)}$ are given by the marginal distribution of X and conditional probability for each arrow of the corresponding graph. There are exactly $1 + 2(k + 1)$ free parameters denoted by $\theta_x(1)$, $\theta_{1|x}(1|i)$, $\theta_{2|1}(1|i)$, \dots , $\theta_{z|x}(1|i)$ for $i = 0, 1$ where for example $\theta_x(1) = P(X = 1)$, $\theta_{1|x}(1|0) = P(Y_1 = 1|X = 0)$ and $\theta_{j|j-1}(1|0) = P(Y_j = 1|Y_{j-1} = 0)$. It follows that $\Theta_1^{(k)} = [0, 1]^{1+2(k+1)}$.

By (Gilula, 1979, Corollary 1), $M_1^{(k)}$ for every $k \geq 1$ is the saturated model and hence it is equivalent to $M_1^{(0)}$. Since we also have that $M_0^{(k)}$ is equivalent to $M_0^{(0)}$ then for every $k \geq 0$ we compare the same models. Although the models $M_0^{(k)}$ are the same, as parametric or Bayesian models they are very different. If $k = 1$ then $M_0^{(1)}$ is a union of two parametric submodels $Y \perp\!\!\!\perp Z$, $X \perp\!\!\!\perp Y$ of $M_1^{(1)}$ as depicted in Figure 4 and hence $\Theta_0^{(1)}$ is isomorphic to a subspace of $\Theta_1^{(1)}$ given as a union of two intersecting components given by equations: $\theta_{1|x}(1|0) - \theta_{1|x}(1|1) = 0$ and $\theta_{z|1}(1|0) - \theta_{z|1}(1|1) = 0$. The common intersection locus is the singularity of $\Theta_0^{(1)}$. More generally, the larger is k the more complicated and more *singular* is the embedding of the parametric model $M_0^{(k)}$ in $M_1^{(k)}$. This gives the geometric intuition why the model selection for large k may differ from $k = 0$.




FIGURE 4. The saturated model and the independence model under the second parametric formulation with $k = 1$.

A statistical reason why IIA may fail follows from the discussion in the introduction. Thus, for fixed $k \geq 0$ consider the class \mathbb{M}_k of Bayesian models $M_1^{(k)}$ and $M_0^{(k)}$ together with the associated families of all prior distributions over the corresponding parameter spaces. First note that \mathbb{M}_0 can be naturally embedded into \mathbb{M}_1 because $M_0^{(0)} = M_0^{(1)}$ and $M_1^{(0)} = M_1^{(1)}$ as models and the parameter space $\Theta_1^{(0)}$ can be associated with a subspace of $\Theta_1^{(1)}$ for which X and Y are functionally related (e.g. $\theta_{y|x}(1|1) = \theta_{y|x}(0|0) = 1$). Hence every prior on $\Theta_1^{(0)}$ can be treated as a prior concentrated on this subspace. There is no statistical way of distinguishing between $M_1^{(0)}$ and its copy embedded in $M_1^{(1)}$. More generally $\mathbb{M}_{k-1} \subset \mathbb{M}_k$ for all $k \geq 1$. This follows from the fact that every prior distribution for parameters of the Bayesian models in \mathbb{M}_{k-1} is a degenerate prior distribution of \mathbb{M}_k . Hence, at least in principle, the statistical analysis should not depend on a particular embedding.

The problem is that the containment $\mathbb{M}_{k-1} \subset \mathbb{M}_k$ breaks down whenever we assume some sort of regularity of the prior distributions ruling out some possible Bayesian models. In this case the analysis may highly depend on k . This is particularly evident in the case of the BIC criterion which is derived under the assumption that prior distributions are diffuse (bounded and bounded away from zero) and hence then cannot be degenerate. The problem with diffuse priors and the BIC criterion has been reported in other contexts. In particular it has been shown that BIC tends to support simpler models which follows from the Jeffrey-Lindley's paradox (see for example [Denison et al. \(2002\)](#)). Recently it has been shown by [Johnson & Rossell \(2010\)](#) and [Consonni & La Rocca \(2011\)](#), that using local priors causes in particular a very slow convergence to the true smaller model. In [Section 4](#) we show that this problem may be particularly important when the choice of the parametric formulation is not clear.

3. FINITE SAMPLE CONJUGATE SELECTION

Perhaps the most common way to set up the prior densities across a large class of graphical models is to ensure each decomposable model within the class has consistent priors and the same strength in the following sense. Let $\pi = (\pi_{\mathbf{x}})$ denote the vector of prior probabilities that a unit from the model population takes level $\mathbf{x} \in \mathcal{X}$ where \mathcal{X} is the sample space of the problem. This joint prior distribution induces marginal and conditional distributions over subvectors of \mathbf{x} .

- Define $\beta_{\mathbf{x}} = \beta\pi_{\mathbf{x}}$ for $\mathbf{x} \in \mathcal{X}$, where the positive real scalar β , called the *effective sample size*, reflects the number of observations the modeler believes her prior beliefs are worth.
- Differentiate each model in the class with the appropriate hyper-Dirichlet prior (see [Dawid & Lauritzen \(1993\)](#)) faithful to its particular sets of conditional independence assumptions defined by its graph. The hyper-Dirichlet priors form the conjugate class for decomposable graphical models. They are widely used in numerous R packages and in Hugin (see [Madsen et al. \(2003\)](#)).

- Ensure the relevant prior expected clique cell probabilities and effective sample size parameters of these different product Dirichlets are consistent across the different models in the class. This is obtained by choosing the associated Dirichlet parameters β_i of the different (marginal or conditional) components C_i of BN model with prior probability π_i so that $\beta_i = \beta\pi_i$. In this way these Dirichlet distributions over various components are consistent with those of a Dirichlet over a saturated model with parameters $\beta_{\mathbf{x}}$ for $\mathbf{x} \in \mathcal{X}$. This in particular assures that the parameter modularity holds.

In what follows we therefore faithfully follow this procedure applying it to the different embeddings below, in addition ensuring that different embedding match in an analogous way. In this section we show that the selection procedure between these Bayesian models will not satisfy the IIA.

First we compare the models for $k = 0, 1$. Let π_{xyz} be the joint prior distribution on (X, Y, Z) . By π_{xz} , π_x , π_z , $\pi_{y|x}$ and $\pi_{z|y}$ we denote the corresponding prior marginal and conditional probabilities. Let first $k = 0$. For $M_1^{(0)}$ we set the prior distribution to be the Dirichlet distribution, $\theta_{xz} \sim \text{Dir}(\beta_{xz})$, where $\beta_{xz} = \beta\pi_{xz}$ and $\beta > 0$. For $M_0^{(0)}$ we have $\theta_x \sim \text{Beta}(\beta_x)$ and $\theta_z \sim \text{Beta}(\beta_z)$, where $\beta_x = \beta\pi_x$ and $\beta_z = \beta\pi_z$. Let now $k = 1$. The standard conjugate analysis for $M_1^{(1)}$ requires setting hyper-Dirichlet priors for the joint distribution of (X, Y, Z) . We set $\theta_x \sim \text{Beta}(\beta_x)$, $\theta_{y|x}(1|i) \sim \text{Beta}(\beta_{y|x}(i))$ and $\theta_{z|y}(\cdot|i) \sim \text{Beta}(\beta_{z|y}(i))$, where $\beta_{y|x}(i) = \beta\pi_{y|x}(\cdot|i)$ and $\beta_{z|y}(i) = \beta\pi_{z|y}(\cdot|i)$ for $i = 0, 1$. If we assume that all five random variables θ_x , $\theta_{y|x}(\cdot|0)$, $\theta_{y|x}(\cdot|1)$, $\theta_{z|y}(\cdot|0)$ and $\theta_{z|y}(\cdot|1)$ are independent, then the variable p_{xyz} , where $p_{xyz}(i, j, k) = \theta_x(i)\theta_{y|x}(j|i)\theta_{z|y}(k|j)$, has a hyper-Dirichlet distribution. This induces a distribution of p_{xz} by

$$(1) \quad p_{xz}(i, k) = p_{xyz}(i, 0, k) + p_{xyz}(i, 1, k) = \theta_x(i) \sum_{j=0}^1 \theta_{y|x}(j|i)\theta_{z|y}(k|j)$$

which is *not* in general Dirichlet as for $M_1^{(0)}$ and so gives a different value of the marginal likelihood in this case.

Since the induced distribution of p_x and p_z is the same as for $k = 0$, the difference in the analyses performed for the Bayes factor follows from the difference in the marginal likelihoods of the saturated models. For $M_0^{(0)}$ and $M_1^{(0)}$ we can easily obtain formulae for the posterior distribution and for the marginal likelihood functions of the sample counts $\mathbf{u} = [u_{ij}]$. So the Bayes factor can be calculated directly. For any table $\alpha = [\alpha_i]$ define the Beta function $B(\alpha) = \frac{\prod_i \Gamma(\alpha_i)}{\Gamma(\sum_i \alpha_i)}$ where $\Gamma(\cdot)$ is the Gamma function. The marginal likelihood for $M_1^{(0)}$ is:

$$(2) \quad L_1^{(0)} = \frac{B(\beta_{xz} + \mathbf{u})}{B(\beta_{xz})},$$

The marginal likelihood for $M_0^{(0)}$ is:

$$(3) \quad L_0^{(0)} = \frac{B(\boldsymbol{\beta}_x + \mathbf{u}_x) B(\boldsymbol{\beta}_z + \mathbf{u}_z)}{B(\boldsymbol{\beta}_x) B(\boldsymbol{\beta}_z)},$$

where \mathbf{u}_x and \mathbf{u}_z denote the marginal counts for X and Z respectively.

To obtain the marginal likelihood $L_1^{(1)}$ of $M_1^{(1)}$ the prior density for p_{xz} can be written explicitly using (1). Exact computations are technically inelegant but some recent developments of Lin et al. (2009) make it possible to compute the marginal integrals exactly at least for simple mixture models. We note that simple Monte Carlo approximations of the integrals below give very good results as well.

As an example consider the following table of counts:

$$\mathbf{u} = \begin{bmatrix} 13 & 9 \\ 4 & 18 \end{bmatrix}.$$

The exact Fisher's test in this case gives the value of the odds ratio 6.2 and the corresponding p -value is 0.012. Hence the data strongly support the alternative hypothesis. The Bayes factor for the first parametric formulation, denoted by BF_0 , is easily computed by dividing $L_1^{(0)}$ in (2) by $L_0^{(0)}$ in (3). For $k = 1$ we calculated scores using Lin et al. (2009) confirming these against good approximate results provided by simple Monte Carlo simulation. Initially assume that $\beta = 4$ and $\pi_{xyz}(i, j, k) = 1/8$ for all $i, j, k = 0, 1$. This gives p_{xz} a Dirichlet $\text{Dir}(1, 1, 1, 1)$ (and hence uniform) distribution if $k = 0$ (but not if $k = 1$) and:

$$(4) \quad \text{BF}_0 = 11.47, \quad \text{BF}_1 = 6.27.$$

So in the second parametric formulation, when $k = 1$, the Bayes factor slightly underestimates the evidence for the saturated model. The reason here is that the induced prior distribution on p_y , where $p_y(j) = \sum_{i,k} p_{xyz}(i, j, k)$, is not uniform and equals $\text{Beta}(2, 2)$. This is an important point because it shows that the uniform prior π may lead to highly informative scenarios, which may then affect our analysis.

The result of the analysis changes if we set $\beta = 4$ and $\pi_{xyz}(i, 0, k) = t/4$ and $\pi_{xyz}(i, 1, k) = (1-t)/4$ for $t < 0.5$. Note that the induced distribution on p_{xz} in the first parametric formulation is still uniform because $\pi_{xz}(i, k) = 1/4$ for all $i, k = 0, 1$ but the prior information on the distribution of the hidden variable is much stronger. This affects the distribution of p_{xz} under the second parametric formulation. The induced distribution of p_y is Beta with parameters $(4(1-t), 4t)$. In particular, for $t \leq 0.25$ the corresponding density is not bounded giving increasingly more probability to the event $p_y(1) > 1 - \epsilon$ as $t \rightarrow 0$. Some Beta density functions for different values of t are given in Figure 5.

If t is small then there is a strong a priori information that the inner vertex is close to being degenerate. Therefore, the Bayesian model $M_1^{(1)}$ is close to the model of independence. This should cause overestimation of the independence model. Indeed, if $t = 0.2$ then for \mathbf{u} given above we have:

$$\text{BF}_0 = 11.47, \quad \text{BF}_1 = 2.68.$$




FIGURE 5. The densities of Beta distributions with parameters $(4(1-t), 4t)$, where $t = 0.5, 0.4, 0.3, 0.2, 0.1$.

This example illustrates a more serious problem. For each t , as above, we compute $\text{BF}_0(t)$ and $\text{BF}_1(t)$ obtaining the following result.

Proposition 3.1. *For every table \mathbf{u} the Bayes Factor $\text{BF}_0(t)$ is constant and does not depend on t . Moreover, for every \mathbf{u} , $\text{BF}_1(t) \rightarrow 0$ almost surely as $t \rightarrow 0$.*

This result may seem obvious. However, what it really shows is that, even though the induced priors on the joint distribution of (X, Z) follow from the same joint prior distribution π on (X, Y, Z) , in the first case the Bayes factor may give evidence in favour of the saturated model and the second in favour of the model of independence.

The situation becomes even more dramatic if k increases. Thus if $k = 2$ we set $\beta = 4$ and $\pi_{x_{12z}}(i, j, k, l) = 1/16$ for all $i, j, k, l = 0, 1$ and again the induced prior on p_{xz} for the case $k = 0$ is uniform. Under this setting we find that $\text{BF}_2 = 2.13$ which is to be compared with (4). This uniform case is easy to generalize and illustrates another serious issue related with the choice of the parametric formulation for the model under consideration.

For general $k \geq 1$, if $\beta = 4$ and $\pi_{x_1 \dots k z}(\mathbf{i}) = 2^{-(k+2)}$ for every $\mathbf{i} = (i_0, \dots, i_{k+1}) \in \{0, 1\}^{k+2}$, then the induced prior distributions are $\theta_x \sim \text{Beta}(2, 2)$, $\theta_{1|x}(\cdot|i), \dots, \theta_{z|k}(\cdot|i) \sim \text{Beta}(1, 1)$ for $i = 0, 1$.

Proposition 3.2. *Let $\beta = 4$ and $\pi_{x_1 \dots k z}(\mathbf{i}) = 2^{-(k+2)}$ for every $\mathbf{i} \in \{0, 1\}^{k+2}$. Then $\text{BF}_k \rightarrow 0$ as $k \rightarrow \infty$.*

The proof is given in Section [A.1](#).

In the case of Proposition [3.1](#) we analysed the discrepancy of the model selection if the prior of the hidden variable becomes degenerate without affecting the prior over the observed variables. In Proposition [3.2](#) however the prior distribution of every hidden variable is symmetric and hence it shows a different aspect of the discussed problem.

4. ASYMPTOTIC ANALYSIS

Perhaps not surprisingly a similar ambiguity also arises if we use the popular BIC rather than Bayes Factor model selection method. Let q denote the true density of (X, Z) , assumed strictly positive everywhere, and let $(X_1, Z_1), \dots, (X_n, Z_n)$ be a random sample from this distribution. In this section we perform an asymptotic analysis. Let \mathcal{Z}_n denote the marginal likelihood function and S_q the entropy function of q . We define $F_n = -\log \mathcal{Z}_n$.

In the natural parametrisation the asymptotic likelihood as $n \rightarrow \infty$ is always maximised over the unique point q . By the result of [Schwarz \(1978\)](#) the asymptotic formula for the marginal likelihood can be obtained using the Laplace approximation. Whenever the prior distributions are bounded and bounded away from zero then, as $n \rightarrow \infty$,

$$(5) \quad \mathbb{E}F_n = nS_q - \frac{d}{2} \log n + O(1),$$

where $d = 2$ for the marginal likelihood given the model $M_0^{(0)}$ and $d = 3$ for model $M_1^{(0)}$. This asymptotic approximation of the marginal likelihood justifies the use of the well known BIC penalty used in routine model selection. If the true distribution lies in the null model then the entropies S_q for both models will be asymptotically equal and the difference in scores $\text{BIC}_0 - \text{BIC}_1 = \frac{1}{2} \log n$ will be always positive, which gives a positive evidence in favour of the null model.

The interpretation of BIC in the presence of hidden variables is more subtle. Under the second proposed parametrisation $k = 1$ both models are unidentifiable hence the asymptotic analysis is much harder. In particular the Laplace approximation is no longer formally valid and the appropriate asymptotic analysis must use the results of singular learning theory developed by [Watanabe \(2009\)](#).

We now compute the marginal likelihood of the data under the parametric model $M_1^{(k)}$ for $k \geq 1$. The asymptotic formula depends on the true distribution q generating the data. If $q \in M_1^{(k)} \setminus M_0^{(k)}$ for $k \geq 1$ then, despite the identifiability issue, the correct asymptotic approximations can be shown to be equal to the classical BIC formula in the case when $k = 0$. The problem occurs when $q \in M_0^{(k)}$. In this case the set of parameters mapping to q is highly singular. Whenever the prior distribution is bounded and bounded away from

zero on the whole parameter space, the asymptotic approximation of $\mathbb{E}F_n$ for the model $M_1^{(k)}$, as $n \rightarrow \infty$, is

$$(6) \quad \mathbb{E}F_n = \begin{cases} nS_q - \frac{3}{2} \log n + O(1) & \text{if } q \in M_1^{(k)} \setminus M_0^{(k)}, \\ nS_q - \frac{3}{2} \log n + k \log \log n + O(1) & \text{if } q \in M_0^{(k)}. \end{cases}$$

The proof of (6) is given in Section A.2.

The marginal likelihood of $M_0^{(k)}$ is equal to the marginal likelihood of $M_0^{(0)}$ and thus in this case

$$\mathbb{E}F_n = nS_q - \log n + O(1).$$

Now we see that a problem might occur if the true data generating distribution lies in the independence model. The bigger k , the harder it gets to distinguish $M_0^{(k)}$ from $M_1^{(k)}$. Since the entropy value will be asymptotically the same in both cases, the difference in scores between those two models is

$$(-\log n) - \left(-\frac{3}{2} \log n + k \log \log n\right) = \frac{1}{2} \log n - k \log \log n.$$

Since the true model is the model of independence we expect this difference to be highly positive like in the case when $k = 0$. However, the component of the penalty $k \log \log n$ distorts this as is shown in Figure 6. The score difference remains negative whenever $k \geq 2$ even for very large n . Hence, for all usual values of the sample size n the use of the standard BIC criterion underestimates the evidence of the null model whenever $k \geq 1$.

5. DISCUSSION AND CONCLUSIONS

In this paper we have shown that we need a foundation for justifying a particular embedding before BF or BIC model selection is unambiguous for binary tree models. This fact can be shown also to apply to all model selection over discrete graphical models, Gaussian graphical models with potential hidden variables and, in particular, Bayesian hierarchical models, where systematically hidden variables are routinely added to the system to articulate various types of dependence structures. Problems are particularly acute when the dimension of the embedding class is itself contextually ambiguous as in our running example. For all these models whenever the appropriate embedding is not transparent, we should be aware this choice could be critical to the result of our selection. In particular the current praxis of paying little attention to the different inferential implications of a chosen embedding, focusing instead on the numerical efficacy of a particular representation is one that should be of great concern to Bayesian modelers. We note that the problems we identify here do not concern just BF or BIC: other Bayesian model selection methods also suffer the same difficulty.

What can we do to address this issue? First, if there exists an associated meaning to a given embedding then we could elicit a prior distribution for each model and then average over this. In our example we could for example elicit the number of potential differently evolved ancestors for which only one direct descendant survived. However, as in our example, it may well be difficult to unambiguously make this association and is certainly




FIGURE 6. The difference in scores $\frac{1}{2} \log n - k \log \log n$ depicted as a function of n for different values of $k = 1, 2, 3, 4, 5$. The n axis is in the log scale.

not in the spirit of current model selection, which tries for the sake of "objectivity" to avoid the incorporation of as little domain knowledge as possible: including much more direct domain knowledge than this. A second possible direction is to systematically check the plausibility of a particular embedding on the associated marginal likelihood of different models to check how the system will learn in various contingencies and calibrate to that.

Finally we could question, as some others do e.g. Dennis Lindley, [Draper \(1999\)](#) whether model selection is actually compatible with Bayesian methodology at all. It would be sad however to discard these techniques which have undoubtedly provided such useful output to scientists endeavouring to understand the processes underpinning their observations. But at least when using Bayesian model selection techniques in conjunction with apparently innocuous homogeneity assumptions like, in the case of BN model selection, parameter modularity we should be aware that the associated inferences could seriously mislead the investigator.

A.1. Proof of Proposition 3.2. Since the marginal likelihood $L_0^{(k)}$ for the null model does not change with k , it suffices to show that the marginal likelihood $L_1^{(k)}$ of the model $M_1^{(k)}$ converges to zero as $k \rightarrow \infty$. Let $\boldsymbol{\theta} = (\theta_x, \theta_{1|x}, \theta_{2|1}, \dots, \theta_{z|k}) \in [0, 1]^{1+2(k+1)}$ and

$$f_k(\boldsymbol{\theta}) = \prod_{i,j=0}^1 \left(\theta_x(i) \sum_{i_1, \dots, i_k} \theta_{1|x}(i_1|i) \theta_{2|1}(i_2|i_1) \cdots \theta_{k|k-1}(i_k|i_{k-1}) \theta_{z|k}(j|i_k) \right)^{u_{ij}}.$$

Since $\theta_x \sim \text{Beta}(2, 2)$, $\theta_{1|x}(\cdot|i), \dots, \theta_{z|k}(\cdot|i) \sim \text{Beta}(1, 1)$ for $i = 0, 1$, we have

$$(7) \quad L_1^{(k)} = \frac{\Gamma(4)}{\Gamma(2)^2} \int_{[0,1]^{1+2(k+1)}} f_k(\boldsymbol{\theta}) \prod_{i=0}^1 \theta_x(i) d\boldsymbol{\theta}.$$

Let $t_{i_1, \dots, i_k}^i = \theta_{1|x}(i_1|i) \theta_{2|1}(i_2|i_1) \cdots \theta_{k|k-1}(i_k|i_{k-1})$ for $i, i_1, \dots, i_k = 0, 1$. Note that $\varphi(x) = x^a$ is a convex function on $[0, \infty)$ whenever $a \geq 1$ or $a = 0$. Since $\sum_{i_1, \dots, i_k} t_{i_1, \dots, i_k}^i = 1$ for $i = 0, 1$ and either $u_{ij} \geq 1$ or $u_{ij} = 0$, by Jensen's inequality

$$\left(\theta_x(i) \sum_{i_1, \dots, i_k} t_{i_1, \dots, i_k}^i \theta_{z|k}(j|i_k) \right)^{u_{ij}} \leq \sum_{i_1, \dots, i_k} t_{i_1, \dots, i_k}^i (\theta_x(i) \theta_{z|k}(j|i_k))^{u_{ij}}$$

and hence

$$(8) \quad f_k(\boldsymbol{\theta}) \leq \sum_{i_1, \dots, i_k} \prod_{i,j=0}^1 \theta_{1|x}(i_1|i) \theta_{2|1}(i_2|i_1) \cdots \theta_{k|k-1}(i_k|i_{k-1}) (\theta_x(i) \theta_{z|k}(j|i_k))^{u_{ij}}.$$

We have

$$\int_{[0,1]^2} \prod_{i,j=0}^1 \theta_{1|x}(i_1|i) d\theta_{1|x} = \int_{[0,1]^2} \theta_{1|x}(i_1|0)^2 \theta_{1|x}(i_1|1)^2 d\theta_{1|x}(1|1) d\theta_{1|x}(1|0) = \left(\frac{\Gamma(3)}{\Gamma(4)} \right)^2$$

and for $l = 2, \dots, k$

$$\int_{[0,1]} \prod_{i,j=0}^1 \theta_{l|l-1}(i_l|i_{l-1}) d\theta_{l|l-1}(1|i_{l-1}) = \int_{[0,1]} \theta_{l|l-1}(i_l|i_{l-1})^4 d\theta_{l|l-1}(1|i_{l-1}) = \left(\frac{\Gamma(5)}{\Gamma(6)} \right)^2.$$

This together with (8) gives

$$L_1^{(k)} \leq C(\mathbf{u}) \sum_{i_1, \dots, i_k} \left(\frac{\Gamma(5)}{\Gamma(6)} \right)^{k-1} \left(\frac{\Gamma(3)}{\Gamma(4)} \right)^2 = C(\mathbf{u}) \frac{5}{9} \left(\frac{2}{5} \right)^k,$$

where $C(\mathbf{u})$ is a constant which depends only on the table of counts \mathbf{u} . This in particular implies that $L_1^{(k)} \rightarrow 0$ as $k \rightarrow \infty$.

A.2. Proof of Equation (6). The set of parameters of the model is the marginal distribution p_x of X , the conditional distribution of Y_1 given $X = i$ denoted by $\theta_1(\cdot|i)$, the conditional distribution of Z given $Y_k = i$, denoted by $\theta_z(\cdot|i)$ and conditional distributions of Y_j given $Y_{j-1} = i$ for every $j = 2, \dots, k$ denoted by $\theta_j(\cdot|i)$. Hence the parameter vector θ lies in $[0, 1]^{1+2(k+1)}$. We have

$$p_{xz}(i, k) = p_x(i) \sum_{j_1, \dots, j_k} \theta_1(j_1|i) \theta_2(j_2|j_1) \cdots \theta_k(j_k|j_{k-1}) \theta_z(k|j_k).$$

Since the true distribution is assumed to lie in the independence model we cannot use the standard Laplace approximation for the marginal likelihood because the asymptotic likelihood is maximised over a singular subset of the parameter space. Assume that the prior distribution on Θ is bounded and bounded away from zero. Then by The Corollary 6.1 of [Watanabe \(2009\)](#) the marginal likelihood is asymptotically, as $n \rightarrow \infty$, approximated by

$$nS - \lambda \log n + (m - 1) \log \log n + O(1),$$

where λ and m are the smallest pole and its multiplicity of an analytic function given by

$$(9) \quad \xi(w) = \int_{\Theta} (f(\theta))^{-w} d\theta,$$

where $f(\theta)$ is the Kullback-Leibler divergence from the true model q . By Theorem 1.2 of [Lin \(2011\)](#) we can also replace $f(\theta)$ with the sum of squares $\sum_{i,j=0}^1 (p_{ij}(\theta) - q_{ij})^2$.

Let $\mu_x(\theta) = p_{10}(\theta) + p_{11}(\theta)$, $\mu_z(\theta) = p_{01}(\theta) + p_{11}(\theta)$ and

$$\mu_{xz}(\theta) = p_{11}(\theta) - (p_{10}(\theta) + p_{11}(\theta))(p_{01}(\theta) + p_{11}(\theta)).$$

It is immediate to see that μ_x, μ_z, μ_{xz} are in one-to-one correspondence with $[p_{ij}]$. Moreover, μ_{xz} is just the covariance between X and Z and hence it is zero if and only if $X \perp\!\!\!\perp Z$. Since the pole of (9) and its multiplicity do not change under isomorphisms, we can further replace the function $\sum_{i,j} (p_{ij}(\theta) - q_{ij})^2$ with

$$(10) \quad (\mu_x(\theta) - \mu_x^*)^2 + (\mu_z(\theta) - \mu_z^*)^2 + (\mu_{xz}(\theta) - \mu_{xz}^*)^2.$$

The asterisks refer to the moments of the true distribution q . Note that since q lies in the independence model then $\mu_{xz}^* = 0$.

Now make an isomorphic change of coordinates of θ to parameters

$$\omega = (\mu_x, \mu_{y_1}, \dots, \mu_{y_k}, \mu_z; \eta_{xy_1}, \eta_{y_1 y_2}, \dots, \eta_{y_k z})$$

where μ_{y_i} is the mean of Y_i and $\eta_{xy} = \mathbb{P}(Y = 1|X = 1) - \mathbb{P}(Y = 1|X = 0)$ is the linear regression coefficient of Y with respect to X . Since

$$\mu_{xz}(\omega) = \frac{1}{4} (1 - \mu_x^2) \eta_{xy_1} \cdots \eta_{y_k z}$$

then the function in (10) expressed in new parameters becomes

$$(\mu_x - \mu_x^*)^2 + (\mu_z - \mu_z^*)^2 + \left(\frac{1}{4} (1 - \mu_x^2) \eta_{xy_1} \cdots \eta_{y_k z}\right)^2.$$

By Remark 7.2 of [Watanabe \(2009\)](#), if ω^* is an interior point of the parameter space, $\lambda = 1 + \lambda'$ and $m = m'$ where (λ', m') are the smallest pole and its multiplicity of the analytic function of w given by

$$\int_{[-\epsilon, \epsilon]^{k+1}} (\eta_{xy_1} \cdots \eta_{y_k z})^{-2w} d\eta_{xy_1} \cdots d\eta_{y_k z},$$

for a sufficiently small $\epsilon > 0$. Finally, this integral is equal to $C(\epsilon) \left(\frac{1}{1-2w}\right)^{k+1}$, where $C(\epsilon)$ is a constant which depends only on ϵ . It follows that the pole of this function is $\lambda' = 1/2$ and the multiplicity of this pole is $m' = k + 1$ and hence $\lambda = 3/2$ and $m = k + 1$.

Note that to use Remark 7.2 of [Watanabe \(2009\)](#) we assume that the parameter space has locally a product structure where μ_x and μ_z are independent of other parameters. This holds only in the interior of the parameter space so the boundary points ω^* need to be checked separately. We omit the details.

ACKNOWLEDGEMENTS

We would like to thank Shaowei Lin for help with setting up the exact computations in Section 3.

REFERENCES

- BERNARDO, J.-M. & SMITH, A. F. M. (1994). *Bayesian theory*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. Chichester: John Wiley & Sons Ltd.
- CONSONNI, G. & LA ROCCA, L. (2011). *Bayesian Statistics 9*, chap. On moment priors for Bayesian model choice with applications to directed acyclic graphs. Oxford University Press.
- DAWID, A. P. & LAURITZEN, S. L. (1993). Hyper-Markov laws in the statistical analysis of decomposable graphical models. *Ann. Statist.* **21**, 1272–1317.
- DENISON, D., HOLMES, C., MALLICK, B. & SMITH, A. (2002). *Bayesian methods for nonlinear classification and regression*. John Wiley & Sons Inc.
- DRAPER, D. (1999). [bayesian model averaging: A tutorial]: Comment. *Statistical Science* **14**, 405–409.
- GILULA, Z. (1979). Singular value decomposition of probability matrices: Probabilistic aspects of latent dichotomous variables. *Biometrika* **66**, 339–344.
- HECKERMAN, D. (1999). A tutorial on learning with Bayesian networks. In *Learning in Graphical Models. Adaptive computation and machine learning.*, M. I. Jordan, ed. MIT Press.
- HECKERMAN, D., GEIGER, D. & CHICKERING, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning* **20**, 197–243.
- JOHNSON, V. E. & ROSSELL, D. (2010). On the use of non-local prior densities in Bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72**, 143–170.

- KEY, J. T., PERICCHI, L. R. & SMITH, A. F. M. (1999). Bayesian model choice: what and why? In *Bayesian statistics, 6 (Alcoceber, 1998)*. New York: Oxford Univ. Press, pp. 343–370.
- LAURITZEN, S. L. (1996). *Graphical models*, vol. 17 of *Oxford Statistical Science Series*. New York: The Clarendon Press Oxford University Press. Oxford Science Publications.
- LIN, S. (2011). Asymptotic Approximation of Marginal Likelihood Integrals. [arXiv:1003.5338](#). Submitted.
- LIN, S., STURMFELS, B. & XU, Z. (2009). Marginal likelihood integrals for mixtures of independence models. *The Journal of Machine Learning Research* **10**, 1611–1631.
- MADSEN, A., LANG, M., KJÆRULFF, U. & JENSEN, F. (2003). The hugin tool for learning bayesian networks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, T. Nielsen & N. Zhang, eds., vol. 2711 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, pp. 594–605.
- NASH, JR., J. F. (1950). The bargaining problem. *Econometrica* **18**, 155–162.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics* **6**, 461–464.
- SEMPLE, C. & STEEL, M. (2003). *Phylogenetics*, vol. 24 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford: Oxford University Press.
- SETTIMI, R. & SMITH, J. Q. (2000). Geometry, moments and conditional independence trees with hidden variables. *Ann. Statist.* **28**, 1179–1205.
- THOMAS, L. C. (1984). *Games, theory and applications*. Ellis Horwood Series: Mathematics and its Applications. Chichester: Ellis Horwood Ltd.
- WATANABE, S. (2009). *Algebraic Geometry and Statistical Learning Theory*. No. 25 in Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press. ISBN-13: 9780521864671.
- YANG, Z. (2006). *Computational Molecular Evolution (Oxford Series in Ecology and Evolution)*. Oxford, USA: Oxford University Press, USA.

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, TU EINDHOVEN, PO Box 513, 5600 MB, EINDHOVEN, THE NETHERLANDS
E-mail address: piotr.zwiernik@gmail.com

DEPARTMENT OF STATISTICS, UNIVERSITY OF WARWICK, COVENTRY, CV4 7AL, UK
E-mail address: j.q.smith@warwick.ac.uk