

University of Puerto Rico
Rio Piedras Campus
Faculty of Natural Sciences
Department of Mathematics

**A UNIVERSAL AND ROBUST BOUND FOR THE
INTRINSIC BAYES FACTORS**

By

Richard Anthony Clare Morales

A DISSERTATION SUBMITTED IN PARTIAL FULFILMENT OF THE
REQUIREMENTS FOR THE DEGREE OF DOCTOR OF
PHILOSOPHY IN MATHEMATICS AT THE UNIVERSITY OF
PUERTO RICO, RIO PIEDRAS CAMPUS

May, 2024

Abstract of Ph.D Thesis Presented to the Graduate School
of the University of Puerto Rico, Rio Piedras Campus in Partial Fulfilment of the
Requirements for the Degree of Doctor of Philosophy in Mathematics

**A UNIVERSAL AND ROBUST BOUND FOR THE INTRINSIC BAYES
FACTORS**

By

Richard Anthony Clare Morales

May 2024

Chair: Luis Raúl Pericchi Guerra Ph.D
María E. Pérez Hernández Ph.D
Major Department: Mathematics

In this work, we undertake a comprehensive reformulation, modification, and extension of Smith & Spiegelhalter's [14] and [15] Bayes Factor work within the evolving subject of Objective Bayes Factors. Our primary focus centers on defining and computing empirical and theoretical bounds for the Intrinsic Bayes Factor (IBF) across various models, including normal, exponential, Poisson, geometric, linear, and ANOVA. We show that our new bounds are useful, feasible, and change with the amount of information. We also propose a methodology to construct the least favorable (for the null model) intrinsic priors that result in the lower and upper bounds of the Intrinsic Bayes Factors under certain conditions. Notably, our lower bounds exhibit superior performance compared to the well-known $-ep \log(p)$ bound proposed by Sellke et al. (2001) [13] based on p-values.

Copyright © 2024

by

Richard Anthony Clare Morales

To my dearest mother, Maria Morales, whose unwavering support and encouragement guided every step of my life. Richard Clare, my father, your belief in my capabilities has been a constant source of inspiration. Marta Ortiz, my beloved wife, your boundless love and patience carried me through every challenge. To my sons, Alaya and Richard Clare, your existence provided joy and purpose in my life.

ACKNOWLEDGMENTS

I hereby express my sincere gratitude to my advisor Luis Pericchi, for his invaluable guidance, support, and insights throughout my research work. His suggestions and supervision during this process were of paramount importance and greatly contributed to the success of this thesis. His mentorship has not only helped me to overcome numerous challenges but also broadened my horizons and enriched my knowledge. It has been an honor and privilege to work with such a knowledgeable, supportive, and encouraging advisor. I am forever grateful for his assistance and dedication to advancing my academic career. Luis Pericchi acknowledges illuminating conversations with D. Spiegelhalter, AFM Smith, and J.O. Berger over many years.

TABLE OF CONTENTS

	<u>page</u>
ABSTRACT	ii
ACKNOWLEDGMENTS	v
LIST OF FIGURES	viii
LIST OF ABBREVIATIONS AND SYMBOLS	ix
1 INTRODUCTION	1
2 LITERATURE REVIEW	3
2.1 Bayes factors	3
2.2 Prior and Posterior Model Probabilities	4
2.3 Jeffreys’s objective priors	5
2.4 Objective Bayes Factors	5
2.5 Minimal training samples	6
2.5.1 A formal definition for the set of all minimal training samples	6
2.5.2 An algorithm to generate the set of all minimal training samples	7
2.6 The Intrinsic Bayes Factor	8
2.7 The Expected Posterior Bayes Factors	10
3 A UNIVERSAL ROBUST BOUND	12
3.1 The Smith and Spiegelhalter’s method	12
3.2 The universal robust bound for the Intrinsic Bayes Factors	13
3.3 A bridge between the universal bound and the IBF	14
3.3.1 A comparison with the p-value based lower bound	15
4 The SP Bayes Factors	20
4.1 The SP Prior	21
4.2 The SP Bayes Factor	21
4.3 Results	22
5 Hypothesis Testing with the Normal Distribution	25
5.1 The Normal Scale Hypothesis Test	25
5.1.1 The universal lower bound	26
5.1.2 Remarks	28
5.1.3 The SP and Intrinsic Priors	29

5.2	The Normal Mean Hypothesis Test	32
5.2.1	The Expected Intrinsic Bayes Factor for the Normal Mean Hypothesis Test	33
5.2.2	The Intrinsic Bayes Factor Bounds for the Normal Mean Hypothesis Test	34
5.2.3	Normal mean hypothesis test with unknown variance.	35
6	The Exponential Model Hypothesis Test	38
6.1	The Intrinsic Bayes Factor Dynamic and Universal Lower Bounds	39
6.2	The SP Bayes Factors	40
6.3	The EP Bayes Factors	43
6.4	The -eplogp lower bound	49
6.5	Conclusions	52
7	Linear regression and ANOVA Hypothesis Testing with Bayes Factors	53
7.1	Linear regression	53
7.2	Two nested normal-linear models	55
7.3	The One-way Layout (ANOVA)	59
7.3.1	The Full Jeffrey's (no independence) $q_k = p_k$, for $k = i, j$ for ANOVA	62
7.3.2	The Reference prior ($q_k = 0, k = i, j$) for ANOVA	63
7.3.3	The Modified Jeffrey's Prior ($q_j = p_j - p_i, q_i = 0$) for ANOVA	63
7.4	Conclusions	64
8	Separated Hypothesis Testing	65
8.1	The Poisson vs. Geometric Separated Hypothesis Test	66
8.2	The Negative Binomial vs. Poisson Separated Hypothesis Test	79
8.3	Conclusions	83
9	Ghosh and Samanta Alternative Approach	84
9.1	The Normal Mean Hypothesis Test Example	84
9.2	The Exponential Hypothesis Test Example	85
9.2.1	Conclusions	86
10	Research Contributions	87

LIST OF FIGURES

<u>Figure</u>	<u>page</u>
5-1 A comparison between the SP Prior and the Intrinsic Prior for the Normal Scale Hypothesis test for different null hypothesis values.	32
6-1 \bar{B}_{10}^{SS} vs B_{10}^{SP} under H_0 for increasing sample size and exponential data. . . .	42
6-2 Simulations with different number of samples for the exponential SP and EP Priors for $\lambda_0 = 1$	44
6-3 Logarithm of the SP and EP Bayes factor using one-hundred exponential simulations (100 samples each) for $\lambda_0 = 1, H_0$ True	47
6-4 Logarithm of the SP and EP Bayes factor using one-hundred exponential simulations (100 samples each) for $\lambda_0 = 1, H_0$ False	48
6-5 The $-ep \log(p)$ and \underline{B}_{01}^I using 100 exponential simulations (100 samples each) for $\lambda_0 = 1$ (Null hypothesis true), and 100 exponential simulations with 10 samples each for $\lambda_0 = 1$ and the simulations rate parameter $r = 5$ (false null hypothesis)	51
8-1 Poisson vs. Geometric IBF 10 bound as the sample size increases when H_0 is true (left) and when H_0 is false (right).	70
8-2 (a) Empirical and Theoretical \bar{B}_{10}^I when H_0 is true (b) Empirical and Theoretical \bar{B}_{01}^I when H_0 is false	72
8-3 $AIBF_{10}$ and The IBF_{10} Lower Bound Comparison for the Poisson vs Geometric Test	75
8-4 Cox (1962) p. 414	76
8-5 We used the randomized Cox data considered a growing vector with those samples and calculated the Bayes Factor as we added more data.	78
8-6 Poisson vs. Negative Binomial IBF lower bound as the sample size increases with samples from H_0	82

LIST OF ABBREVIATIONS AND SYMBOLS

ε	Random error vector
$\bar{\mathbf{y}}$	Mean of the random vector
\mathbf{A}	Design matrix
\mathbf{A}^t	transpose of \mathbf{X}
Y	Random variable
$\log_e(p) = \log(p)$	Logarithm in base e
\mathbf{R}	Statistical software
\mathbf{Y}	Random vector
\mathbf{y}	Vector of observations
\mathbf{I}	Fisher Information Matrix
$\pi(\theta)$	Prior distribution for parameters vector θ
$\pi(\theta \mathbf{y})$	Posterior distribution for parameters vector θ given the observations \mathbf{x}
$\pi^N(\theta)$	Objective Non-informative prior distribution for parameters vector θ
$f(\mathbf{y} \theta)$	Density of \mathbf{y} given θ
$F(\mathbf{y} \theta)$	Distribution of \mathbf{y} given θ
$\frac{\partial f}{\partial \theta}$	Partial derivative of function f with respect to θ
m_i	Marginal for model i
B_{ij}	Bayes factor for Model i vs Model j
B_{ij}^N	Objective (non-informative) Bayes factor for Model i vs Model j
\mathbf{I}_n	$n \times n$ identity matrix
θ	Vector of parameters

CHAPTER 1

INTRODUCTION

The misuse of classical hypothesis testing methods, specifically p-values, has been the subject of substantial critique in the statistical literature (Wasserstein and Lazar, 2016 [17]; Wasserstein et al., 2019 [16]). Despite these warnings, the practice of p-value null hypothesis significance testing (NHST) continues to be a predominant technique for model selection. However, alternatives do exist. Among them, the Bayesian testing methodology developed by Sir Harold Jeffreys (1961) [7] stands as one of the most widely accepted and used. Jeffreys's methodology, premised on the concept that the null hypothesis carries a positive point mass, quantifies the support for or against a model based on the data. Additionally, it offers practical applications through derived approximations for complex or analytically infeasible solutions. A significant contribution of Jeffreys's framework is the Bayes factor, a measure of support for one model over another, irrespective of the absolute correctness of these models. Its effectiveness spans various inference desiderata, including interpretability, adherence to Occam's razor, and consistency in large samples. Kass and Raftery (1995) [8] offer an extensive introduction to Bayes factors. Our research primarily concentrates on the pivotal role of training samples within statistical methodologies. These methodologies include classification and discrimination, robustness, model selection, and cross-validation. Recent advancements in Bayesian model selection, namely the intrinsic Bayes factor by Berger and Pericchi (1996a) [3] and the expected posterior prior by Perez and Berger (2002) [10], have employed training samples to transition improper objective priors into more suitable distributions for model selection.

A challenge often faced is the indeterminacy of classical Bayes Factors, due to the dependence on an undefined ratio of constants derived from the improper priors. To circumnavigate this, Jeffrey suggested conventional proper priors for the extra-parameters under the larger hypothesis, albeit improper for common parameters. Various techniques have been proposed to address this issue, including those by Smith and Spiegelhalter (1980) [14] and Spiegelhalter and Smith (1982) [15]. Though approximate, these techniques facilitate sensible scaling for Bayes Factors and bear a direct relationship with more recent approaches like the Intrinsic Bayes Factor, the Intrinsic Priors, and EP-priors. To elucidate these connections, we will re-conceptualize Smith and Spiegelhalter's approach and the associated Bayes Factor, hereby referred to as SSBF's (B^{SS}), drawing parallels to other methodologies.

Chapter three specifically focuses on the quasi-Bayes factor, which emerges as a universal lower bound for the Intrinsic Bayes factor. In Chapter four we introduce the concept of SP Priors (Superior Posterior priors) that either generate the Intrinsic Bayes Factors bound under specific assumptions or serve as an asymptotic approximation of the bounds. These priors are the least favorable priors for the null hypothesis and through chapter three to chapter eight we performed numerical simulations and verified the consistency of the SP, the empirical and the theoretical Intrinsic Bayes Factor bounds across various essential distributions such as the normal, exponential, Poisson, geometric, and negative binomial. Furthermore, we calculated the bounds for two nested normal-linear models and ANOVA. We discussed model selection problems across different distributions and hypothesis tests such as nested and separated hypotheses, and for those cases we provided closed forms for the Intrinsic Bayes factors Bounds, the SP Priors, and the SP Bayes Factors.

CHAPTER 2

LITERATURE REVIEW

2.1 Bayes factors

Traditional point estimation methods, such as the maximum likelihood estimator (MLE) or the method of moments (MoM), typically involve the utilization of specific estimates for each model parameter based on the available data. In contrast, Bayesian inference takes a different approach by assigning a probability distribution to each model parameter.

Let $\mathbf{y} = (y_1, \dots, y_n)$ represent a set of n random samples obtained from a random variable Y with a density $f_i(y|\theta_i)$, where θ_i denotes a vector of parameters associated with model M_i . Employing the Bayes' Theorem, the posterior distribution for model M_i can be expressed as:

$$\pi(\theta_i|\mathbf{y}) = \frac{f_i(\mathbf{y}|\theta_i)\pi_i(\theta_i)}{m_i(\mathbf{y})} = \frac{f_i(\mathbf{y}|\theta_i)\pi_i(\theta_i)}{\int f_i(\mathbf{y}|\theta_i)\pi_i(\theta_i)d\theta_i}$$

where $m_i(\mathbf{y})$ is called the marginal, evidence, or global likelihood, and $\pi_i(\theta)$ is the prior distribution of the parameters for M_i . Suppose that we are comparing two models for the data \mathbf{y}

$$M_k : \mathbf{Y} \text{ has a density } f_k(\mathbf{y}|\theta_k) \quad k = i, j$$

where θ_k are unknown model parameters associated with model M_k . Suppose we have available prior distributions $\pi_k(\theta_k)$ for the unknown parameters. Define the marginal or predictive densities of \mathbf{y} as

$$m_i(\mathbf{y}) = \int f_i(\mathbf{y}|\theta_i)\pi_i(\theta_i)d\theta_i.$$

The *Bayes factor* of M_i to M_j is given by

$$B_{ij} = \frac{m_i(\mathbf{y})}{m_j(\mathbf{y})} = \frac{\int f_i(\mathbf{y}|\theta_i)\pi_i(\theta_i)d\theta_i}{\int f_j(\mathbf{y}|\theta_j)\pi_j(\theta_j)d\theta_j}.$$

The *Bayes factor* is often interpreted as the evidence provided by the data for the model M_i vs. the alternative model M_j , but the *Bayes factor* depends also on the *priors*. Alternatively, B_{ij} can be interpreted as the weighted likelihood ratio of M_i to M_j . with the priors being the "weighting functions."

2.2 Prior and Posterior Model Probabilities

Suppose that prior probabilities $\pi(M_i), i = 1, \dots, n$, of the n models, are available, then we can obtain for each model M_i the posterior probability from the Bayes factor given the data \mathbf{y} by calculating

$$P(M_i|\mathbf{y}) = \frac{\pi(M_i)m_i(\mathbf{y})}{\sum_{k=1}^n \pi(M_k)m_k(\mathbf{y})}$$

This equation can also be written as

$$P(M_i|\mathbf{y}) = \left[\sum_{k=1}^n \frac{\pi(M_k)}{\pi(M_i)} B_{ki}(\mathbf{y}) \right]^{-1}$$

where $O_{ki} = \pi(M_k)/\pi(M_i)$ is known as the prior odds of the model M_k over the model M_i , and the posterior probabilities odds are given by

$$\frac{P(M_i|\mathbf{y})}{P(M_j|\mathbf{y})} = \frac{\pi(M_i)m_i(\mathbf{y})}{\pi(M_j)m_j(\mathbf{y})} = O_{ij}B_{ij}$$

Hence, the posterior odds are given the product of the prior odds and the Bayes factor. A common unbiased choice of the prior model probabilities is given by $\pi(M_i) = 1/n$ where each model has the same prior probability. Note that using those uniform priors we can

obtain the re-normalized marginal distributions from the posterior model probabilities as

$$\bar{m}_i(\mathbf{y}) = \frac{m_i(\mathbf{y})}{\sum_{k=1}^n m_k(\mathbf{y})}$$

and also the Bayes Factor can be in terms of the re-normalized marginals as

$$B_{ij}(\mathbf{y}) = \frac{\bar{m}_i(\mathbf{y})}{\bar{m}_j(\mathbf{y})}$$

2.3 Jeffreys's objective priors

Consider the Fisher information matrix when there are n parameters, so that θ is a $n \times 1$ vector $\theta = [\theta_1, \theta_2, \dots, \theta_n]^T$, then the Fisher information takes the form of an $n \times n$ matrix with typical element

$$I_{ij}(\theta) = E \left[\left(\frac{\partial}{\partial \theta_i} \log f(\mathbf{y}|\theta) \right) \left(\frac{\partial}{\partial \theta_j} \log f(\mathbf{y}|\theta) \right) \middle| \theta \right].$$

Jeffreys prior is a non-informative (objective) prior distribution for a parameter space; it is proportional to the square root of the determinant of the Fisher information matrix

$$\pi^N(\theta) \propto \sqrt{\det I(\theta)}$$

and this prior is invariant under reparameterization of the parameter vector θ and this property makes it of special interest for use with location and scale parameters.

2.4 Objective Bayes Factors

The unscaled Bayes factors are obtained using the non-informative objective priors $\pi_i^N(\theta_i)$ and $\pi_j^N(\theta_j)$ for the models M_i and M_j respectively and can be calculated as

$$B_{ij}^N(\mathbf{y}) = \frac{\int f_i(\mathbf{y}|\theta_i) \pi_i^N(\theta_i) d\theta_i}{\int f_j(\mathbf{y}|\theta_j) \pi_j^N(\theta_j) d\theta_j}$$

Non-informative priors can be used for one-sided hypothesis testing but unfortunately, when we are using improper priors for point null hypotheses the Bayes Factor depends on an undetermined ratio of constants C_i/C_j , which comes from the improper

priors. Some techniques are widely used to overcome this issue which can approximate the desired ratio of constants and those will be studied and extended in the next sections.

2.5 Minimal training samples

There are different approaches to calculate Bayes Factors based on the concept of a proper 'minimal training sample' which is a subset of the entire data \mathbf{y} .

Definition: A training sample $\mathbf{y}(\ell) \subset \mathbf{y}$, is called proper if $0 < m(\mathbf{y}(\ell)) < \infty$, and minimal if it is proper and no subset is proper.

2.5.1 A formal definition for the set of all minimal training samples

The formal definition and formula for the set of all possible minimal training samples can be obtained using the definition above and extending the concept. Let $y(\ell_i) = y(i_1, i_2, \dots, i_k) = \{y_{i_1}, y_{i_2}, \dots, y_{i_k}\}$ be any minimal training sample with $k \leq n$. The set $\{y(\ell_1), \dots, y(\ell_L)\}$ of all possible minimal training samples (non-repeated subsets of size k) can be defined as the set

$$D_{n,k} = \left\{ \left\{ \left\{ \dots \{y(i_1, i_2, \dots, i_k)\}_{i_k=i_{(k-1)}}^{i_n} \dots \right\}_{i_2=i_1}^{i_n} \right\}_{i_1=1}^{i_n} : 1 \leq i_1 < i_2 < i_3 < \dots < i_k \leq n \right\}$$

Hence,

$$|D_{n,k}| = \binom{n}{k} = L \quad \text{and} \quad D_{\infty,k} = \lim_{n \rightarrow \infty} D_{n,k}$$

Therefore,

$$|D_{\infty,k}| = \lim_{n \rightarrow \infty} |D_{n,k}| = \lim_{n \rightarrow \infty} \binom{n}{k} = \lim_{n \rightarrow \infty} \frac{(n)(n-1)(n-(k+1))}{k!} = \infty$$

For ease of notation, we will denote the set $D_{\infty,k}$ as D throughout this thesis.

The determination of the set $D_{n,k}$, encompassing all minimal training samples of size k from a sample of size n , will be facilitated by an original, exhaustive, albeit sub-optimal algorithm. This algorithm yields the computation of Bayes factors involving these terms?such as the Arithmetic Intrinsic Bayes Factor, Median Intrinsic Bayes Factor,

EP-Bayes factors, and SS-Bayes factors?which will be expounded upon in subsequent chapters.

Similarly, considering $L = \binom{n}{k}$, we can derive the following computationally intuitive formula applicable to the Arithmetic Intrinsic Bayes Factor (AIBF)

$$\frac{1}{L} \sum_{i_1=1}^{i_n} \sum_{i_2=i_1}^{i_n} \cdots \sum_{i_k=i_{(k-1)}}^{i_n} \mathbb{1}(1 \leq i_1 < i_2 < \cdots < i_k \leq n) B_{01}(y(i_1, \dots, i_k)) = \frac{1}{L} \sum_{d \in D_{n,k}} B_{01}(d).$$

In the next section, we will present the general definition algorithm that can receive a vector of observations, and will generate the set of all possible minimal training samples.

2.5.2 An algorithm to generate the set of all minimal training samples

We present here an original algorithm devised for the purpose of generating the set of all minimal training samples. It is important to note that the algorithm proposed in this study may not be suitable for processing very large sample sizes due to its computational complexity, which amounts to approximately $O(n^k)$ operations. For future investigations, it is recommended to explore more efficient algorithms for generating the matrix \mathbf{M} based on the algorithm discussed herein.

Alternatively, a generalized version of this algorithm can be constructed using a recurrence relation instead of nested for loops. Such an approach would enable the development of a versatile algorithm applicable to any given values of n and k . In cases where the sample sizes are substantial, it is advisable to employ randomization techniques to generate a representative subset of minimal training samples, which can yield satisfactory approximations of the Intrinsic Bayes Factors (IBFs).

Algorithm 1 An algorithm to produce the set of all possible minimal training samples

Input: $n, k \in \mathbb{N}, n \geq k$

Output: A matrix with k columns and $\binom{n}{k}$ rows with the indexes of all of the possible minimal training samples of size k from a sample of size n .

function MTS((n, k))

Initialize an empty matrix M

for i_1 in $1 : n$ **do**

for i_2 in $i_1 : n$ **do**

\vdots

for i_k in $i_{k-1} : n$ **do**

if $1 \leq i_1 < i_2 < \dots < i_k \leq n$ **then**

 Add the row with values (i_1, \dots, i_k) to the bottom of the matrix M .

end if

end for

\vdots

end for

end for

Return M

end function

2.6 The Intrinsic Bayes Factor

Suppose that non-informative (usually improper) priors $\pi_i^N(\theta_i), i = 1, \dots, n$ are available for the models M_1, \dots, M_n respectively. The general recommendation is that we choose those priors as "reference priors" (Berger and Bernardo 1992 [2]), but there are many other choices and approximations with good results and advantages over the reference priors in terms of computational and mathematical complexity.

The *Intrinsic Bayes Factor* (IBF) is calculated using a minimal training sample $\mathbf{y}(\ell)$ which converts the improper prior $\pi_i^N(\theta_i)$ to proper posteriors $\pi_i^N(\theta_i|\mathbf{y}(\ell))$. The Bayes factor is calculated using this proper posterior as the prior distribution of the remaining of the data $\mathbf{y}(-\ell) = \mathbf{y} \setminus \mathbf{y}(\ell)$ as follows

$$B_{ij}^I = B_{ij}(\ell) = \frac{\int f_i(\mathbf{y}(-\ell)|\theta_i, \mathbf{y}(\ell))\pi_i(\theta_i|\mathbf{y}(\ell))d\theta_i}{\int f_j(\mathbf{y}(-\ell)|\theta_j, \mathbf{y}(\ell))\pi_j(\theta_j|\mathbf{y}(\ell))d\theta_j}$$

where $B_{ij}(\ell) = B_{ij}(\mathbf{y}(-\ell)|\mathbf{y}(\ell))$. There are $L = \binom{n}{k}$ different minimal training samples for n i.i.d samples, where k is the number of parameters of the density under consideration. Denote by $B_{ij}^N(\mathbf{y})$, the Bayes Factor with the whole sample \mathbf{y} and un-scaled objective (Non-Informative) priors $\pi_k^N(\boldsymbol{\theta}_k), k = i, j$, and, $B_{ij}(\mathbf{y}(-\ell)|\mathbf{y}(\ell))$, the "trained" Bayes Factor after the training sample has been employed to get a proper scaling of the Bayes Factor. A fundamental Lemma with Bayes Factors is the following (Berger and Pericchi(1996) [3])

Lemma 2.6.1. *Let $\mathbf{y} = \{y_1, \dots, y_n\}$ be independent and identically distributed random samples, and assume that $\mathbf{y}(\ell) \in D$ with $\mathbf{y}(-\ell) \cup \mathbf{y}(\ell) = \mathbf{y}$, then*

$$B_{ij}(\ell) = B_{ij}^N(\mathbf{y})B_{ji}^N(\mathbf{y}(\ell)), \forall \ell. \quad (2.1)$$

Proof. The Bayes Factor of the remaining samples given the training samples can be written as

$$B_{ij}(\mathbf{y}(-\ell)|\mathbf{y}(\ell)) = \frac{m_i(\mathbf{y}(-\ell)|\mathbf{y}(\ell))}{m_j(\mathbf{y}(-\ell)|\mathbf{y}(\ell))} = \frac{\int f_i(\mathbf{y}(-\ell)|\boldsymbol{\theta}_i)\pi(\boldsymbol{\theta}_i|\mathbf{y}(\ell))d\boldsymbol{\theta}_i}{\int f_i(\mathbf{y}(-\ell)|\boldsymbol{\theta}_j)\pi(\boldsymbol{\theta}_j|\mathbf{y}(\ell))d\boldsymbol{\theta}_j}$$

The proper posterior distribution of the parameters given the MTS can be written as

$$\pi(\boldsymbol{\theta}_k|\mathbf{y}(\ell)) = \frac{f_k(\boldsymbol{\theta}_k|\mathbf{y}(\ell))\pi_k^N(\boldsymbol{\theta}_k)}{m_k^N(\mathbf{y}(\ell))}, \text{ for } k = i, j$$

Hence,

$$\begin{aligned} \frac{\int f_i(\mathbf{y}(-\ell)|\boldsymbol{\theta}_i)\pi(\boldsymbol{\theta}_i|\mathbf{y}(\ell))d\boldsymbol{\theta}_i}{\int f_j(\mathbf{y}(-\ell)|\boldsymbol{\theta}_j)\pi(\boldsymbol{\theta}_j|\mathbf{y}(\ell))d\boldsymbol{\theta}_j} &= \frac{\int f_i(\mathbf{y}(-\ell)|\boldsymbol{\theta}_i)\frac{f_i(\boldsymbol{\theta}_i|\mathbf{y}(\ell))\pi_i^N(\boldsymbol{\theta}_i)}{m_i^N(\mathbf{y}(\ell))}d\boldsymbol{\theta}_i}{\int f_j(\mathbf{y}(-\ell)|\boldsymbol{\theta}_j)\frac{f_j(\boldsymbol{\theta}_j|\mathbf{y}(\ell))\pi_j^N(\boldsymbol{\theta}_j)}{m_j^N(\mathbf{y}(\ell))}d\boldsymbol{\theta}_j} \\ &= \frac{m_j^N(\mathbf{y}(\ell))}{m_i^N(\mathbf{y}(\ell))} \frac{\int f_i(\mathbf{y}(-\ell)|\boldsymbol{\theta}_i)f_i(\boldsymbol{\theta}_i|\mathbf{y}(\ell))\pi_i^N(\boldsymbol{\theta}_i)d\boldsymbol{\theta}_i}{\int f_j(\mathbf{y}(-\ell)|\boldsymbol{\theta}_j)f_j(\boldsymbol{\theta}_j|\mathbf{y}(\ell))\pi_j^N(\boldsymbol{\theta}_j)d\boldsymbol{\theta}_j} \end{aligned} \quad (2.2)$$

Since $f_k(\mathbf{y}(\ell)|\boldsymbol{\theta}_k)f_k(\mathbf{y}(-\ell)|\boldsymbol{\theta}_k) = f_k(\mathbf{y}|\boldsymbol{\theta}_k)$ for $k = i, j$, the ratio of the marginals is

$$\frac{m_j^N(\mathbf{y}(\ell))}{m_i^N(\mathbf{y}(\ell))} = B_{ji}^N(\mathbf{y}(\ell))$$

Hence, equation (2) becomes

$$B_{ji}^N(\mathbf{y}(\ell)) \frac{\int f_i(\mathbf{y}|\boldsymbol{\theta}_i)\pi_i^N(\boldsymbol{\theta}_i)d\boldsymbol{\theta}_i}{\int f_j(\mathbf{y}|\boldsymbol{\theta}_j)\pi_j^N(\boldsymbol{\theta}_j)d\boldsymbol{\theta}_j} = B_{ji}^N(\mathbf{y}(\ell))B_{ij}^N(\mathbf{y}).$$

□

By employing the methodology elucidated in the aforementioned lemmas, we have successfully mitigated the reliance of B_{ij} on the scales of the priors, thereby resolving the issue of undetermined constants associated with the Bayes factor. However, a new dependence arises, contingent upon the arbitrary selection of the minimal training sample denoted as $\mathbf{y}(\ell)$. To address this dependency and enhance stability, we propose to perform an "average" of $B_{ij}(\ell)$ over all feasible training samples $\mathbf{y}(\ell)$, where $\ell = 1, \dots, L$.

2.7 The Expected Posterior Bayes Factors

Suppose that non-informative (typically improper) priors $\pi_i^N(\boldsymbol{\theta}_i)$, where $i = 1, \dots, n$, associated with the models M_1, \dots, M_n , respectively are available. The concept of Expected Posterior (EP) priors, introduced by Perez and Berger in 2002 [10], allows for the calculation of proper posteriors $\pi_i^N(\boldsymbol{\theta}_i|\mathbf{y}(\ell))$ using a minimal training sample $\mathbf{y}(\ell)$. The EP prior for model M_i is then defined as:

$$\pi_i^{EP}(\boldsymbol{\theta}_i) = \frac{1}{L} \sum_{\ell=1}^L \pi_i^N(\boldsymbol{\theta}_i|\mathbf{y}(\ell))$$

The *EP Bayes factor*, which enables the comparison between models M_i and M_j , is computed as follows:

$$B_{ij}^{EP}(\mathbf{y}) = \frac{\int f_i(\mathbf{y}|\boldsymbol{\theta}_i)\pi_i^{EP}(\boldsymbol{\theta}_i)d\boldsymbol{\theta}_i}{\int f_j(\mathbf{y}|\boldsymbol{\theta}_j)\pi_j^{EP}(\boldsymbol{\theta}_j)d\boldsymbol{\theta}_j}$$

In the above equations, $f_i(\mathbf{y}|\boldsymbol{\theta}_i)$ and $f_j(\mathbf{y}|\boldsymbol{\theta}_j)$ represent the likelihood functions associated with models M_i and M_j , respectively. The EP prior, derived from improper priors using the minimal training sample $\mathbf{y}(\ell)$, allows for the calculation of the EP Bayes factor, facilitating model comparison and selection.

CHAPTER 3

A UNIVERSAL ROBUST BOUND

3.1 The Smith and Spiegelhalter's method

This chapter presents an innovative methodology for computing Bayes Factors, building upon the seminal work of Spiegelhalter and Smith (1982) [15]. Their approach, which encompassed examining an alternative prior specification, could be interpreted as a genuine subjective Bayesian analysis based on specific prior beliefs, or alternatively, as a formal analysis serving as a theoretical tool for comparing a model M_0 with a localized subset of models within another model M_1 . Spiegelhalter and Smith employed the concept of "imaginary training samples" to maximize support for the simpler model M_0 . While their paper primarily focused on assigning the constant c_0/c_1 in equation (5), Lempers (1971) [9] and Atkinson (1978) [1] emphasized that assigning vague prior information directly to c_0 and c_1 is inherently arbitrary without incorporating a training sample to determine appropriate variable scaling. Consequently, they proposed partitioning a portion of the data for a pre-comparison inference phase, enabling the formation of suitable priors for the parameters within each model.

The method suggested by Spiegelhalter and Smith involved determining \mathbf{y}_0 as the argument that maximizes $m_0^N(\mathbf{y}(\ell))/m_1^N(\mathbf{y}(\ell))$ and subsequently solving the equation

$$cB_{01}^N(\mathbf{y}_0) = 1.$$

This approach yielded the Bayes factor expressed as

$$B_{10}^{SS}(\mathbf{y}) = cB_{10}^N(\mathbf{y}) = B_{01}^N(\mathbf{y}_0)B_{10}^N(\mathbf{y}).$$

While Spiegelhalter and Smith’s technique focused exclusively on nested point hypotheses, the upcoming sections will expand upon this method to encompass more general cases and examples. Furthermore, a novel lower bound will be introduced, leveraging the concept of "imaginary training samples" to maximize support for M_0 . However, there is a distinction between this proposed method and the aforementioned approach by Smith and Spiegelhalter. Our method does not concern itself with the constant c ; instead, we aim to establish a lower bound for various forms of Intrinsic Bayes Factors, including the arithmetic, geometric, expected and harmonic. This bound encompasses both theoretical and empirical versions, with the empirical version providing the tightest bound. The subsequent section will outline the procedures for obtaining both theoretical and empirical bounds.

3.2 The universal robust bound for the Intrinsic Bayes Factors

In this study, we propose a novel method that generates a lower bound for the Bayes factor and corresponding priors, employing an approach that is closely related to the concept of Intrinsic Bayes Factors and the technique employed by Spiegelhalter and Smith. Let $\mathbf{y}(\ell)$ denote a minimal training sample, and $\pi(\theta_i|\mathbf{y}(\ell))$ represent the posterior distribution of parameter θ_i given the minimal training sample for two models, M_i and M_j . Our aim is to compare models M_i and M_j . To achieve this comparison, we define the set D as the aggregate of all possible minimal training samples with size k , where k is determined as the maximum of $\{\dim(\theta_i), \dim(\theta_j)\}$.

Definition 1. *The Intrinsic Bayes Factor Lower Bound (\underline{B}_{01}^I), is the theoretical lower bound of all possible Intrinsic Bayes Factors obtained by computing the infimum over the set of all possible imaginary training samples D .*

Definition 2. *The Empirical Intrinsic Bayes Factor Lower Bound (\underline{B}_{01}^{I*}) is the empirical lower bound obtained by computing the minimum over the set of all possible empirical training samples $D_{n,k}$.*

3.3 A bridge between the universal bound and the IBF

In this section, we will explore various approaches to construct lower bounds for the Intrinsic Bayes factor, ultimately contributing to the field of model comparison and selection. There is a bridge between IBF and the Least Favorable Intrinsic Bayes Factor approach that we want to explore by considering the AIBF, defined as,

$$B_{10}^{AI}(\mathbf{y}) = B_{10}^N(\mathbf{y}) \times \frac{1}{L} \sum_{\ell=1}^L B_{01}^N(\mathbf{y}(\ell)) \quad (3.1)$$

where $L = \binom{n}{m}$, where m is the minimal training sample size. The AIBF is asymmetric the more complex model has to be placed above to guarantee the convergence.

$$\bar{B}_{10}^I := B_{10}^N \sup_{\mathbf{y}(\ell) \in D} B_{01}(\mathbf{y}(\ell)) \geq B_{10}^N \frac{1}{L} \sum B_{01}(\mathbf{y}(\ell)) = B_{10}^A. \quad (3.2)$$

By forming the reciprocals, we have the important consequence for a bound on the probability of a null hypothesis:

$$\underline{B}_{01}^I := B_{01}^N \inf_{\ell} B_{10}(\mathbf{y}(\ell)) = B_{01}^N \frac{1}{\sup_{\mathbf{y}(\ell) \in D} B_{01}(\mathbf{y}(\ell))} \leq B_{01}^N \frac{1}{\frac{1}{L} \sum B_{01}(\mathbf{y}(\ell))} = B_{01}^{AI}. \quad (3.3)$$

where $\mathbf{y}(\ell)$ is a theoretical training sample from the set of all possible training samples. By construction, all Bayes Factors obey: $B_{01} = 1/B_{10}$ and $\bar{B}_{10}^I \geq B_{10}^{AI}$, since the supremum is necessarily bigger or equal than the arithmetic average. On the other hand, for generalizing the bound we may consider another bound based on empirical training samples,

$$B_{01}^{AI} \geq B_{01}^N \inf_{\mathbf{y}(\ell) \in D_{n,k}} B_{10}(\mathbf{y}(\ell)) = \underline{B}_{01}^{I*}, \quad (3.4)$$

where $\mathbf{y}(\ell)$ are empirical real training samples, and \underline{B}_{01}^{I*} is formed over all the real training samples, of a minimal size m . Clearly, $\bar{B}_{10}^{I*} \leq \bar{B}_{10}^I$ and $B_{01}^{AIBF} \geq \underline{B}_{01}^{I*} \geq \underline{B}_{01}^I$. Thus, \underline{B}_{01}^{I*} and \underline{B}_{01}^I are the least favorable lower bounds on the Bayes Factor. In addition, the results obtained with the AIBF are valid for other related types of Intrinsic Bayes Factors.

3.3.1 A comparison with the p-value based lower bound

In this section, we aim to illustrate the superiority of our bound in comparison to the well-known p-value-based robust lower bound introduced by Sellke et al. (2001) [13]. To commence our discussion, we will first establish the definition of a p-value.

Definition 3. A p-value $p(\mathbf{Y})$ is a statistic satisfying $0 \leq p(\mathbf{y}) \leq 1$ for every sample point \mathbf{y} . Small values of $p(\mathbf{Y})$ give evidence that $H_1 : \theta \in \Theta_0^c$ is true, where Θ_0 is some subset of the parameter space and Θ_0^c is its complement. A p-value is valid if, for every $\theta \in \Theta_0$ and every $0 \leq \alpha \leq 1$,

$$P_\theta(p(\mathbf{Y}) \leq \alpha) \leq \alpha.$$

The asymptotic behavior of the p-value as the sample size tends toward infinity is contingent upon both the null hypothesis and the fundamental distribution of the test statistic. Generally speaking, in the context of a two-sided hypothesis test, if the null hypothesis holds true, the p-value follows a uniform distribution within the range $[0,1]$. This implies that, irrespective of the sample size, the p-value will be uniformly distributed provided that the null hypothesis is valid.

The *Robust Lower Bound* offers a framework for calibrating p-values and is defined as follows:

$$B_{01}^L(p) = \begin{cases} -ep \log(p) & \text{if } p < \frac{1}{e} \\ 1 & \text{if } p \geq \frac{1}{e} \end{cases} \quad (3.5)$$

We postulate that given any sample size n and for a fixed p , our IBF Lower Bound as a function of the p-value, $\underline{B}_{01}^I(p)$, is superior to the robust lower bound $B_{01}^L(p)$. It can be mathematically expressed as

$$\underline{B}_{01}^I(p) \geq \underline{B}_{01}^L(p) \geq B_{01}^L(p).$$

The last part of this inequality was proven by Sellke et. al (2001) [13] and it should be noted that $B_{01}^J(p)$ is dynamic, changing with the sample size, akin to a real Bayes Factor, whereas $B_{01}^L(p)$ remains static with n . In the next theorem, we use the Laplace approximation to show that the robust lower bound stays constant under H_0 while our lower bound converges to the correct decision as the sample size increases. The Laplace approximation is a method used to approximate a probability distribution with a normal distribution and it's valid under the following regularity conditions:

1. **Smoothness of the Logarithm of the Density:** For a density function $f(y)$, if $\log(f(y))$ is twice continuously differentiable in a neighborhood around the mode of the distribution, then the Laplace approximation tends to work well.
2. **Existence of a Unique Mode:** The distribution should have a single, well-defined mode within the region of interest.
3. **Curvature of the Log Density:** The curvature (second derivative) of $\log(f(y))$ at the mode should not be zero. A non-zero curvature ensures that the distribution can be well-approximated by a quadratic form, which is the characteristic shape of a normal distribution.
4. **Convergence to Normality:** As the sample size tends to infinity, the distribution should converge to a normal distribution in a suitable sense (like the Central Limit Theorem).

These conditions, particularly the smoothness and curvature of the log density, are crucial for the Laplace approximation to accurately represent the shape of the distribution around its mode.

However, it's important to note that while the Laplace approximation is often effective for many distributions that meet these conditions, it might not work well for highly skewed or heavy-tailed distributions where the normality assumption might not hold. In such cases, other approximation methods or numerical techniques might be more appropriate.

Theorem 3.3.1. *Let $f_0(\mathbf{y}|\theta_0)$ and $f_1(\mathbf{y}|\theta_1)$ be the densities that satisfies regularity conditions (1)-(4). Let $H_0 \subset H_1$, $k_0 = \dim(\theta_0)$ and $k_1 = \dim(\theta_1)$ where $k_0 < k_1$, then under H_0*

$$\frac{B_{01}^L(p)}{\underline{B}_{01}^I} \rightarrow 0 \text{ as } n \rightarrow \infty$$

Proof. For a proper minimal training sample $\mathbf{y}^*(\ell)$, $\exists b, c$ such that $0 < b \leq c < \infty$ and

$$0 < b \leq B_{10}^N(\mathbf{y}^*(\ell)) \leq c$$

The Laplace approximation for $B_{01}^N(\mathbf{y})$ (Berger & Pericchi (2001) [4]) yields

$$\left(\frac{f_0(\mathbf{y}|\hat{\theta}_0)|\hat{I}_0|^{-1/2}}{f_1(\mathbf{y}|\hat{\theta}_1)|\hat{I}_1|^{-1/2}} \cdot \frac{\pi_0^N(\hat{\theta}_0)}{\pi_1^N(\hat{\theta}_1)} \right) \cdot b \leq B_{01}^N(\mathbf{y})B_{10}^N(\mathbf{y}^*(\ell)) \leq \left(\frac{f_0(\mathbf{y}|\hat{\theta}_0)|\hat{I}_0|^{-1/2}}{f_1(\mathbf{y}|\hat{\theta}_1)|\hat{I}_1|^{-1/2}} \cdot \frac{\pi_0^N(\hat{\theta}_0)}{\pi_1^N(\hat{\theta}_1)} \right) \cdot c$$

where \hat{I}_i and $\hat{\theta}_i$ are the observed information matrix and m.l.e., respectively under model M_i . The BIC approximation for \underline{B}_{01}^I gives us

$$b \cdot \frac{f_0(\mathbf{y}|\hat{\theta}_0)}{f_1(\mathbf{y}|\hat{\theta}_1)} \cdot (n)^{(k_1-k_0)/2} \leq \underline{B}_{01}^I \leq c \cdot \frac{f_0(\mathbf{y}|\hat{\theta}_0)}{f_1(\mathbf{y}|\hat{\theta}_1)} \cdot (n)^{(k_1-k_0)/2}$$

In the nested hypothesis case under the null hypothesis, the likelihood ratio in the equation above is bounded because $-2 \cdot \log(LR_{01})$ is bounded (it converges to a central Chi-square distribution by Wilks theorem). Hence, $\underline{B}_{01}^I(\mathbf{y}) \rightarrow \infty$ and $\bar{B}_{10}^I(\mathbf{y}) \rightarrow 0$ as $n \rightarrow \infty$ and by definition, $0 \leq B_{01}^L(p) \leq 1$. Hence,

$$0 \leq \frac{B_{01}^L(p)}{\underline{B}_{01}^I} \leq \frac{1}{\underline{B}_{01}^I} = \bar{B}_{10}^I \rightarrow \lim_{n \rightarrow \infty} 0 = 0 \leq \lim_{n \rightarrow \infty} \frac{B_{01}^L(p)}{\underline{B}_{01}^I} \leq \lim_{n \rightarrow \infty} \bar{B}_{10}^I = 0$$

By the Squeezing Theorem,

$$\lim_{n \rightarrow \infty} \frac{B_{01}^L(p)}{\underline{B}_{01}^I} = 0.$$

□

Although this proof is for nested hypotheses, we believe that this is true in general and we support this claim with examples for separated hypotheses (see Geometric vs. Poisson, Negative Binomial vs Poisson examples).

A study of the p-value and the RBL vs. the SS Bayes factor under H_1

We want to study the following limit under H_1

$$\lim_{n \rightarrow \infty} \frac{B_{01}^L(p)}{B_{01}^I}$$

Using an approach similar to the one used in the previous theorem, we can see that

$$\frac{f_1(\mathbf{y}|\hat{\theta}_1)}{c \cdot n^{k_1-k_0} f_0(\mathbf{y}|\hat{\theta}_0)} \leq \frac{B_{01}^L(p)}{B_{01}^I} \leq \frac{f_1(\mathbf{y}|\hat{\theta}_1)}{b \cdot n^{k_1-k_0} f_0(\mathbf{y}|\hat{\theta}_0)}$$

To explore this limit we need the asymptotic distribution of the likelihood ratio under the alternative hypothesis and related results by Self, S. G., & Liang, K. Y. (1987) [12] known as the generalized likelihood ratio test (GLRT) can be used to handle the case where the alternative hypothesis is true. The GLRT is a modification of the likelihood ratio test that allows for testing of both null and alternative hypotheses, and it is based on the same asymptotic theory as Wilks' theorem. Specifically, as the sample size goes to infinity, the distribution of the GLRT under the alternative hypothesis approaches a non-central chi-squared distribution with a non-centrality parameter proportional to the squared distance between the true parameter value and the null parameter value, and degrees of freedom equal to the difference in the number of parameters between the null and alternative hypotheses. Mathematically, we can write this as:

$$LR \sim \chi^2(p, \lambda) \text{ as } n \rightarrow \infty \text{ under } H_1$$

where $\chi^2(p, \lambda)$ denotes the non-central chi-squared distribution with p degrees of freedom, the non-centrality parameter λ and $LR = 2 * \log(L(\hat{\theta}_1)/L(\hat{\theta}_0))$. The GLRT can be used to test hypotheses in a wide range of statistical models, including linear regression, logistic regression, and ANOVA, among others. So, while Wilks' theorem applies only to the null hypothesis, the related GLRT provides a way to handle testing of the alternative hypothesis using the same underlying asymptotic theory.

In some cases, it may be possible to write the asymptotic distribution of the likelihood ratio test statistic as a function of n .

However, note that this formula applies only to specific cases, and the asymptotic distribution of the likelihood ratio test statistic may be different for other models and hypotheses. Determining the non-centrality parameter λ for the asymptotic distribution of the likelihood ratio test statistic can be challenging and requires knowledge of the true parameter values under the alternative hypothesis. In general, there is no simple formula that can be used to calculate λ for any hypotheses test. In some cases, simulation methods may be used to estimate the non-centrality parameter λ . This involves generating simulated data sets from the assumed distribution under the alternative hypothesis, calculating the likelihood ratio test statistic for each data set, and then estimating the distribution of the test statistic based on the simulated values. The estimated distribution can then be compared to the theoretical distribution to estimate the non-centrality parameter.

In general, determining the non-centrality parameter and its rate of convergence for the asymptotic distribution of the likelihood ratio test statistic is an active area of research, and the methods used to estimate λ depend on the specific model and hypothesis being tested. We conjecture that if there exist $r > 0$ such that $\lim_{n \rightarrow \infty} \frac{LR}{n^r}$ is constant under H_1 , then

$$\lim_{n \rightarrow \infty} \frac{B_{01}^L(p)}{B_{01}^I} = 0 \text{ if } k_1 - k_0 > r \quad (3.6)$$

because if this limit is zero, then similarly to the case when H_0 is true, we can show that the limit of the ratio is also zero. Further research is needed to determine if **Theorem 3.1** is true under the alternative hypothesis in general, or if this theorem only holds under certain conditions.

CHAPTER 4

THE SP BAYES FACTORS

In this chapter, we introduce a novel prior known as the SP Prior and leverage it to establish what we term the SP Bayes Factor. The designation 'SP' denotes the 'superior posterior' approach, which underpins the creation of these priors. The underlying concept involves transforming an improper prior into a proper posterior distribution by identifying either the theoretical or empirical minimal training sample that maximizes evidence for the null hypothesis (or yields the least favorable prior for the alternative).

Both the prior and the Bayes factor are presented in two versions-empirical and theoretical-similar to the universal lower bound discussed in the preceding chapter. These priors exhibit a close relationship with IBF bounds, which we elucidate within this chapter.

The empirical prior proves valuable in cases where theoretical maximization poses challenges. Remarkably, the empirical SP Bayes Factor closely approximates the actual Bayes factor value compared to the theoretical counterpart. Nevertheless, in scenarios where theoretical maximization remains viable, the theoretical prior offers computational efficiency over the empirical one. Employing both priors, the SP Bayes factors converge towards accurate decisions with increasing sample sizes-a point we substantiate both theoretically and through numerical examples in subsequent sections.

4.1 The SP Prior

Definition 4 (The Theoretical SP Prior). *The Theoretical SP prior for B_{10} is defined as*

$$\pi_k^{SP}(\theta_k) = \pi_k(\theta_1 | \mathbf{y}_T^*(\ell)) = \frac{f_k(\mathbf{y}_T^*(\ell) | \theta_k) \pi_k^N(\theta_k)}{m_k(\mathbf{y}_T^*(\ell))}, k = 0, 1$$

$$\text{where } \mathbf{y}_T^*(\ell) = \arg \left[\sup_{\mathbf{y}(\ell) \in D} B_{01}(\mathbf{y}(\ell)) \right].$$

Definition 5 (The Empirical SP Prior). *The Empirical SP prior for B_{10} is defined as*

$$\pi_k^{SP*}(\theta_k) = \pi_k(\theta_k | \mathbf{y}_E(\ell)) = \frac{f_k(\mathbf{y}_E^*(\ell) | \theta_k) \pi_k^N(\theta_k)}{m_k(\mathbf{y}_E^*(\ell))}, k = 0, 1$$

$$\text{where } \mathbf{y}_E^*(\ell) = \arg \left[\sup_{\mathbf{y}(\ell) \in D_{n,k}} \frac{m_0(\mathbf{y}(\ell))}{m_1(\mathbf{y}(\ell))} \right].$$

4.2 The SP Bayes Factor

Definition 6 (The Theoretical SP Bayes Factor).

If $\arg \sup_{\mathbf{y}(\ell) \in D} B_{01}(\mathbf{y}(\ell)) < \infty$, then

$$B_{10}^{SP} = \frac{m_1^{SP}(\mathbf{y})}{m_0^{SP}(\mathbf{y})} = \frac{\int f_1(\mathbf{y} | \theta_1) \pi_1^{SP}(\theta_1) d\theta_1}{\int f_0(\mathbf{y} | \theta_0) \pi_0^{SP}(\theta_0) d\theta_0}.$$

Definition 7 (The Empirical SP Bayes Factor).

Let $\mathbf{y}_T^*(\ell) = \arg \sup_{\mathbf{y}(\ell) \in D_{n,k}} B_{01}(\mathbf{y}(\ell))$, and $\mathbf{y}(-\ell) = \mathbf{y} \setminus \mathbf{y}_T^*(\ell)$, then

$$B_{10}^{SP*} = \frac{m_1^{SP*}(\mathbf{y})}{m_0^{SP*}(\mathbf{y})} = \frac{\int f_1(\mathbf{y}(-\ell) | \theta_1) \pi_1^{SP*}(\theta_1) d\theta_1}{\int f_0(\mathbf{y}(-\ell) | \theta_0) \pi_0^{SP*}(\theta_0) d\theta_0}.$$

4.3 Results

Theorem 4.3.1. *The Empirical IBF upper and lower bounds can be obtained using the Empirical SP Prior, that is*

$$B_{10}^{SP*} = B_{10}^N[\mathbf{y}] \sup_{\ell=1,\dots,L} B_{01}[\mathbf{y}(\ell)] = \overline{B}_{10}^{I*}(\mathbf{y}).$$

Proof. Let $\mathbf{y}(-\ell) = \mathbf{y} \setminus \{\mathbf{y}_E^*(\ell)\}$, without loss of generality, we will prove it for the upper-bound with the supremum since the proof for the infimum lower bound is identical. The Empirical SP Bayes factor is

$$B_{10}^{SP*} = \frac{m_1^{SP*}(\mathbf{y})}{m_0^{SP*}(\mathbf{y})} = \frac{\int f_1(\mathbf{y}(-\ell)|\boldsymbol{\theta}_1) \frac{f(\mathbf{y}_E^*(\ell)|\boldsymbol{\theta}_1)\pi_i^N(\boldsymbol{\theta}_1)}{m_1(\mathbf{y}_E^*(\ell))} d\boldsymbol{\theta}_1}{\int f_0(\mathbf{y}(-\ell)|\boldsymbol{\theta}_0) \frac{f(\mathbf{y}_E^*(\ell)|\boldsymbol{\theta}_0)\pi_j^N(\boldsymbol{\theta}_0)}{m_0(\mathbf{y}_E^*(\ell))} d\boldsymbol{\theta}_0}$$

The marginal distributions m_0 and m_1 do not depend on $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$, hence we can take those terms out of the integrals to obtain

$$\begin{aligned} \frac{m_0(\mathbf{y}_E^*(\ell))}{m_1(\mathbf{y}_E^*(\ell))} \frac{\int f_1(\mathbf{y}(-\ell)|\boldsymbol{\theta}_1) f(\mathbf{y}_E^*(\ell)|\boldsymbol{\theta}_1) \pi_i^N(\boldsymbol{\theta}_1) d\boldsymbol{\theta}_1}{\int f_0(\mathbf{y}(-\ell)|\boldsymbol{\theta}_0) f(\mathbf{y}_E^*(\ell)|\boldsymbol{\theta}_0) \pi_0^N(\boldsymbol{\theta}_0) d\boldsymbol{\theta}_0} &= \frac{m_0(\mathbf{y}_E^*(\ell))}{m_1(\mathbf{y}_E^*(\ell))} \frac{m_1(\mathbf{y}(-\ell), \mathbf{y}_E^*(\ell))}{m_0(\mathbf{y}(-\ell), \mathbf{y}_E^*(\ell))} \\ &= B_{01}[\mathbf{y}_E^*(\ell)] B_{10}^N[\mathbf{y}(-\ell), \mathbf{y}_E^*(\ell)] \end{aligned}$$

By definition, we have that

$$B_{01}[\mathbf{y}_E^*(\ell)] = B_{01}[\arg(\sup_{\ell=1,\dots,L} B_{01}(\mathbf{y}(\ell)))] = \sup_{\ell=1,\dots,L} B_{01}[\mathbf{y}(\ell)]$$

and $\{\mathbf{y}_E^*(\ell)\} \cup \{\mathbf{y}(-\ell)\} = \mathbf{y}$ and $m_k(\mathbf{y}(-\ell), \mathbf{y}_E^*(\ell)) = m_k(\mathbf{y})$, for $k = 0, 1$

Therefore,

$$B_{10}^{SP*} = B_{10}^N[\mathbf{y}] \sup_{\ell=1,\dots,L} B_{01}^N[\mathbf{y}(\ell)] = \overline{B}_{10}^{I*}(\mathbf{y})$$

□

Theorem 4.3.2. *Let $f(\mathbf{y}|\theta)$ be a p.d.f or p.m.f a random variable Y with support $S = \{\mathbf{y} : f(\mathbf{y}|\theta) > 0\}$. Suppose we want to calculate the Bayes factor for the following nested hypothesis test $H_0 : \theta = \theta_0$ vs $H_1 : \theta \neq \theta_0$, then the empirical SP-Prior and SP-BF converges to the Theoretical SP-Prior and SP-BF.*

Proof. Without loss of generality we can assume that $\dim(\theta) = 1$, so that the minimal training sample size is one. In this case, the set $D_{n,k} = D_{n,1} = \mathbf{y}_n = (y_1, \dots, y_n) \subseteq S$, $\lim_{n \rightarrow \infty} D_{n,1} = D$ and $L = \binom{n}{k} = \binom{n}{1} = n$. The proof will be divided in the discrete and continuous support cases.

Case 1: The support S is discrete

$$\lim_{n \rightarrow \infty} D_{n,1} = \lim_{L \rightarrow \infty} \{\mathbf{y}(\ell)\}_{\ell=1}^L = \{\mathbf{y}(\ell)\}_{\ell=1}^{\infty} = D = S$$

Therefore,

$$\lim_{n \rightarrow \infty} B_{01}^{SP*}(\mathbf{y}) = B_{01}^N[\mathbf{y}] \lim_{L \rightarrow \infty} \sup_{\{\mathbf{y}(\ell)\}_{\ell=1}^L} B_{10}[\mathbf{y}(\ell)] = B_{01}^N[\mathbf{y}] \sup_D B_{10}[\mathbf{y}(\ell)]$$

Hence,

$$\lim_{n \rightarrow \infty} B_{01}^{SP*}(\mathbf{y}) = B_{01}^{SP}(\mathbf{y})$$

Clearly, in the discrete support case, we can eventually (for some n) obtain the minimal training sample $\mathbf{y}_n^*(\ell)$ that generates the theoretical supremum, and the empirical SS-BF will be exactly equal to the theoretical SSBF.

Case 2: The support S is continuous

In the case of a continuous support S , we can only find a sequence of minimal training samples $\{\mathbf{y}_n^*(\ell)\}$ that gets arbitrarily close to a minimal imaginary sample $\mathbf{y}^*(\ell)$ that produces the theoretical supremum. We will construct such a sequence and prove its

existence. Assume that $\mathbf{y}^*(\ell) = \arg \sup_D B_{10}(\mathbf{y}(\ell))$ and for each n ,

$$\mathbf{y}_n^*(\ell) = \mathbf{y}(\ell_{i_n}), \text{ where } i_n = \arg \min_{i \in \{1, \dots, \binom{n}{1}\}} |\mathbf{y}^*(\ell) - \mathbf{y}(\ell_i)|$$

We claim that $\mathbf{y}_n^*(\ell) \rightarrow \mathbf{y}^*(\ell)$ for any $\mathbf{y}^*(\ell) \in D$

Proof:

Assume that $\mathbf{y}^*(\ell) \in D$ and $\exists \varepsilon > 0$ such that

$$\forall n \in \mathbb{N}, |\mathbf{y}^*(\ell) - \mathbf{y}_n^*(\ell)| > \varepsilon$$

Then the interval $(\mathbf{y}^*(\ell) - \frac{\varepsilon}{2}, \mathbf{y}^*(\ell) + \frac{\varepsilon}{2}) \cap D = \emptyset$ which is a contradiction since $\mathbf{y}^*(\ell) \in D$, thus the intersection can not be empty. Similarly, if such a point does exist in D , then

$$\int_{\mathbf{y}^*(\ell) - \frac{\varepsilon}{2}}^{\mathbf{y}^*(\ell) + \frac{\varepsilon}{2}} f(\mathbf{y}|\theta) d\mathbf{y} = 0 \implies f \text{ is not continuous in } \mathbf{y}^*(\ell) \implies \mathbf{y}^*(\ell) \notin D$$

□

CHAPTER 5

HYPOTHESIS TESTING WITH THE NORMAL DISTRIBUTION

This chapter explores hypothesis testing for mean and precision (the inverse of scale) parameters within the Normal Distribution. It covers essential concepts such as the universal lower bound, the SP/Intrinsic Priors, and the SP/Intrinsic Bayes factor. Through numerical experiments and visual charts, we compare these approaches and visualize the results. By blending theoretical foundations with practical applications, this chapter aims to provide readers with a robust understanding of conducting and interpreting these crucial statistical tests, leveraging our novel theoretical framework..

5.1 The Normal Scale Hypothesis Test

Assume $\mathbf{y} = \{y_1, y_2, \dots, y_n\}$ are i.i.d samples where $y_i \sim N(y_i|\mu, h)$ for every i , and that the mean μ is unknown. We wish to decide if the precision parameter $h = 1/\sigma^2$ is a specified value. If we want to make a decision, we should perform the following hypothesis test

$$H_0 : h = h_0 \text{ vs } H_1 : h \neq h_0,$$

The likelihood is $f(\mathbf{y}|\mu, h) = (\frac{h}{2\pi})^{n/2} \exp[-\frac{h}{2}(S^2 + n(\bar{\mathbf{y}} - \mu)^2)$, where $S^2 = \sum_{i=1}^n (y_i - \bar{\mathbf{y}})^2$. As a first step, use the so-called "Jeffrey's rule" equating the prior to the square root of the determinant of the Fisher Information. In the motivating example this leads to: $\pi^J(\mu, h) = C_0 \cdot \frac{d\mu dh}{\sqrt{h}}$. For the so-called "independence" Jeffrey's we have instead

$$\pi_I^J(\mu, h) = \pi^J(\mu)\pi^J(h) = C_1 \cdot \frac{d\mu dh}{h}$$

Putting together both priors $\pi^J(\mu, h) = C/h^r$ with $r = 1/2, 1$ for the dependent or independent Jeffreys prior respectively. It turns out that the Bayes Factor is,

$$B_{01}^N(\mathbf{y}) = \frac{h_0^{(n-1)/2} \exp(-h_0 S^2/2)}{(2/S^2)^{(n-r)/2} \Gamma((n-r)/2)} \times \frac{C_0}{C_1}, \quad (5.1)$$

where r is equal to $1/2$ for the dependent prior, and for the independent prior equal to one. In this thesis, we assume the independent prior, as put forward by Jeffreys. It is apparent from this expression that the Bayes Factor is undetermined since it depends on the undefined ratio C_0/C_1 , the undefined constants that come from the improper priors. It can be argued, from several points of view, that the constants regarding the location μ cancel out, leaving only the indeterminacy related to the hypothesis parameter h . Some of these points of view are among others:

1. Mean and Precision parameters are orthogonal in the Fisher Information Matrix, and thus "cancel-out" in the ratio of marginal likelihoods
2. It turns out that two location models are predictively matched (see Pericchi (2005) [11]) and when the scale is integrated into the denominator, models under both hypotheses become location models, and the corrections cancel out.

5.1.1 The universal lower bound

It is evident from the previous example that, unfortunately, the Bayes factor remains defined only up to the constant that pertains to the precision parameter under test. This is why Jeffrey's suggested conventional proper priors for the extra-parameters under the larger hypothesis (but improper for common parameters like μ here). There have been suggested techniques around this problem, as in Smith and Spiegelhalter (1980) [14] and (1982) [15]. These techniques, although approximate, are useful in devising sensible scaling for Bayes Factors. In fact, we proved there for the first time that this approach has a very direct relationship with more recent approaches that have been studied in detail, particularly the Intrinsic Bayes Factor, the Intrinsic Priors, and EP-priors.

In the example of the section, the minimal training sample consists of two observations $(y(\ell_1), y(\ell_2))$, and the expression (2) turns out to be

$$B_{01}^N(\mathbf{y}) \times B_{10}^N(\mathbf{y}(\ell)) = \frac{h_0^{(n-1)/2} \exp(-h_0 S^2/2)}{(2/S^2)^p \frac{C_0}{C_1} \sqrt{h_0/\pi} \cdot \exp(-h_0 \mathbf{D}(\ell)^2/4) |\mathbf{D}(\ell)| \cdot \frac{C_1}{C_0}},$$

where $\mathbf{D}(\ell) = y(\ell_1) - y(\ell_2)$ and $p = (n-1)/2\Gamma((n-1)/2)$. It is apparent from (3) that the undefined constants cancel out. As a specific example, take $H_0 : h_0 = 1$. The crucial step is: what to do about the (theoretical or imaginary) training samples summary statistics $\mathbf{D}(\ell)$? SSBF's practical and simplifying approach is to take:

$$\sup_{\mathbf{y}(\ell) \in D} B_{10}^N(\mathbf{D}(\ell)), \quad (5.2)$$

which is attained at $\hat{D} = \sqrt{2}$ and the "correction" (4) is equal to 0.484. The *IBF* bound is then $\bar{B}_{01}^I = B_{01}^N \times \sup_{\mathbf{y}(\ell) \in D} B_{10}^N(\mathbf{y}(\ell))$. To see the relationship with more recent approaches like Intrinsic Bayes Factors the sup above is replaced by the arithmetic mean or by the theoretical expectation, on $\mathbf{y}(\ell)$, under $H_1 \supset H_0$. The **theoretical** expressions in (4) and (5), which take the formal maximum possible value for all possible training samples, can be refined by the **empirical** observed maximum for the observed $\mathbf{y}(\ell)$'s:

$$\sup_{\ell=1, \dots, L} B_{10}^N(\mathbf{y}(\ell)) \leq \sup_{\mathbf{y}(\ell) \in D} B_{10}^N(\mathbf{y}(\ell)), \quad (5.3)$$

where the distinction arises as $\mathbf{y}(\ell)$ is the set of empirically observed training samples and D which is the set of all possible minimal training samples from the whole support. The empirical bound is then sharper than the theoretical bound, and we can refer to the theoretical bound and the empirical bound of the Intrinsic Bayes Factors. In the Expected IBF case, for any pair of i.i.d. training samples $\{y(\ell_1), y(\ell_2)\}$, assumed (when taking expectations) to be drawn from $H_1 \supset H_0$, results in $\mathbf{D}(\ell) \sim N(0, 2/h)$. Then under H_1 :

$$\frac{B_{10}^{EIBF}(\mathbf{y})}{B_{10}^N(\mathbf{y})} = \int \frac{|\mathbf{D}|}{\sqrt{\pi}} \exp\left(-\frac{\mathbf{D}^2}{4}\right) \cdot \frac{\sqrt{h} \cdot \exp(-h \cdot \mathbf{D}^2/4)}{2\sqrt{\pi}} d\mathbf{D} = \frac{2\sqrt{h}}{\pi(h+1)}.$$

In an applied problem with real data, the parameter h will be estimated by its MLE. In this theoretical exercise, we study it as a function of h , and computation yields that its maximum is attained at the null hypothesis $h = h_0 = 1$ and its value at its maximum is $1/\pi = 0.318$.

5.1.2 Remarks

1. Our ratio of constants is 0.484 and the Maximum of the Expected IBF of the Training Sample is 0.318. This may seem like a big difference (the ratio is about .66). These values are attained close to the null hypothesis, the difference far from the null, may be much bigger but there, the dominant factor $B_{01}^N(\mathbf{y})$, will overwhelm that difference for a moderate sample, and even small but larger than the training sample sizes.
2. The lower bound is robust, it does not depend on the fact that all the training samples belong to the alternative hypothesis H_1 and that all belong to the same population. This bound achieves a very important robust effect. The bound has Universal validity: It is still valid under violations of distributional assumptions, for example, Normality or homogeneity of the population of training samples. It can also be used to check the computations of IBFs and with Intrinsic Priors.
3. The SS's bound on Intrinsic Bayes Factors is typically easy to compute and generalize to realistic scenarios.
4. In terms of p-values, if it can be found a "tight" bound $g(p)$ so that $g(p) \leq B_{01}^N$ then

$$g(p) \times \frac{1}{\sup B_{01}^N(\mathbf{y}(\ell))} \leq \underline{B}_{01}^I \leq B_{01} \quad (5.4)$$

which is a Universal Bound that adapts to the sample size n .

5.1.3 The SP and Intrinsic Priors

In this section, we turn our attention to the calculation of Superior Posterior (SP) and Intrinsic priors within the context of the Normal Scale Hypothesis testing scenario. An in-depth exploration will be conducted to illuminate the nuanced attributes and similarities shared between these two types of priors. To enhance the comprehension and tangibility of the underlying principles, we shall include illustrative representations in the form of charts. These graphical elucidations will serve as effective tools for visualizing the properties and similarities of both SP and Intrinsic priors.

Moreover, to ensure results and facilitate the application of the presented methods, we will accompany our discussion with executable **R** code snippets. These will illustrate the computational procedures required to carry out the prior calculations and generate the corresponding charts. By making this code available, we aim to provide a practical and accessible means for readers to engage with our methodology. Consider the Bayes factor for the minimal training sample of the normal scale example with an unknown mean

$$B_{10}^N(\mathbf{y}(\ell)) \propto \frac{1}{\exp\{-h_0\mathbf{D}(\ell)^2/4\}|\mathbf{D}(\ell)|} = \frac{\exp\{h_0\mathbf{D}(\ell)^2/4\}}{|\mathbf{D}(\ell)|}$$

To find the value of $\mathbf{D}(\ell)$ that generates the sup of $B_{10}^N(\mathbf{y}(\ell))$, we must take logarithms on both sides, and differentiate the equation $B_{10}^N(\mathbf{y}(\ell)) = 0$ with respect to $\mathbf{D}(\ell)$,

$$\frac{d}{d\mathbf{D}(\ell)} \log\left(\frac{\exp\{h_0\mathbf{D}(\ell)^2/4\}}{|\mathbf{D}(\ell)|}\right) = 2h_0\frac{\hat{D}}{4} - \frac{1}{\hat{D}} = 0 \implies h_0\frac{\hat{D}^2}{2} = 1 \implies \hat{D} = \sqrt{\frac{2}{h_0}}$$

After calculating this value, we must consider the prior of μ and h given $\mathbf{y}(\ell)$ which is

$$\pi(\mu, h|\mathbf{y}(\ell)) = \frac{f(\mathbf{y}(\ell)|\mu, h)\pi(\mu)\pi(h)}{\int f(\mathbf{y}(\ell)|\mu, h)\pi(\mu)\pi(h)d\mu dh} = \frac{\frac{1}{h} \frac{h}{2\pi} \exp\{-\frac{h}{2}[(y(1) - \mu)^2 + (y(2) - \mu)^2]\}}{\frac{1}{2\pi} \int \exp\{-\frac{h}{2}\left(\frac{\mathbf{D}(\ell)^2}{2} + 2(\mu - \bar{\mathbf{y}}(\ell))^2\right)\}d\mu dh} \quad (5.5)$$

Simplifying this equation, we obtain

$$\frac{\exp\{-\frac{h}{2}[\frac{\mathbf{D}(\ell)^2}{2} + 2(\mu - \bar{y}(\ell))^2]\}}{\int_0^\infty \exp\{-\frac{h}{4}\mathbf{D}(\ell)^2\} \frac{\sqrt{2\pi}}{\sqrt{2h}} dh}$$

To solve the last integral, we should consider the following change of variable

$$u = \frac{h\mathbf{D}(\ell)^2}{4}, \quad h = \frac{4u}{\mathbf{D}(\ell)^2}, \quad dh = \frac{4}{\mathbf{D}(\ell)^2} du$$

$$\frac{4\sqrt{\pi}}{\mathbf{D}(\ell)^2} \int_0^\infty \left(\frac{4u}{\mathbf{D}(\ell)^2}\right)^{\frac{1}{2}-1} \exp\{-u\} du = \Gamma\left(\frac{1}{2}\right) \frac{4\sqrt{\pi} D}{\mathbf{D}^2} \frac{1}{2} = \frac{2\pi}{\mathbf{D}}$$

Hence,

$$\pi(\mu, h|D = \hat{D}) = \frac{\exp\{-\frac{h}{2}[\frac{\hat{D}^2}{2} + 2(\mu - \bar{y}(\ell))^2]\} \hat{D}}{2\pi} \quad (5.6)$$

Integrating with respect to μ , we obtain the prior for h

$$\begin{aligned} \pi(h|\mathbf{D}) &= \int \pi(\mu, h|\mathbf{D}) d\mu = \int \frac{\mathbf{D}}{2\pi} \exp\left\{-\frac{h}{2}\left(\frac{\mathbf{D}^2}{2} + 2(\mu - \bar{y}(\ell))^2\right)\right\} d\mu \\ &= \frac{\mathbf{D}}{2\pi} \exp\left\{-\frac{h\mathbf{D}^2}{2}\right\} \frac{\sqrt{2\pi}}{\sqrt{2h}} = \frac{\mathbf{D} \exp\left\{-\frac{h\mathbf{D}^2}{4}\right\}}{\sqrt{2\pi}\sqrt{2h}} \end{aligned}$$

and by substituting \mathbf{D} by \hat{D} , we can obtain the SP prior for h

$$\boxed{\pi^{SP}(h|\hat{D} = \sqrt{\frac{2}{h_0}}) = \frac{1}{\sqrt{2\pi h h_0}} \exp\left\{-\frac{h}{2h_0}\right\}}$$

Evidently, the prior $\pi^{SP}(h)$ is *Gamma*($\alpha = 1/2, \beta = 2h_0$). To find the prior for μ , we must integrate with respect to h as follows

$$\pi^{SP}(\mu|\mathbf{D}) = \int_0^\infty \exp\left\{-\frac{h}{2}\left[\frac{\mathbf{D}^2}{2} + 2(\mu - \bar{y}(\ell))^2\right]\right\} \frac{\mathbf{D}}{2\pi} dh = \frac{8\pi}{\mathbf{D}^2 + 4(\mu - \bar{y}(\ell))^2}$$

$$\pi^{SP}(\mu|\hat{D} = \sqrt{\frac{2}{h_0}}) = \frac{8\pi}{\frac{2}{h_0} + 4(\mu - \bar{y}(\ell))^2} = \frac{2\pi}{\frac{1}{2h_0} + (\mu - \bar{y}(\ell))^2} = \frac{4h_0\pi}{1 + \left(\frac{\mu - \bar{y}(\ell)}{\sqrt{2h_0}}\right)^2}$$

The distribution of μ for B_{01} is proportional to a *Cauchy*(\bar{y} , γ) where $\gamma = \sqrt{2h_0}$.

$$\pi^{SP}(\mu) \propto \frac{1}{\pi\gamma\left(1 + \left(\frac{\mu - \bar{y}(\ell)}{\gamma}\right)^2\right)}$$

Next, we will compute the intrinsic priors for the same hypothesis test above to compare those with the SP priors. The improper intrinsic prior for μ is proportional to a constant $\pi^I(\mu) \propto c$. In the other hand, the intrinsic prior for the scale h is

$$\pi^I(\mu, h) = \pi^N(\mu, h) \int \frac{m_0^N(\mathbf{y}(\ell))}{m_1^N(\mathbf{y}(\ell))} f(\mathbf{y}(\ell)|\mu, h) d\mathbf{y}(\ell) = \int \pi(\mu, h|\mathbf{y}(\ell)) m_0(\mathbf{y}(\ell)) d\mathbf{y}(\ell)$$

$$\pi^I(h) \propto \frac{1}{\pi} \frac{\sqrt{\frac{h}{h_0}}}{\frac{h+1}{h_0}} \frac{1}{h} = \frac{1}{\pi h_0} \frac{\left(\frac{h}{h_0}\right)^{1/2-1}}{\left(\frac{h}{h_0} + 1\right)}$$

Evidently, the prior $\pi^I(h)$, is a *SBeta2*($h|p = 1/2, q = 1/2, b = h_0$) (Scale beta2 distribution). The Scale beta two density is given by

$$f(y|p, q, b) = \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} \frac{(y/b)^{p-1}}{\left(\frac{y}{b} + 1\right)^{(p+q)}, \quad y, p, q, b > 0$$

The subsequent script is written in the **R** programming language that facilitates the generation of a comparative graphical representation of both priors. These priors are superimposed on the same set of axes to foster a direct comparison. This visual comparison is executed across distinct values of the null hypothesis ($H_0 = 1$ and $H_0 = 10$).

R program

```

1 sp_prior = function(h,h0) { (1/sqrt(2*pi*h*h0))*exp(-h/(2*h0)) }
2 i_prior = function(h,h0) { (1/(pi*h0))*(h/h0)^(-0.5)/(h/h0+1) }
3 par(mfrow=c(1,2)); h = seq(0.1,5,0.01);h0=1; iprior = i_prior(h,h0); spprior = sp_prior(h,h0)
4 ymin = min(spprior, iprior); ymax = max(spprior, iprior)
5 plot(h, spprior, type="l", main="Comparison of the SP and Intrinsic Priors of h", ylab="Priors")
6 mtext(text="Hypothesis test: H0 = 1 vs H0 != 1", side=3); lines(h, iprior, type="l", col="red")
7 legend(3.4, 1.2, legend=c("SP", "Intrinsic"), col=c("black", "red"), lty=c(1,1), cex=0.8, title="Priors", text.
      font=4, bg='lightblue')
8 h = seq(1,15,0.01);h0=10; iprior = i_prior(h,h0); spprior = sp_prior(h,h0)
9 ymin = min(spprior, iprior); ymax = max(spprior, iprior)
10 plot(h, spprior, type="l", main="Comparison of the SP and Intrinsic Priors of h", ylim=c(ymin, ymax), ylab="Priors")
11 mtext(text="Hypothesis test: H0 = 10 vs H0 != 10", side=3); lines(h, iprior, type="l", col="red")
12 legend(10.5, 0.12, legend=c("SP", "Intrinsic"), col=c("black", "red"), lty=c(1,1), cex=0.8, title="Priors", text.
      font=4, bg='lightblue')

```





Figure 5–1: A comparison between the SP Prior and the Intrinsic Prior for the Normal Scale Hypothesis test for different null hypothesis values.

As illustrated in the preceding figures, a noteworthy characteristic of the Superior Posterior (SP) Prior emerges. This prior assigns a heightened probability to the null hypothesis in comparison to the Intrinsic Prior, specifically when the values of h are less than that of the null hypothesis. It is furthermore observed that the probability ascribed by the SP is invariably equal to or exceeds that of the Intrinsic Prior under these conditions. This pattern elucidates the differential weights that these priors place on the null hypothesis, offering valuable insights for our investigation.

5.2 The Normal Mean Hypothesis Test

This section is dedicated to exploring the domain of Normal Mean Hypothesis Testing. Here, we uncover the core principles and methodologies that drive inference regarding the mean parameter within the Normal Distribution. Assume that $y_i \sim N(\mu, \sigma_0^2)$ and consider the following hypothesis test

$$H_0 : \mu = \mu_0 \text{ vs } H_1 : \mu \neq \mu_0, \sigma_0, \text{ known}$$

Hence, the non-informative prior is $\pi^N(\mu) = 1$ and a minimal training sample is only one observation $y(\ell)$, for any $\ell = 1, \dots, n$. We will use that information to calculate different Bayes factors and their bounds in the next sections.

5.2.1 The Expected Intrinsic Bayes Factor for the Normal Mean Hypothesis Test

Consider the general formula for the Expected Intrinsic Bayes Factor (EIBF)

$$B_{10}^{EI}(\mathbf{y}) = B_{10}^N(\mathbf{y}) E_{\mathbf{y}(\ell)|H_0}^{\hat{\theta}} \left(\frac{m_0^N(\mathbf{y}(\ell))}{m_1^N(\mathbf{y}(\ell))} \right)$$

For this specific hypothesis test, the EIBF can be obtained by solving the following integral

$$E_{\mathbf{y}(\ell)|H_0}^{\hat{\theta}} \left(\frac{m_0^N(\mathbf{y}(\ell))}{m_1^N(\mathbf{y}(\ell))} \right) = \int_{-\infty}^{\infty} \left(\frac{1}{\sqrt{2\pi}\sigma_0} \right)^2 \exp\left\{-\frac{1}{2\sigma_0^2}(\mathbf{y}(\ell) - \mu_0)^2\right\} \exp\left\{-\frac{1}{2\sigma_0^2}(\mathbf{y}(\ell) - \mu)^2\right\} d\mathbf{y}(\ell)$$

This integral can be easily solved to obtain the following equality

$$E_{\mathbf{y}(\ell)|H_0}^{\hat{\theta}} \left(\frac{m_0^N(\mathbf{y}(\ell))}{m_1^N(\mathbf{y}(\ell))} \right) = N(\mu|\mu_0, 2\sigma_0^2) = \frac{1}{\sqrt{2\pi}\sqrt{2\sigma_0^2}} \exp\left\{\frac{-1}{4\sigma_0^2}(\mu - \mu_0)^2\right\}$$

In this case, the parameter μ is unknown and we should estimate it using the maximum likelihood estimator \bar{y} to obtain

$$E_{\mathbf{y}(\ell)|H_0}^{\hat{\theta}} \left(\frac{m_0^N(\mathbf{y}(\ell))}{m_1^N(\mathbf{y}(\ell))} \right) = \frac{1}{\sqrt{2\pi}\sqrt{2\sigma_0^2}} \exp\left\{\frac{-1}{4\sigma_0^2}(\bar{y} - \mu_0)^2\right\}$$

After finding this term, we can proceed to calculate the Expected Intrinsic Bayes Factor

$$B_{10}^{EI} = \frac{1}{\sqrt{2\pi}\sqrt{2\sigma_0^2}} \exp\left\{\frac{-1}{4\sigma_0^2}(\bar{y} - \mu_0)^2\right\} \frac{\sqrt{2\pi}\frac{\sigma_0}{\sqrt{n}}}{\exp\left\{\frac{-1}{2\sigma_0^2}n(\bar{y} - \mu_0)^2\right\}}$$

Simplifying the equation above yields

$$B_{10}^{EI} = \frac{1}{\sqrt{2n}} \exp\left\{\frac{1}{2\sigma_0^2}\left(n - \frac{1}{2}\right)(\bar{y} - \mu_0)^2\right\}$$

5.2.2 The Intrinsic Bayes Factor Bounds for the Normal Mean Hypothesis Test

Let us consider a scenario where the random vector $\mathbf{Y} = (Y_1, \dots, Y_n)$ is composed of independent and identically distributed (i.i.d) random variables Y_i , following a normal distribution with mean μ_0 and standard deviation σ_0 (known) under the null hypothesis M_0 , and a normal distribution with mean μ and standard deviation σ_0 (known) under the alternative hypothesis M_1 . We adopt the non-informative prior, specifically $\pi_1^N(\mu) = 1$.

$$B_{01}^N(\mathbf{y}(\ell)) = \frac{\frac{1}{\sqrt{2\pi}\sigma_0} \exp\left\{-\frac{1}{2\sigma_0^2}(\mathbf{y}(\ell) - \mu_0)^2\right\}}{\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left\{-\frac{1}{2\sigma_0^2}(\mathbf{y}(\ell) - \mu)^2\right\} d\mu}$$

From the equation above, it is evident that the maximum value of the Bayes factor is obtained when $\mathbf{y}(\ell) = \mu_0$, and the maximum value is

$$\boxed{\sup_{\mathbf{y}(\ell) \in D} B_{01}^N(\mathbf{y}(\ell)) = \frac{1}{\sqrt{2\pi}\sigma_0}}.$$

By using the result above, we can compute the Bayes factor bound

$$\begin{aligned} \bar{B}_{10}^I &= B_{10}^N(\mathbf{y}) \times \sup B_{01}^N(\mathbf{y}(\ell)) \\ &= \frac{\int_0^\infty \left(\frac{1}{\sqrt{2\pi}\sigma_0}\right)^n \exp\left\{-\frac{1}{2\sigma_0^2} \sum [(\mathbf{y}(\ell) - \bar{\mathbf{y}})^2 + n(\bar{\mathbf{y}} - \mu)^2]\right\} d\mu}{\left(\frac{1}{\sqrt{2\pi}\sigma_0}\right)^n \exp\left\{-\frac{1}{2\sigma_0^2} \sum (\mathbf{y}(\ell) - \mu_0)^2\right\}} \times \sup B_{01}^N(\mathbf{y}(\ell)) \\ &= \frac{\sqrt{2\pi} \frac{\sigma_0}{\sqrt{n}}}{\exp\left\{\frac{-1}{2\sigma_0^2} n(\bar{\mathbf{y}} - \mu_0)^2\right\}} \sup B_{01}^N(\mathbf{y}(\ell)) = \frac{\sqrt{2\pi} \frac{\sigma_0}{\sqrt{n}}}{\exp\left\{\frac{-1}{2\sigma_0^2} n(\bar{\mathbf{y}} - \mu_0)^2\right\}} \frac{1}{\sqrt{2\pi}\sigma_0} \end{aligned}$$

Simplifying the equation above yields

$$\boxed{\bar{B}_{10}^I = \frac{\exp\left\{\frac{1}{2\sigma_0^2} n(\bar{\mathbf{y}} - \mu_0)^2\right\}}{\sqrt{n}}}$$

Similarly, we can obtain the IBF lower bound

$$\boxed{\underline{B}_{01}^I = \sqrt{n} \left(\exp - \left\{ \frac{1}{2\sigma_0^2} n(\bar{\mathbf{y}} - \mu_0)^2 \right\} \right)}.$$

5.2.3 Normal mean hypothesis test with unknown variance.

Let us consider a scenario where the random vector $\mathbf{Y} = (Y_1, \dots, Y_n)$ is composed of independent and identically distributed (i.i.d) random variables Y_i , following a normal distribution with mean μ_0 and standard deviation σ_0 under the null hypothesis M_0 , and a normal distribution with mean μ and standard deviation σ_1 under the alternative hypothesis M_1 .

$$H_0 : \mu = \mu_0 \text{ vs } H_1 : \mu \neq \mu_0, \sigma \text{ unknown}$$

We adopt non-informative priors, specifically $\pi_0^N(\sigma_0) = 1/\sigma_0$ and $\pi_1^N(\mu, \sigma_1) = 1/\sigma_1^2$. It is important to note that the choice of π_1^N differs from that of π_0^N due to a seminal work by Berger, Pericchi, and Varshavsky in 1998 [5], wherein they calculated the marginal of $\mathbf{y}(\ell)$ using $\pi_1^N(\mu, \sigma_1) = 1/\sigma_1$. In this context, the expression for the density function under the alternative hypothesis takes the form:

$$m_1(\mathbf{y}(\ell)) = \frac{1}{2|y_1 - y_2|}$$

where $\mathbf{y}(\ell)$ represents the observed data. However, we face a challenge in evaluating the Bayes Factor due to the nature of the above density function. Specifically, the infimum and supremum of this function can potentially become either zero or infinity, leading to indeterminate bounds for the Bayes Factor. To circumvent this issue and ensure well-defined bounds, we opt for the aforementioned choice of priors. This strategic choice of priors enables us to derive meaningful marginals and subsequently evaluate the Bayes Factor.

$$m_0^N(\mathbf{y}(\ell)) = \frac{1}{2\pi(y_1^2 + y_2^2)}, \quad m_1^N(\mathbf{y}(\ell)) = \frac{1}{\sqrt{\pi}(y_1 - y_2)^2}$$

and the Bayes factor is

$$B_{01}(\mathbf{y}(\ell)) = \frac{m_0(\mathbf{y}(\ell))}{m_1(\mathbf{y}(\ell))} = \frac{(y_1 - y_2)^2}{2\sqrt{\pi}(y_1^2 + y_2^2)}$$

To find the IBF bounds, we found the supremum of the Bayes factor above

$$\sup_{\mathbf{y}(\ell) \in D} B_{01}(\mathbf{y}(\ell)) = \frac{m_0(\mathbf{y}(\ell))}{m_1(\mathbf{y}(\ell))} = \frac{(y_1 - y_2)^2}{2\sqrt{\pi}(y_1^2 + y_2^2)} = \frac{1}{\sqrt{\pi}}$$

The IBF upper bound can be obtained as

$$\bar{B}_{10} = B_{10}(\mathbf{y}) \sup B_{01}(\mathbf{y}(\ell)) = \frac{m_1(\mathbf{y})}{m_0(\mathbf{y})} \frac{1}{\sqrt{\pi}}$$

To find the upper bound, we must calculate the marginals m_0 and m_1 . To obtain m_0 , we must solve the following integral

$$m_0(\mathbf{y}) = \int_0^\infty \left(\frac{1}{\sqrt{2\pi}}\right)^n \exp\left\{\frac{-1}{2\sigma_0^2}[\sum(y_i - \mu_0)^2]\right\} \frac{1}{\sigma_0^{(n+1)}} d\sigma_0$$

Consider the change of variable $\beta = \frac{1}{2}[\sum(y_i - \mu_0)^2]$, $\frac{1}{\sigma^2} = \eta$, then

$$\frac{-2}{\sigma^3} d\sigma = d\eta \implies d\sigma = -\frac{\sigma^3}{2} d\eta$$

$$\frac{1}{\eta^{1/2}} = \sigma \implies d\sigma = -\eta^{3/2}/2 d\eta$$

$$m_0(\mathbf{y}) = \frac{1}{2} \left(\frac{1}{\sqrt{2\pi}}\right)^n \int_0^\infty \exp\{-\beta\eta\} \eta^{n/2-1} d\eta = \frac{1}{2} \left(\frac{1}{\sqrt{2\pi}}\right)^n \frac{\Gamma(n/2)}{\beta^{n/2}}$$

Similarly, we can calculate m_1

$$m_1(\mathbf{y}) = \left(\frac{1}{\sqrt{2\pi}}\right)^n \int_0^\infty \frac{1}{\sigma_1^{n+2}} \int_{-\infty}^\infty \exp\left\{-\frac{1}{2\sigma_1^2}(\sum(y_i - \mu)^2)\right\} d\mu d\sigma_1$$

Adding and subtracting $\bar{\mathbf{y}}$ inside the exponential function yields

$$\sum(y_i - \mu)^2 = \sum((y_i + \bar{\mathbf{y}}) - (\mu + \bar{\mathbf{y}}))^2 = \sum(y_i - \bar{\mathbf{y}})^2 + n(\mu - \bar{\mathbf{y}})^2$$

Inserting this result inside the integral with respect to μ results in

$$\begin{aligned} \int_{-\infty}^\infty \exp\left\{-\frac{1}{2\sigma_1^2}(\sum(y_i - \mu)^2)\right\} d\mu &= \exp\left\{\frac{-1}{2\sigma_1^2} \sum(y_i - \bar{\mathbf{y}})^2\right\} \int_{-\infty}^\infty \exp\left\{\frac{-1}{2(\sigma_1/\sqrt{n})^2}(\mu - \bar{\mathbf{y}})^2\right\} d\mu \\ &= \sigma_1 \sqrt{2\pi/n} \exp\left\{-\frac{1}{2\sigma_1^2}(\sum(y_i - \bar{\mathbf{y}})^2)\right\} \end{aligned}$$

To calculate m_1 , we insert this result to solve the integral with respect to σ_1

$$\frac{1}{\sqrt{n}} \left(\frac{1}{\sqrt{2\pi}} \right)^{n-1} \int_0^\infty \frac{1}{\sigma_1^{n+1}} \exp\left\{-\frac{1}{2\sigma_1^2} (\sum (y_i - \bar{y})^2)\right\} d\sigma_1$$

Let $\beta = \frac{\sum (y_i - \bar{y})^2}{2}$ and $\eta = \frac{1}{\sigma_1^2}$, then

$$\frac{1}{2\sqrt{n}} \left(\frac{1}{\sqrt{2\pi}} \right)^{n-1} \int_0^\infty \eta^{n/2-1} \exp\{-\beta\eta\} d\eta$$

This integral is the kernel of a Gamma distribution which results in

$$m_1(\mathbf{y}) = \frac{1}{2\sqrt{n}} \left(\frac{1}{\sqrt{2\pi}} \right)^{n-1} \frac{\Gamma(n/2)}{\left(\frac{\sum (y_i - \bar{y})^2}{2} \right)^{n/2}}$$

After obtaining m_0 and m_1 , we can calculate the Bayes Factor upper bound

$$\bar{B}_{10} = \sqrt{2/n} \left(\frac{\sum (y_i - \mu_0)^2}{\sum (y_i - \bar{y})^2} \right)^{n/2}.$$

Similarly, the Bayes Factor Lower bound is

$$\underline{B}_{01} = \sqrt{n/2} \left(\frac{\sum (y_i - \bar{y})^2}{\sum (y_i - \mu_0)^2} \right)^{n/2}.$$

CHAPTER 6

THE EXPONENTIAL MODEL HYPOTHESIS TEST

In the exploration of statistical hypothesis testing, this chapter embarks on a comprehensive investigation into exponential hypothesis testing. Specifically, we delve into the comparison of two hypotheses $H_0 : \lambda = \lambda_0$ and $H_1 : \lambda \neq \lambda_0$ where λ is the parameter for an exponential density $f(y|\lambda) = \lambda e^{-\lambda y}$. We will employ several robust methodologies to assess this hypothesis test from different but consistent approaches.

Our focus centers on evaluating and contrasting several pivotal methodologies: the EP Approach, the $-ep \log(p)$ lower bound, the Universal Robust Bound for the Intrinsic Bayes Factor, and the SP Approach. These methodologies serve as our guiding compasses in the pursuit of discerning their effectiveness in hypothesis validation. We rely on simulations to shed light on the performance and consistency of these methodologies. By employing these simulations, we aim to craft illustrative charts that will serve as visual aids, elucidating the comparative efficacy and nuances of each approach. This chapter unfolds as a meticulous exploration, aiming not only to identify the most effective methodology but also to unravel the intricacies underlying their application.

The exponential density is defined as

$$f(y|\lambda) = \lambda e^{-\lambda y}, \text{ where } y \geq 0 \text{ and } \lambda > 0$$

The likelihood, Jeffrey's prior, and posterior distribution for an exponential density with parameter λ are given by

$$L(\lambda|\mathbf{y}) = \lambda^n \exp(-\lambda \sum_{i=1}^n y_i), \quad \pi^N(\lambda) \propto \frac{1}{\lambda}, \quad \pi(\lambda|\mathbf{y}) = \frac{L(\lambda|\mathbf{y})\pi_J(\lambda)}{m(\mathbf{y})}$$

After calculating the likelihood and Jeffrey's prior for the exponential density we can easily obtain the following marginal

$$m(\mathbf{y}) \propto \int_0^\infty \lambda^{(n-1)} \exp(-\lambda \sum_{i=1}^n y_i) d\lambda = \frac{(\sum_{i=1}^n y_i)^n}{\Gamma(n)}$$

Finally, we can obtain the posterior distribution for λ

$$\pi(\lambda|\mathbf{y}) \propto \frac{\Gamma(n)\lambda^{n-1} \exp\left(-\lambda \sum_{i=1}^n y_i\right)}{(\sum_{i=1}^n y_i)^n}$$

In the following sections, we'll tackle a hypothesis test:

$$H_0 : \lambda = \lambda_0 \text{ vs. } H_1 : \lambda \neq \lambda_0$$

We'll use different methods to calculate Bayes Factors, bounds, and priors. This helps us compare traditional approaches with our new methods, giving us a clearer understanding of their outcomes.

6.1 The Intrinsic Bayes Factor Dynamic and Universal Lower Bounds

In this section, we're computing our new bound for the Intrinsic Bayes Factor. To derive this bound, we start by examining the Bayes factor for the minimal training sample:

$$B_{01}^N(y(\ell)) = \frac{\mathbf{y}(\ell)\lambda_0}{\exp^{\lambda_0 \mathbf{y}(\ell)}}$$

Additionally, we need the full Bayes factor for this test

$$B_{01}^N(\mathbf{y}) = \frac{\lambda_0^n \exp(-\lambda_0 \sum y_i)}{\int_0^\infty \lambda^n \exp(-\lambda \sum y_i) \frac{1}{\lambda} d\lambda} = \frac{\lambda_0^n \exp(-\lambda_0 \sum y_i) (\sum y_i)^n}{\Gamma(n)}$$

By minimizing the first Bayes factor, we obtain the Empirical IBF lower bound

$$\underline{B}_{01}^{I*} = \frac{(\sum y_i)^n \lambda_0^{n-1}}{\mathbf{y}^*(\ell) \Gamma(n)} \exp(\lambda_0 [\mathbf{y}^*(\ell) - \sum y_i]), \quad \mathbf{y}^*(\ell) = \arg \min_{y_i \in \{y_1, \dots, y_n\}} \left\{ \frac{\mathbf{y}_i \lambda_0}{\exp^{\lambda_0 y_i}} \right\}$$

Our next goal is to obtain the theoretical IBF bound, and to achieve this we need to find $\sup_{y(\ell) \in \mathcal{D}} B_{01}^N(y(\ell))$. This maximization can be attained by solving the following equation

$$\begin{aligned} \frac{d}{dy(\ell)} [\log(B_{01}^N(y(\ell)))] &= \frac{d}{dy(\ell)} [\log(y(\ell)) + \log(\lambda_0) - \lambda_0 y(\ell)] &&= 0 \\ &= \frac{1}{y^*(\ell)} + \lambda_0 &&= 0 \\ y^*(\ell) &= \frac{1}{\lambda_0} \end{aligned}$$

To verify that this is a maximum, we take the second derivative with respect to $y(\ell)$

$$\frac{d^2}{dy(\ell)^2} [\log(B_{01}^N(y(\ell)))] = \frac{-1}{y(\ell)^2} < 0, \forall y(\ell) \in \mathbb{R}$$

Therefore, $y^*(\ell)$ the supremum in $(0, \infty)$ and $B_{01}^N(y^*(\ell)) = 1/e$. By using these results, we can obtain the following Intrinsic Bayes Factor upper bound

$$\boxed{\bar{B}_{10}^I = 1/e \frac{\Gamma(n) \exp(\lambda_0 \sum y_i)}{(\sum y_i)^n \lambda_0^n}}$$

Similarly, the IBF lower bound is given by

$$\boxed{B_{01}^I = \frac{e(\sum y_i)^n \lambda_0^n}{\Gamma(n) \exp(\lambda_0 \sum y_i)}}$$

6.2 The SP Bayes Factors

In this section, our focus lies in computing the SP priors and the SP Bayes factor for the exponential hypothesis test. We aim to contrast the outcomes from the SP approach with the intrinsic Bayes factor bounds derived earlier. To bolster our findings, we'll conduct simulations and generate illustrative charts. These visual aids will fortify and elucidate our conclusions.

The theoretical SP Priors for the exponential model for B_{10} are given by

$$\boxed{\pi_0^{SP}(\lambda) = 1, \pi_1^{SP}(\lambda) = \frac{1}{\lambda_0} \exp^{-\lambda/\lambda_0}, \lambda > 0}$$

It's noteworthy that $\pi_1^{SS}(\lambda)$ follows an exponential distribution with the parameter $1/\lambda_0$. Let's assume, without loss of generality, that $y_1 = 1/\lambda_0$. The SP marginals for the hypothesis test $H_0 : \lambda = \lambda_0$ versus $H_1 : \lambda \neq \lambda_0$ are

$$\begin{aligned} m_0^{SP}(\mathbf{y}) &\propto \lambda_0^n e^{-\lambda_0 \sum_{i=1}^n y_i} \\ &= \lambda_0^n e^{-\lambda_0 \sum_{i=1}^n y_i} \end{aligned}$$

$$\begin{aligned} m_1^{SP}(\mathbf{y}) &\propto \frac{1}{\lambda_0} \int_0^\infty \lambda^n e^{-\lambda(\sum_{i=1}^n y_i + \frac{1}{\lambda_0})} \\ &= \frac{1}{\lambda_0} [\Gamma(n+1) (\sum_{i=1}^n y_i + \frac{1}{\lambda_0})^{-(n+1)}] \\ &= \frac{1}{\lambda_0} [\Gamma(n+1) (\sum_{i=1}^n y_i + \frac{1}{\lambda_0})^{-(n+1)}] \end{aligned}$$

Utilizing these marginals leads to the results of the SP Bayes factor as follows:

$$B_{10}^{SP} = \frac{\Gamma(n+1) e^{\lambda_0 \sum_{i=1}^n y_i}}{(\sum_{i=1}^n y_i + \frac{1}{\lambda_0})^{n+1} \lambda_0^{n+1}}$$

In order to verify the validity of our calculations, and to compare the numerical values of B_{10}^{SP} with \bar{B}_{10}^I , a simulation was conducted and the corresponding chart was obtained. The simulation was implemented using the following code snippet:

R Program

```

1 b10_theoretical = function(n,h0,lambda=lambda0) {
2   y = c(); b10_t = c(); b10_sp = c(); s = c(); e = exp(1)
3   for (i in seq(1,n,1)) {
4     y = c(y, rexp(1, rate=1/lambda))
5     s = sum(y)
6     val = -log(e,e) + -n*log(s,e) - n*log(h0,e) + log(gamma(n),e)+ h0*s
7     b10_t = c(b10_t, val)
8     val = log(gamma(n+1),e) + h0*s - (n+1)*log(s + 1/h0,e) - (n+1)*log(h0,e)
9     b10_sp = c(b10_sp, val)
10  }
11  plot(seq(1,n,1), b10_t, type="l", col="red",
12       xlab="Number of samples",
13       ylab="log(BF10)", ylim=c(min(b10_t, b10_sp), max(b10_t, b10_sp)),
14       main="Theoretical SS Bayes Factor (Red) vs SP Bayes Factor (Blue) (H0 TRUE)")
15  lines(b10_sp, col="blue")
16 }

```

Theoretical SS and SP Bayes Factor Lower Bounds comparison under H_0




Figure 6-1: \bar{B}_{10}^{SS} vs B_{10}^{SP} under H_0 for increasing sample size and exponential data.

This comparison serves to analyze the numerical relationship between B_{10}^{SP} and \bar{B}_{10}^I , thereby contributing to the evaluation and understanding of their respective characteristics. The obtained chart provides a visual representation of the convergence of the \bar{B}_{10}^I to B_{10}^{SP} as the number of samples increases.

6.3 The EP Bayes Factors

The exponential Expected posterior prior under the alternative (EP-Prior) for the exponential hypothesis test $H_0 : \lambda = \lambda_0$ vs $H_1 : \lambda \neq \lambda_0$ can be calculated as

$$\pi_1^{EP}(\lambda) = \frac{\sum_{i=1}^L \pi(y^{(i)}|\theta_i)}{n} = \frac{\sum_{i=1}^n y_i e^{-\lambda y_i}}{n}$$

It can be verified that the prior in question integrates unity with respect to λ by performing integration overall positive values of λ . Additionally, the SP prior, which is an exponential distribution with a parameter of $1/\lambda_0$, also integrates one.

The objective of this study is to compare the SP prior with the EP prior under different sample sizes, specifically for the cases of $n = 10$ and $n = 100$. To achieve this, simulations were conducted utilizing the following R code:



R Program

```

1 prior = function(data, lambda0, lambda) {
2   data = c(); object = c(); lambda0 = 0; ep1 = c(); ss1 = rep(0, length(lambda)); ss1 = (1/lambda0)*exp(-
      lambda/lambda0)
3   for (j in 1:length(lambda)) {
4     epnew = 0
5     for (i in 1:length(data)) { epnew = epnew + data_[i]*exp(-data[i]*lambda[j]) }
6     ep1 = c(ep1, epnew)
7   }
8   object$ss = ss1; object$ep = ep1/length(data)
9   return(object)
10 }
11 data = rexp(n=100, r=1); lambda0=1; lambda = seq(0.1, 2, 0.1)
12 object = prior(data = data, lambda0=lambda0, lambda=lambda);
13 ss = exp(-lambda); ep = object$ep
14 plot(lambda, ep, col="red", ylim=c(min(ss, ep), max(ss, ep)), type="b", ylab="Prior",
15 main="EP Prior under H1 (Red) / SP Prior under H1 (Blue)");
16 lines(lambda, ss, col="blue")

```

By employing these simulations, we aim to investigate and evaluate the performance of the SP prior relative to the EP prior in the context of various sample sizes. The obtained results will contribute to enhancing our understanding and analysis of the behavior and characteristics of these priors.

(a) Exponential simulation with 10 samples for $\lambda_0 = 1$ (b) Exponential simulation with 100 samples for $\lambda_0 = 1$ **Figure 6–2:** Simulations with different number of samples for the exponential SP and EP Priors for $\lambda_0 = 1$

The exponential EP-Bayes factor for $\lambda = \lambda_0$ vs $\lambda \neq \lambda_0$ can be obtained by calculating the following expression

$$B_{10}^{EP}(\mathbf{y}) = \frac{m_1^{EP}(\mathbf{y})}{m_0^{EP}(\mathbf{y})} = \frac{\frac{1}{n} \int_0^{\infty} \lambda^n \exp\{-\lambda \sum_{i=1}^n y_i\} \sum_{i=1}^n y_i \exp\{-\lambda y_i\} d\lambda}{\lambda_0^n \exp\{-\lambda_0 \sum_{i=1}^n y_i\}}$$

Our objective is to resolve the integral mentioned earlier. To facilitate this, let's define $S = \sum_{i=1}^n y_i$. Subsequently, the integral transforms into:

$$\sum_{i=1}^n y_i \int_0^{\infty} \lambda^n \exp\{-\lambda(S + y_i)\} d\lambda$$

Consider the change of variable $\alpha = n + 1$ and $\beta_i = \frac{1}{S+y_i}$, which leads to the transformation:

$$\sum_{i=1}^n y_i \int_0^{\infty} \lambda^n \exp\{-\lambda(S+y_i)\} d\lambda = \sum_{i=1}^n y_i \int_0^{\infty} \lambda^{\alpha-1} \exp\{-\lambda/\beta_i\} d\lambda$$

This new form represents the kernel of a Gamma density with parameters α and β_i , resulting in the integral:

$$\Gamma(n+1) \sum_{i=1}^n y_i \left(\frac{1}{S+y_i}\right)^{n+1}$$

Hence, the EP Bayes Factor, in this case, becomes:

$$B_{10}^{EP}(\mathbf{y}) = \frac{\frac{1}{n} \Gamma(n+1) \sum_{i=1}^n y_i \left(\frac{1}{S+y_i}\right)^{n+1}}{\lambda_0^n \exp\{-\lambda_0 \sum_{i=1}^n y_i\}}$$

The principal objective of this numerical experiment is to compare the EP and SP Bayes factors, previously derived, by creating 100 simulations, each consisting of 100 samples. The process entails averaging the Bayes factor for each sample size, ranging from 1 to 100. Furthermore, we intend to plot the logarithm of each Bayes factor on the same axis to discern their respective behaviors under two distinct scenarios: when the null hypothesis is valid and when it is not. The forthcoming R program has been devised to generate the anticipated output.

R Program

```

1 #Initialize vectors and variables
2 finallogbf10 = c(); finallogbf10sp=c(); nsim = 100; n = 100;
3 data = c(); s = 0; sumdata = 0; logbf10=c(); logbf10sp=c();
4 #H0 TRUE
5 true_lambda = 1; lambda0 = 1;
6 #Use this for H0 FALSE
7 #true_lambda = 1; lambda0 = 0.3;
8 for (k in 1:nsim) {
9   #Generate random samples
10  data =rexp(n=n,rate=true_lambda)
11  for (i in 1:n) {
12    s = sum(data[1:i])
13    for (j in 1:i) { sumdata = sumdata + data[j]*((1/(s+data[j]))^(i+1)) }
14    #EP marginals
15    m1_ep = (gamma(i+1)*sumdata)/i; m0_ep = (lambda0)^i*exp(-lambda0*s)
16    #Log EP Bayes factor
17    logbf10 = c(logbf10,log(m1_ep) - log(m0_ep))
18    #SP Marginals
19    m1_sp = gamma(i+1)*exp(lambda0*s); m0_sp = ((s + 1/lambda0)^(i+1)) *(lambda0^(i+1))
20    #Log SP Bayes Factor
21    logbf10sp = c(logbf10sp,log(m1_sp) - log(m0_sp))
22    sumdata = 0
23  }
24  #Storing simulations
25  finallogbf10 = cbind(finallogbf10,logbf10); finallogbf10sp = cbind(finallogbf10sp,logbf10sp)
26 }
27 logbf10_avg = rep(0,n); logbf10sp_avg = rep(0,n);
28 #Averaging Bayes factors for all the simulations
29 for (i in 1:n) {
30   logbf10_avg[i] = mean(finallogbf10[i,]); logbf10sp_avg[i] = mean(finallogbf10sp[i,])
31 }
32 x = 1:n
33 y1 = logbf10_avg[x]; y2 = logbf10sp_avg[x]
34 if (lambda0 == true_lambda) { h0 =TRUE } else { h0 = FALSE }
35 plot(x,y1,type="l",col="red",ylim=c(min(y1,y2),max(y1,y2)), main="log(EP Bayes Factor 10) (Red) vs log(SP Bayes
    Factor 10) (Black)",xlab="Sample size",ylab="Log Bayes Factor 10")
36 lines(x,y2,type="l",col="black")
37 mtext(paste("H0 is ",h0," / Number of simulations: ",nsim," / Number of samples: ",n," / lambda = ",true_lambda
    , " / lambda_0 = ",lambda0, sep=""),side=3)

```

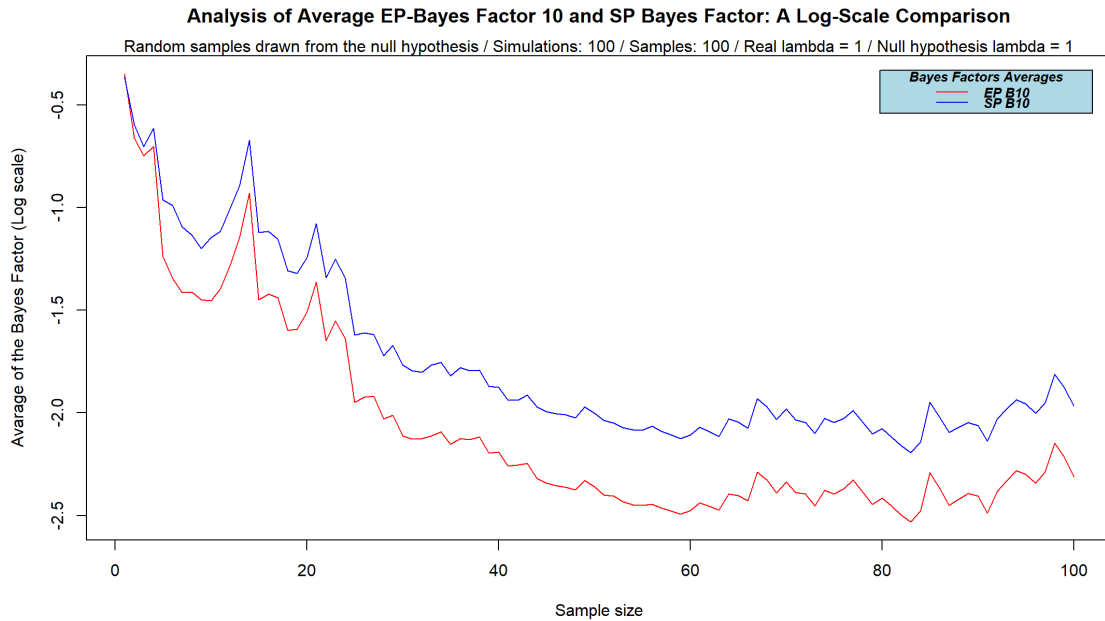


Figure 6–3: Logarithm of the SP and EP Bayes factor using one-hundred exponential simulations (100 samples each) for $\lambda_0 = 1, H_0$ True

Figure 6–3 demonstrates the consistency of both the SP and EP Bayes factors. The figure illustrates that both Bayes factors approach zero when calculating the Bayes factor of the alternative hypothesis over the null hypothesis. This observation is expected since the data used for the calculations are random samples generated from the null hypothesis. The decreasing trend of the Bayes factors indicates their convergence towards zero, indicating consistency in the evidence provided by both methods.




Figure 6-4: Logarithm of the SP and EP Bayes factor using one-hundred exponential simulations (100 samples each) for $\lambda_0 = 1, H_0$ False

Figure 6-4 demonstrates the consistency of both the SP and EP Bayes factors. The figure illustrates that both Bayes factors increase in value with the sample size when calculating the Bayes factor of the alternative hypothesis over the null hypothesis. This observation is expected since the data used for the calculations are random samples generated from the alternative hypothesis. The increasing trend of the Bayes factors indicates consistency in the evidence provided by both methods.

6.4 The $-ep\log p$ lower bound

In this section we revisit the Sellke, Bayarri, and Berger (2001) [13] lower bound for the Bayes factor that employs the classical p-value of the hypothesis test, represented as $-ep\log(p) \leq B_{01}$. We will perform a simulation to compare this p-value approximation of the Bayes factor to contrast it with the intrinsic Bayes factor lower bound for the exponential test. In an attempt to compute the $-ep\log(p)$ lower bound of B_{01} , we determined the likelihood ratio for the hypotheses H_0 against H_1 . The approximation of the p-value was accomplished through the utilization of Wilk's theorem and the chi-squared distribution. The graphical representation provided demonstrates the lower bound and the $-ep\log(p)$ in tandem as the sample size increases.

The forthcoming illustration elucidates how both the empirical and theoretical bounds converge toward the appropriate decision as the sample size escalates. This is demonstrated by drawing samples from an exponential distribution characterized by a known rate parameter, while simultaneously altering the null hypothesis.

R Program

```

1
2 exponential_simulations=function(n_sim,n_data,h0,rate) {
3   lambda = h0; sim = c(); bfbound = c(); bfssbound = c(); pval_seq= c(); epbound = c(); ssbound = c();
4   for (j in 1:n_sim) {
5     for (n in 1:n_data) {
6       sim = c(sim, rexp(n=1,rate=rate))
7       lrt_null = n*log(lambda) - lambda*sum(sim)
8       lrt_alt = n*(log(n) - log(sum(sim)) - 1)
9       lrt = exp(lrt_null-lrt_alt)
10      ts = -2*log(lrt)
11      #p-value approximation wilks theorem
12      pval = 1-pchisq(ts,1)
13      pval_seq = c(pval_seq,pval)
14      #selke bound
15      if (pval > (1/exp(1))) { bfbound = c(bfbound,1) }
16      else { bfbound = c(bfbound,-exp(1)*pval*log(pval)) }
17      bfssbound = c(bfssbound,(exp(1)*(sum(sim)*lambda)^n)/gamma(n)*exp(-lambda*sum(sim)))
18    }
19    epbound = cbind(epbound,bfbound); ssbound = cbind(ssbound,bfssbound)
20  }
21  epbound_avg = c(); ssbound_avg = c()
22  for (i in 1:n_data) {
23    epbound_avg = c(epbound_avg,mean(epbound[i,]))
24    ssbound_avg = c(ssbound_avg,mean(ssbound[i,]))
25  }
26  object = c(); object$epbound_avg = epbound_avg; object$ssbound_avg = ssbound_avg
27  return(object)
28 }
29
30 #100 Simulations and 100 Samples; H0 True
31 n_sim = 100; n_data = 100; h0 = 1; rate = 1
32
33 simulations = exponential_simulations(n_sim=n_sim,n_data = n_data,h0=h0,rate=rate)
34
35 par(mfrow=c(1,2))
36 plot(1:n_data,simulations$epbound_avg,type="l",col="red",xlab="Number of samples",ylab="Average of the Lower
      bounds",ylim=c(min(simulations$epbound_avg,simulations$ssbound_avg),max(simulations$ssbound_avg,
      simulations$epbound_avg)),
37      main="Comparison between the -eplogp and IBF lower bounds")
38 lines(simulations$ssbound_avg,col="blue")
39
40 mtext(paste("Simulations: lambda = ",rate, " / Null hypothesis: lambda = ",h0, sep=""),side=3)
41 legend(0, 10, legend=c("-eplogp", "SS"),col=c("red", "blue"), lty=c(1,1), cex=0.8,title="Lower bounds", text.
      font=4, bg='lightblue')

```

By employing the aforementioned code, a simulation was carried out to calculate both the empirical and theoretical lower bounds of B_{01}^I for the exponential distribution. This was executed to test the hypothesis H_0 against H_1 under two distinct circumstances: when H_0 holds true and when it does not. The objective was to examine the convergence of the Bayes factor lower bound toward the accurate decision as the sample size expands. The outcomes of this simulation are presented in the following figures.




Figure 6-5: The $-ep\log(p)$ and B_{01}^I using 100 exponential simulations (100 samples each) for $\lambda_0 = 1$ (Null hypothesis true), and 100 exponential simulations with 10 samples each for $\lambda_0 = 1$ and the simulations rate parameter $r = 5$ (false null hypothesis)

In **Figure 6-5** we illustrated the consistency of the IBF lower bound in both scenarios. On the other hand, the second figure (right) illustrates that both Bayes factor's lower bounds approach zero when calculating the Bayes factor of the null hypothesis over the alternative hypothesis. This observation is expected since the data used for the calculations are random samples generated from the alternative hypothesis. The decreasing trend of the lower bounds indicates their convergence towards zero, indicating consistency in the evidence provided by both methods under the alternative, however, the $-ep\log(p)$ bound does not provide decisive evidence to accept the null hypothesis when we draw random samples from it, because the evidence is bounded by one.

6.5 Conclusions

Our research indicates that the Theoretical SP Bayes Factor, derived using the SP-Priors, is congruent with the Intrinsic Bayes factor Bound \bar{B}_{10}^I , albeit with the inclusion of an additional observation being the supremum of $B_{01}^N(y(\ell))$. We demonstrated that the SS-prior yields comparable results and that this methodology facilitates the generation of a lower bound for the theoretical SS Bayes factor.

As the sample size increases, the discrepancy between these two methods diminishes, given that we are more likely to encounter samples closer to the supremum (under H_0). In the context of comparing the EP and SP Priors for H_1 , it becomes evident that the Theoretical SP prior remains unchanged for any sample, while the EP-prior fluctuates in accordance with the samples.

Moreover, it can be inferred from the simulations and the chart that the SP prior assigns a higher probability to the null hypothesis $\lambda_0 = 1$, owing to its wider area from 0 to 1 compared to the EP-prior for both $n = 10$ and $n = 100$. Both the EP-Bayes factor and the SP-Bayes factors exhibit similar patterns of convergence and values. However, the SP-Bayes factor presents higher values under the null and lower values under the alternative.

The consistencies observed in the results validate the results of the SP-Bayes factors and confirm the accuracy of the computations performed in this study.

CHAPTER 7

LINEAR REGRESSION AND ANOVA HYPOTHESIS TESTING WITH BAYES FACTORS

In the realm of statistical modeling, the avenues of linear regression and analysis of variance (ANOVA) stand as cornerstones, offering valuable insights into relationships between variables and differences among groups. Traditional frequentist approaches have long governed these methodologies, relying on p-values and significance testing. However, the burgeoning field of Bayesian statistics introduces a paradigm shift, emphasizing coherent inference and robust evidence assessment through Bayes factors.

This chapter embarks on an exploration of linear regression and ANOVA within the Bayesian framework. Departing from traditional methods, we delve into the application of Bayes factor approximations, building upon the discussions from our previous chapters. Moreover, we extend the foundations laid by Smith and Spiegelhalter, broadening our insights by calculating these approximations through the utilization of a generalized Jeffrey's prior.

ANOVA, a stalwart technique in comparing group means, undergoes a transformation in this chapter as we explore its foundations through Bayesian inference. The comparison of multiple groups finds new light as Bayes factors unveil nuanced insights into the evidence supporting differing hypotheses.

7.1 Linear regression

Linear regression is a statistical model that is commonly utilized in several disciplines, ranging from economics to engineering. At its core, linear regression is a conditional model where the outcome variable is predicated on a linear combination of the

predictor variables, in conjunction with an unobserved error term that introduces variability into the relationship between the input and output variables.

The fundamental form of a simple linear regression model is as follows:

$$Y = \theta_0 + \theta_1 X + \varepsilon \quad (7.1)$$

where Y is the dependent or outcome variable, X is the independent or predictor variable, θ_0 is the Y-intercept, which represents the expected value of Y when all X are 0, θ_1 is the slope of the regression line, indicating the degree to which Y changes for each unit change in X , and ε is the error term, embodying the difference between the observed and predicted values of Y .

The goal of linear regression analysis is to estimate the coefficients θ_0 and θ_1 that minimize the sum of the squared residuals, thus providing the "best fit" line for the observed data.

In a multiple linear regression scenario, the model is expanded to include more than one independent variable:

$$Y = \theta_0 + \theta_1 X_1 + \theta_2 X_2 + \dots + \theta_p X_p + \varepsilon \quad (7.2)$$

where p represents the number of predictor variables. Here, each θ_i (for $i = 1, 2, \dots, p$) signifies the change in the expected value of Y for each unit change in the corresponding predictor X_i , holding all other predictors constant.

Linear regression models assume linearity, independence, homoscedasticity (constant variance), and normality of residuals. Violations of these assumptions may necessitate the application of more complex models or transformations of the data. Linear regression serves as a fundamental building block in understanding more intricate statistical models.

7.2 Two nested normal-linear models

Consider the comparison of two nested normal-linear models $M_0 \subset M_1$, defined as in Smith and Spiegelhalter (1980) [14] by

$$M_i : \mathbf{y} \sim N(\mathbf{A}_i \boldsymbol{\theta}_i, \sigma^2 \mathbf{I}_n), \quad i = 0, 1$$

where \mathbf{A}_i is a full rank p_i known matrix, \mathbf{y} is a vector with dimension n , $\boldsymbol{\theta}_i = (\theta_{i1}, \dots, \theta_{ip_i})$ is a vector of p_i unknown parameters and σ^2 is unknown. This can be written in matrix notation as

$$\mathbf{y} = \mathbf{A}_i \boldsymbol{\theta}_i + \boldsymbol{\varepsilon}_i, \quad i = 0, 1$$

where $\boldsymbol{\varepsilon}_i \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_n)$ for $i = 0, 1$. Let

$$\hat{\boldsymbol{\theta}}_i = (\mathbf{A}_i^T \mathbf{A}_i)^{-1} \mathbf{A}_i^T \mathbf{y} \text{ and } \mathbf{R}_i = |\mathbf{y} - \mathbf{A}_i \hat{\boldsymbol{\theta}}_i|^2$$

denote the least squares estimator of $\boldsymbol{\theta}_i$ and the residual sum of squares, respectively. The Bayes factor in this case is given by

$$B_{01} = \frac{\int \int p(\mathbf{y} | \mathbf{A}_0, \boldsymbol{\theta}_0, \sigma) p(\boldsymbol{\theta}_0, \sigma | \mathbf{A}_0) d\boldsymbol{\theta}_0 d\sigma}{\int \int p(\mathbf{y} | \mathbf{A}_1, \boldsymbol{\theta}_1, \sigma) p(\boldsymbol{\theta}_1, \sigma | \mathbf{A}_1) d\boldsymbol{\theta}_1 d\sigma}$$

Smith and Spiegelhalter showed that for the case when $i = 0, 1$, $p(\boldsymbol{\theta}_i, \sigma | \mathbf{A}_i)$ has an improper limit form representing vague prior information for all parameters inside each linear model. The improper limiting version of the normal-inverse- χ^2 conjugate prior can be written in the form

$$p(\boldsymbol{\theta}_i, \sigma | \mathbf{A}_i) = p(\boldsymbol{\theta}_i | \mathbf{A}_i, \sigma) p(\sigma) = c_i (2\pi\sigma^2)^{-p_i/2} \sigma^{-1}$$

and the likelihood is given by

$$p(\mathbf{y} | \mathbf{A}_i, \boldsymbol{\theta}_i, \sigma) = \left(\frac{1}{\sqrt{2\pi\sigma}} \right)^n \exp \left\{ -\frac{1}{2\sigma^2} \left[\mathbf{R}_i(\mathbf{y}) + (\boldsymbol{\theta}_i - \mathbf{y})^T \mathbf{A}_i^T \mathbf{A}_i (\boldsymbol{\theta}_i - \mathbf{y}) \right] \right\}.$$

The Bayes factor in this case is given by:

$$B_{01} = \frac{c_0}{c_1} [|\mathbf{A}_1^T \mathbf{A}_1| / |\mathbf{A}_0^T \mathbf{A}_0|]^{\frac{1}{2}} \left[1 + \frac{(p_1 - p_0)}{(n - p_1)} F \right]^{-(n/2)} \quad (7.3)$$

Here, F represents the F -test statistic for comparing models M_0 and M_1 . Due to the ratio of undefined constants, this Bayes factor is inherently indeterminate. The formula for the F -test statistic is:

$$F = \frac{\left(\frac{\mathbf{R}_0 - \mathbf{R}_1}{p_1 - p_0} \right)}{\left(\frac{\mathbf{R}_1}{n - p_1} \right)},$$

Moreover, the F value used in the Bayes Factor equation can be derived from the following identity:

$$\frac{\mathbf{R}_0}{\mathbf{R}_1} = 1 + \left(\frac{p_1 - p_0}{n - p_1} \right) F$$

Spiegelhalter and Smith proposed a satisfactory solution to the problem of determining the ratio c_0/c_1 by introducing the concept of an imaginary training sample. Let $\mathbf{A}_0(\ell)$ and $\mathbf{A}_1(\ell)$ be the design matrix of M_0 and M_1 occurring in the "thought experiment" generating the imaginary training sample, they obtained the following result

$$\frac{c_0}{c_1} = [|\mathbf{A}_1(\ell)^T \mathbf{A}_1(\ell)| / |\mathbf{A}_0(\ell)^T \mathbf{A}_0(\ell)|]^{-\frac{1}{2}} \quad (7.4)$$

Our goal is to calculate the SS Bayes Factor, in general, using the generalized prior $\pi(\theta_i, \sigma) \propto \sigma^{-(1+q_i)}$ where different values of q will result in well-known priors such as the reference, complete and modified Jeffrey's priors. Additionally, we will calculate this Bayes factor for ANOVA, to compare it with the results obtained by Spiegelhalter and Smith (1980)[14].

Suppose that we want to generalize the result above and compare M_0 vs M_1 where M_0 has p_0 unknown parameters, and M_1 has p_1 unknown parameters. Assume that our prior has the general form

$$\pi(\theta_i, \sigma | A_i) \propto \sigma^{-(1+q_i)}, i = 0, 1$$

where there are many possible choices for q_i . According to Berger and Pericchi (1996) [3], $q_i = 0$ is the reference prior (example above), and $q_i = p_1 - p_0$ is the modified Jeffrey's prior. In general, the marginal distribution for n samples and $i = 0, 1$ is given by

$$m_i(\mathbf{y}) = \int_{-\infty}^{\infty} \int_0^{\infty} \frac{c_i}{\sigma^{1+q_i}} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^n \exp\left\{-\frac{1}{2\sigma^2} \left[\mathbf{R}_i(\mathbf{y}) + (\boldsymbol{\theta}_i - \mathbf{y})^T \mathbf{A}_i^T \mathbf{A}_i (\boldsymbol{\theta}_i - \mathbf{y})\right]\right\} d\boldsymbol{\theta}_i d\sigma$$

Where $\mathbf{R}_i(\mathbf{y}) = (\mathbf{y} - \mathbf{A}_i \hat{\boldsymbol{\theta}}_i)^T (\mathbf{y} - \mathbf{A}_i \hat{\boldsymbol{\theta}}_i)$. Since the second term inside the exponential function is a multivariate normal distribution, we can integrate it with respect to $\boldsymbol{\theta}_i$. Let $\dim(\boldsymbol{\theta}_i) = p_i$, then the marginal can be written as:

$$m_i(\mathbf{y}) = c_i \left(\frac{1}{\sqrt{2\pi}}\right)^{n-p_i} |\mathbf{A}_i^T \mathbf{A}_i|^{-1/2} \int_0^{\infty} \left(\frac{1}{\sigma}\right)^{n+q_i-p_i+1} \exp\left\{-\frac{\mathbf{R}_i(\mathbf{y})}{2\sigma^2}\right\} d\sigma.$$

We can solve this integral by employing the following change of variables

$$\eta = \frac{\mathbf{R}_i(\mathbf{y})}{2\sigma^2}, \quad \sigma = \frac{\mathbf{R}_i^{1/2}(\mathbf{y})}{\sqrt{2}\eta^{1/2}}, \quad d\sigma = -\frac{\mathbf{R}_i^{1/2}(\mathbf{y})}{\sqrt{2}} \frac{1}{2} \frac{1}{\eta^{3/2}} d\eta$$

$$\begin{aligned} m_i(\mathbf{y}) &= \left(\frac{1}{\sqrt{2\pi}}\right)^{n-p_i} c_i |\mathbf{A}_i^T \mathbf{A}_i|^{-1/2} \int_0^{\infty} \exp^{-\eta} \eta^{\frac{n+q_i+1-p_i}{2}} \left(\frac{2}{\mathbf{R}_i(\mathbf{y})}\right)^{(n+q_i+1-p_i)/2} \left(\frac{\mathbf{R}_i(\mathbf{y})}{2}\right)^{1/2} \frac{1}{2} \frac{1}{\eta^{3/2}} d\eta \\ &= 2^{\frac{n+q_i-p_i-2}{2}} \left(\frac{1}{\sqrt{2\pi}}\right)^{n-p_i} \left(\frac{1}{\mathbf{R}_i(\mathbf{y})}\right)^{\frac{n+q_i-p_i+1}{2}} \left(\mathbf{R}_i(\mathbf{y})\right)^{1/2} |\mathbf{A}_i^T \mathbf{A}_i|^{-1/2} \frac{c_i}{2} \int_0^{\infty} \exp^{-\eta} \eta^{\frac{n+q_i-p_i+3}{2}-1} d\eta \end{aligned}$$

Solving the integral above yields

$$\int_0^{\infty} \exp^{-\eta} \eta^{\frac{n+q_i-p_i+3}{2}-1} d\eta = \Gamma\left(\frac{n+q_i-p_i+3}{2}\right)$$

Therefore, the marginal for M_i is

$$\begin{aligned} m_i(\mathbf{y}) &= c_i 2^{\frac{n+q_i-p_i}{2}} \left(\frac{1}{\sqrt{2\pi}}\right)^{n-p_i} \left(\frac{1}{\mathbf{R}_i(\mathbf{y})}\right)^{\frac{n+q_i-p_i+1}{2}} \left(\frac{\mathbf{R}_i(\mathbf{y})}{2}\right)^{1/2} |\mathbf{A}_i^T \mathbf{A}_i|^{-1/2} \Gamma\left(\frac{n-p_i+q_i+3}{2}\right) \\ &\propto c_i \sqrt{|\mathbf{A}_i^T \mathbf{A}_i|^{-1} [4\pi]^{p_i} \mathbf{R}_i(\mathbf{y})^{-(n+q_i-p_i)}} \end{aligned}$$

After we obtain the marginals, we can calculate the Bayes factor as

$$B_{01}^N(\mathbf{y}) = \frac{m_0(\mathbf{y})}{m_1(\mathbf{y})} = \frac{c_0}{c_1} \sqrt{[4\pi]^{(p_0-p_1)} \frac{|\mathbf{A}_1^T \mathbf{A}_1| [\mathbf{R}_1(\mathbf{y})]^{n+q_1-p_1}}{|\mathbf{A}_0^T \mathbf{A}_0| [\mathbf{R}_0(\mathbf{y})]^{n+q_0-p_0}}}$$

The following step involves computing the Bayes factor for the minimal training sample. For this situation, where at least one observation per unknown parameter is necessary, the minimal training sample size can be found as:

$$n_{01} = \max\{p_0 + 1, p_1 + 1\}.$$

In this case, the Bayes factor for the minimal training sample takes the form:

$$B_{10}^N(\mathbf{y}(\ell)) = \frac{m_1(\mathbf{y}(\ell))}{m_0(\mathbf{y}(\ell))} = \frac{c_1}{c_0} \sqrt{\frac{[4\pi]^{(p_1-p_0)} |\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| [\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_0-p_0}}{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| [\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_1-p_1}}}$$

The previously discussed **Lemma 1** establishes that:

$$B_{01}(\mathbf{y}(-\ell)|\mathbf{y}(\ell)) = B_{01}^N(\mathbf{y})B_{10}^N(\mathbf{y}(\ell))$$

Consequently,

$$B_{01}(\mathbf{y}(-\ell)|\mathbf{y}(\ell)) = \sqrt{\frac{|\mathbf{A}_1^T\mathbf{A}_1| |\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| [\mathbf{R}_1(\mathbf{y})]^{n+q_1-p_1} [\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_0-p_0}}{|\mathbf{A}_0^T\mathbf{A}_0| |\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| [\mathbf{R}_0(\mathbf{y})]^{n+q_0+p_0} [\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_1-p_1}}}$$

By leveraging these results, we derive the Empirical Intrinsic Bayes Factor (IBF) lower bound for General Linear Models:

$$\underline{B}_{01}^{GLI*} = \sqrt{\frac{|\mathbf{A}_1^T\mathbf{A}_1| [\mathbf{R}_1(\mathbf{y})]^{n+q_1-p_1}}{|\mathbf{A}_0^T\mathbf{A}_0| [\mathbf{R}_0(\mathbf{y})]^{n+q_0-p_0}}} \min_{\ell=1, \dots, L} \sqrt{\frac{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| [\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_0-p_0}}{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| [\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_1-p_1}}}$$

Additionally, the Theoretical Intrinsic Bayes Factor (IBF) Lower bound for General Linear Models can be expressed as:

$$\underline{B}_{01}^{GLI} = \sqrt{\frac{|\mathbf{A}_1^T\mathbf{A}_1| [\mathbf{R}_1(\mathbf{y})]^{n+q_1-p_1}}{|\mathbf{A}_0^T\mathbf{A}_0| [\mathbf{R}_0(\mathbf{y})]^{n+q_0-p_0}}} \inf_{\mathbf{y}(\ell) \in D} \sqrt{\frac{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| [\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_0-p_0}}{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| [\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_1-p_1}}} \quad (7.5)$$

7.3 The One-way Layout (ANOVA)

The One-way Layout in analysis of variance (ANOVA) involves m groups of observations, each with n_i observations ($j = 1, \dots, n_i$) independently drawn from a normal distribution $y_{ij} \sim N(\mu_i, \sigma^2)$, given $\mu_1, \dots, \mu_m, \sigma^2$.

The models considered are:

$$M_0 : \mu_1 = \dots = \mu_m \text{ versus } M_1 : \mu_i \neq \mu_j \text{ for some } i \neq j.$$

In this scenario, $p_1 = m$, $p_0 = 1$, and the general form of the matrices \mathbf{A}_0 and \mathbf{A}_1 can be derived. Let $n = \sum_{i=1}^m n_i$. Notably, the design matrices A_0 and A_1 for the one-way ANOVA can generally be written as follows:

$$A_0 = \begin{bmatrix} 1 \\ \cdot \\ \cdot \\ \cdot \\ 1 \end{bmatrix}, A_0^T = \begin{bmatrix} 1 & \dots & 1 \end{bmatrix}, A_1 = \begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 1 \\ \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}$$

The matrices A_0 and A_1 are $(n \times m)$ matrices, and each block of ones in A_1 corresponds to each of the m groups. In this case, every block i has n_i rows and $\det(\mathbf{A}_0^T \mathbf{A}_0) = n$. The matrix product of \mathbf{A}_1^T and \mathbf{A}_1 can be expressed as:

$$\mathbf{A}_1^T \mathbf{A}_1 = \begin{bmatrix} n_1 & 0 & \dots & 0 \\ \cdot & n_2 & \dots & \cdot \\ \cdot & \cdot & \ddots & \cdot \\ 0 & 0 & \dots & n_m \end{bmatrix}$$

The matrix resulting from the product of these two matrices is itself diagonal, with its determinant expressed as:

$$\det(\mathbf{A}_1^T \mathbf{A}_1) = \prod_{i=1}^m n_i \quad (7.6)$$

Thus, we can compute the term $|\mathbf{A}_1^T \mathbf{A}_1|/|\mathbf{A}_0^T \mathbf{A}_0|$ in equation (7.3) as:

$$|\mathbf{A}_1^T \mathbf{A}_1|/|\mathbf{A}_0^T \mathbf{A}_0| = \frac{\prod_{i=1}^m n_i}{n} \quad (7.7)$$

In this case, the minimal training sample requires at least one observation in each group plus one extra observation in any of the groups to estimate σ^2 . Thus, we require that $n_i = 1, i = 1, \dots, j-1, j+1, \dots, m$, and $n_j = 2$, for some $j \in [1, m]$.

For the case mentioned above, we can use equation (7.4), and the result obtained in (7.6) to get c_1 as:

$$\det(\mathbf{A}_1(\ell)^T \mathbf{A}_1(\ell)) = (1 \cdots 1 \cdot 2 \cdot 1 \cdots 1) = 2 \quad (7.8)$$

Similarly, c_0 can be derived as:

$$\det(\mathbf{A}_0(\ell)^T \mathbf{A}_0(\ell)) = n = \sum_{j=1}^m n_j = (1 + \cdots + 1 + 2 + 1 + \cdots + 1) = m + 1 \quad (7.9)$$

Thus, utilizing the outcomes from (7.8) and (7.9), we can evaluate the expression in equation (7.4) as follows:

$$\frac{c_0}{c_1} = \left(\frac{\det(\mathbf{A}_0(\ell)^T \mathbf{A}_0(\ell))}{\det(\mathbf{A}_1(\ell)^T \mathbf{A}_1(\ell))} \right)^{\frac{1}{2}} = \left(\frac{m+1}{2} \right)^{\frac{1}{2}} \quad (7.10)$$

Now, by substituting (7.10) and (7.7) into (7.3), derive the following expression for the Bayes factor:

$$B_{01}^{SS} = \left(\frac{m+1}{2} \right)^{\frac{1}{2}} \left[\prod_{i=1}^m n_i / n \right]^{\frac{1}{2}} \left[1 + \frac{(m-1)}{n-m} F \right]^{-n/2} \quad (7.11)$$

The SS Bayes Factor can alternatively be expressed compactly as:

$$B_{01}^{SS} = \left(\frac{\left(\frac{m+1}{2} \right) \left[\prod_{i=1}^m n_i / n \right]}{\left[1 + \frac{(m-1)}{n-m} F \right]^n} \right)^{\frac{1}{2}} \quad (7.12)$$

Furthermore, relating the F statistics to the p-value can be achieved through:

$$F_{v_1, v_2} = qf(1 - p, v_1, v_2)$$

Here, qf denotes the quantile function for the F distribution, while v_1 and v_2 represent the degrees of freedom. For this case, $v_1 = p_1 - p_0$ and $v_2 = n - p_1$. Consequently, expressing the SS Bayes factor in terms of the p-value yields:

$$B_{01}^{SS}(p) = \left(\frac{\left(\frac{m+1}{2}\right) [\prod_{i=1}^m n_i/n]}{\left[1 + \frac{(m-1)}{n-m} qf(1 - p, p_1 - p_0, n - p_1)\right]^n} \right)^{\frac{1}{2}}$$

This formula offers a straightforward method to convert p-values into Bayes factors. It's akin to the Sellke et al. (2001) [13] bound, but distinguishes itself by being dynamic, enhancing accuracy with increasing information. This technique is applicable to any Bayes factor that relies on F-statistics.

Our subsequent aim involves computing B_{01}^{GSS} as in equation (7.5) for ANOVA models and comparing it with their findings. Considering that $A_0(\ell)$ and $A_1(\ell)$ remain independent of observations in ANOVA models, and the square root is a monotonic function, we can rewrite the expression as:

$$\sup_{y(\ell) \in D_n} \sqrt{\frac{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| [\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_i-p_0}}{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| [\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q-p_1}}} = \sqrt{\frac{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)|}{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)|}} \sup_{y(\ell) \in D_n} \frac{[\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_i-p_0}}{[\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_i-p_1}}$$

The General Theoretical and Empirical SS Bayes factors for ANOVA can be expressed respectively as:

$$B_{10}^{GSSA} = \sqrt{\frac{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| |\mathbf{A}_0^T\mathbf{A}_0| [\mathbf{R}_0(\mathbf{y})]^{n+q_0-p_0}}{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| |\mathbf{A}_1^T\mathbf{A}_1| [\mathbf{R}_1(\mathbf{y})]^{n+q_1-p_1}}} \sup_{y(\ell) \in D} \frac{[\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_1-p_1}}{[\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_0-p_0}}$$

$$B_{10}^{GESSA} = \sqrt{\frac{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| |\mathbf{A}_0^T\mathbf{A}_0| [\mathbf{R}_0(\mathbf{y})]^{n+q_0-p_0}}{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| |\mathbf{A}_1^T\mathbf{A}_1| [\mathbf{R}_1(\mathbf{y})]^{n+q_1-p_1}}} \max_{\ell=1, \dots, L} \frac{[\mathbf{R}_1(\mathbf{y}(\ell))]^{n_{01}+q_1-p_1}}{[\mathbf{R}_0(\mathbf{y}(\ell))]^{n_{01}+q_0-p_0}}$$

7.3.1 The Full Jeffrey's (no independence) $q_k = p_k$, for $k = i, j$ for ANOVA

In accordance with our earlier calculations pertaining to the General Empirical SS Bayes factor for ANOVA, we apply the methodology proposed by Spiegelhalter and Smith (1982) [15]. This involves the selection of the full Jeffrey's prior where $q_i = p_i$. Through the utilization of an imaginary training sample, which renders maximum support for M_0 , we ascertain that $F_y(\ell) = 0$. Consequently, we obtain the General Theoretical SS Bayes Factor for ANOVA under the Full Jeffrey's prior

$$B_{01}^{GESSA_{FJP}} = \sqrt{\frac{m+1}{2} \left[\prod_{i=1}^m n_i/n \right] \left(\frac{\mathbf{R}_1(\mathbf{y})}{\mathbf{R}_0(\mathbf{y})} \right)^n \min_{\ell=1, \dots, L} \left(\frac{\mathbf{R}_0(\mathbf{y}(\ell))}{\mathbf{R}_1(\mathbf{y}(\ell))} \right)^{m+1}}.$$

Now, expressing the ratio of the residual sum of squares as an F-statistics results in:

$$B_{01}^{GESSA_{FJP}} = \sqrt{\frac{m+1}{2} \left[\prod_{i=1}^m n_i/n \right] \left(1 + \frac{(m-1)}{(n-m)} F_y \right)^{-n} \max_{\ell=1, \dots, L} \left(1 + \frac{(p_1 - p_0)}{(n_{01} - p_1)} F_{y(\ell)} \right)^{m+1}}$$

The General Theoretical SS Bayes factor for ANOVA results in:

$$B_{10}^{GSSA_{FJP}} = \sqrt{\frac{2}{m+1} \left[n / \prod_{i=1}^m n_i \right] \left(1 + \frac{(m-1)}{(n-m)} F_y \right)^n \sup_{y(\ell) \in D} \left[\frac{\mathbf{R}_1(\mathbf{y}(\ell))}{\mathbf{R}_0(\mathbf{y}(\ell))} \right]^{m+1}}$$

Observing that $\mathbf{R}_0(\mathbf{y}(\ell)) \geq \mathbf{R}_1(\mathbf{y}(\ell))$, it follows that:

$$\sup_{y(\ell) \in D} \frac{[\mathbf{R}_1(\mathbf{y}(\ell))]}{[\mathbf{R}_0(\mathbf{y}(\ell))]} = 1 \implies \sup_{y(\ell) \in D} \left[\frac{\mathbf{R}_1(\mathbf{y}(\ell))}{\mathbf{R}_0(\mathbf{y}(\ell))} \right]^{m+1} = 1, \quad \forall m > 0$$

Consequently, the General Theoretical SS Bayes Factor, employing the full Jeffrey's prior for ANOVA models, simplifies to:

$$B_{10}^{GSSA_{FJP}} = \sqrt{\frac{2}{m+1} \left[n / \prod_{i=1}^m n_i \right] \left(1 + \frac{(m-1)}{(n-m)} F_y \right)^n} = B_{10}^{SS}.$$

Alternatively, this Bayes factor can be expressed as:

$$B_{10}^{GSSA_{FJP}} = \sqrt{\frac{2}{m+1} \left[n / \prod_{i=1}^m n_i \right] \left[\frac{\mathbf{R}_0(\mathbf{y})}{\mathbf{R}_1(\mathbf{y})} \right]^n}.$$

7.3.2 The Reference prior ($q_k = 0, k = i, j$) for ANOVA

In this section, we delve into the theoretical SS Bayes factor for the ANOVA scenario, employing the reference prior obtained by setting $q_k = 0$ for $k = i, j$. The theoretical SS Bayes Factor for ANOVA, specifically under the *reference prior* context, is formulated as:

$$B_{10}^{GSSA_{RP}} = \sqrt{\frac{|\mathbf{A}_1^T(\ell)\mathbf{A}_1(\ell)| |\mathbf{A}_0^T\mathbf{A}_0| [\mathbf{R}_0(\mathbf{y})]^{n+1}}{|\mathbf{A}_0^T(\ell)\mathbf{A}_0(\ell)| |\mathbf{A}_1^T\mathbf{A}_1| [\mathbf{R}_1(\mathbf{y})]^{n-(m+1)}} \sup_{\mathbf{y}(\ell) \in D} \left(\frac{1}{\mathbf{R}_0(\mathbf{y}(\ell))} \right)^m}$$

Upon inspection of the aforementioned equation, it becomes clear that under this particular prior, the theoretical SS Bayes factor tends to either zero or infinity. Consequently, the computation of the theoretical SS BF with this prior is rendered unfeasible. This implies that for reference priors, the SS bounds do not exist. This limitation is not previously highlighted in the works of Smith and Spiegelhalter (1980)[14] and (1982) [15], thereby underscoring an important consideration for future research in this domain.

7.3.3 The Modified Jeffrey's Prior ($q_j = p_j - p_i, q_i = 0$) for ANOVA

The General Empirical SS Bayes Factor for ANOVA, under the Modified Jeffrey's prior results in:

$$B_{10}^{GESSA_{MJP}} = \sqrt{\frac{2}{m+1} \left[n / \prod_{i=1}^m n_i \right] \left(\frac{\mathbf{R}_0(\mathbf{y})}{\mathbf{R}_1(\mathbf{y})} \right)^{n-1} \max_{\ell=1, \dots, L} \left(\frac{\mathbf{R}_1(\mathbf{y}(\ell))}{\mathbf{R}_0(\mathbf{y}(\ell))} \right)^m}.$$

Similarly, the general Theoretical SS Bayes Factor for ANOVA, under the Modified Jeffrey's prior results in:

$$B_{10}^{GSSA_{MJP}} = \sqrt{\frac{2}{m+1} \left[n / \prod_{i=1}^m n_i \right] \left(\frac{\mathbf{R}_0(\mathbf{y})}{\mathbf{R}_1(\mathbf{y})} \right)^{n-1} \sup_{\mathbf{y}(\ell) \in D} \left(\frac{\mathbf{R}_1(\mathbf{y}(\ell))}{\mathbf{R}_0(\mathbf{y}(\ell))} \right)^m}.$$

Alternatively, this Bayes Factor can be expressed in terms of the F statistics as:

$$B_{10}^{GSSA_{MJP}} = \sqrt{\frac{2}{m+1} \left[n / \prod_{i=1}^m n_i \right] \left(1 + \frac{(m-1)}{(n-m)} F_y \right)^{n-1}}.$$

7.4 Conclusions

In conclusion, the ANOVA case analysis has provided valuable insights regarding the use of different priors in the context of Full Jeffreys, Modified Jeffreys, and Reference Prior. The SS bounds obtained for Full Jeffreys and Modified Jeffreys are informative, useful, and exhibit close proximity to each other. Importantly, it has been demonstrated that the Bayes Factor can be expressed as a function of the F-Statistics under both priors. However, in the case of the Reference Prior, the SS bound fails to provide informative results, highlighting its lack of usefulness. Notably, it is intriguing to observe the sensitivity of the SS bound to the initial objective prior, which further supports the conclusions drawn by Berger and Pericchi (1996) [3] and others that the Modified Jeffreys prior is a superior choice, particularly within the Linear Gaussian Model and potentially in a broader context. Additionally, it may be of interest to explore the extent to which the Modified Jeffreys prior possesses additional properties, such as matching, besides its inherent simplicity and the one-on-one relationship it establishes with the F-Statistics.

CHAPTER 8

SEPARATED HYPOTHESIS TESTING

Traditional hypothesis testing typically involves comparing a single hypothesis to a pre-determined null hypothesis, assuming a single underlying distribution. However, in certain scenarios, there arises a need to examine and compare two distinct distributions simultaneously. This is the essence of separated hypothesis testing, a concept that expands upon conventional approaches by accommodating situations where the hypothesis being tested corresponds to two separate populations or data sources. By addressing this unique requirement, separated hypothesis testing enables researchers to gain valuable insights into the potential disparities or similarities between two distinct distributions, thereby enhancing the accuracy and comprehensiveness of statistical inference. In this chapter, we delve into the theoretical underpinnings, methodology, and applications of separated hypothesis.

We explored the concept of separated hypothesis testing through practical examples involving the comparison of different probability distributions. Specifically, we consider the scenarios of Poisson vs. Geometric and Poisson vs. Negative Binomial distributions. We computed IBF bounds for both comparisons, ensuring the reliability and consistency of our results. To validate our findings, we conducted simulations where we progressively augmented the sample size, observing a numerical convergence toward accurate decisions in our experiments. Each example is accompanied by a comprehensive discussion of the methodologies and code used in our analyses. The figures presented serve as compelling visual evidence substantiating the validity of our conclusions.

8.1 The Poisson vs. Geometric Separated Hypothesis Test

Let's consider comparing two distinct models: the Poisson and the Geometric distributions, where $\mathbf{y} = (y_1, \dots, y_n)$ represents independently and identically distributed (i.i.d.) observations. This frames our model selection problem as follows:

$$M_0 : Y \sim \text{Poisson}(\lambda) \text{ vs } M_1 : Y \sim \text{Geometric}(\theta)$$

The densities for each model are:

$$f_0(y|\lambda) = \frac{e^{-\lambda} \lambda^y}{y!}, \lambda > 0, \text{ and } f_1(y|\theta) = \theta(1-\theta)^y, \theta \in (0,1), y = 0, 1, 2, \dots$$

The Jeffreys's priors and likelihoods associated with M_0 and M_1 yield the following outcomes:

$$\begin{aligned} \pi_0^J(\lambda) &= \frac{c_0}{\lambda^{1/2}}, \quad \pi_1^J(\theta) = \frac{c_1}{\theta(1-\theta)^{1/2}} \\ f_0(\mathbf{y}|\lambda) &= \frac{e^{-n\lambda} \lambda^{\sum_{i=1}^n y_i}}{\prod_{i=1}^n y_i!}, \text{ and } f_1(\mathbf{y}|\theta) = \theta^n (1-\theta)^{\sum_{i=1}^n y_i} \end{aligned}$$

Given that both models entail a single parameter, the minimal training sample size required is one. Utilizing the priors and likelihoods, we proceed to compute the marginals and subsequently derive the Bayes factors as follows:

$$\begin{aligned} m_0(\mathbf{y}) &= \frac{c_0}{\prod_{i=1}^n y_i!} \int_0^\infty e^{-n\lambda} \lambda^{\sum_{i=1}^n y_i} \frac{1}{\lambda^{1/2}} d\lambda \\ m_1(\mathbf{y}) &= c_1 \int_0^1 \theta^n (1-\theta)^{\sum_{i=1}^n y_i} \frac{1}{\theta(1-\theta)^{1/2}} d\theta \end{aligned}$$

After straightforward algebraic manipulation, we find that the integral in $m_0(\mathbf{y})$ follows a Gamma distribution, while the integral in $m_1(\mathbf{x})$ corresponds to a Beta distribution, wherein:

$$m_0(\mathbf{y}) = \frac{c_0 \Gamma(\sum y_i + \frac{1}{2})}{\prod_{i=1}^n y_i!} (1/n)^{(\sum y_i + \frac{1}{2})}$$

$$m_1(\mathbf{y}) = c_1 \int_0^1 \theta^{n-1} (1-\theta)^{\sum_{i=1}^n y_i - 1/2} d\theta = c_1 \frac{\Gamma(n)\Gamma(\sum y_i + 1/2)}{\Gamma(n + \sum y_i + 1/2)}$$

Thus, the Bayes Factor results in:

$$B_{01}^N(\mathbf{y}) = \frac{c_0 \Gamma(n + \sum y_i + 1/2)}{c_1 \prod_{i=1}^n y_i! \Gamma(n) n^{(\sum y_i + \frac{1}{2})}}$$

The Empirical IBF upper bound can be obtained by computing:

$$\bar{B}_{10}^{I*}(\mathbf{y}) = B_{10}^N(\mathbf{y}) \sup_{y(\ell) \in D_{n,k}} \frac{m_0(y(\ell))}{m_1(y(\ell))}$$

Evidently, the following marginals are required:

$$m_0(y(\ell)) = c_0 \frac{\Gamma(y(\ell) + 1/2)}{y(\ell)!}, \quad m_1(y(\ell)) = c_1 \frac{\Gamma(1)\Gamma(y(\ell) + 1/2)}{\Gamma(1 + y(\ell) + 1/2)}$$

By using those marginals, we can calculate the Bayes factor for $y(\ell)$ as:

$$B_{01}^N(y(\ell)) = \frac{c_0}{c_1} \frac{1}{\Gamma(y(\ell) + 1 + 1/2)}$$

The above Bayes factor, as a function of $y(\ell)$ is decreasing. Therefore, calculating the theoretical supremum, which occurs at $y(\ell) = 0$, results in:

$$\sup_{y(\ell) \in D} B_{01}^N(y(\ell)) = \frac{c_0}{c_1} \frac{1}{\Gamma(1 + 1/2)}$$

Utilizing the identity $\Gamma(n + 1/2) = (2n! \sqrt{\pi}) / (4^n n!)$, we derived:

$$\sup_{y(\ell) \in D} B_{01}^N(y(\ell)) = \frac{c_0}{c_1} \frac{2}{\sqrt{\pi}}$$

This results in the Theoretical IBF upper bound:

$$\bar{B}_{10}^I(\mathbf{y}) = \frac{\prod_{i=1}^n y_i! (n-1)! n^{(n\bar{y} + \frac{1}{2})}}{\Gamma(n(1 + \bar{y}) + 1/2)} \frac{2}{\sqrt{\pi}}$$

And the Theoretical IBF lower bound:

$$\underline{B}_{01}^I(\mathbf{y}) = \frac{\Gamma(n(1 + \bar{y}) + 1/2) \sqrt{\pi}}{\prod_{i=1}^n y_i! (n-1)! n^{(n\bar{y} + \frac{1}{2})}} \frac{1}{2}$$

Similarly, the empirical IBF upper bound is:

$$\overline{B}_{10}^{I*}(\mathbf{y}) = \frac{\prod_{i=1}^n y_i! (n-1)! n^{(n\bar{y} + \frac{1}{2})}}{\Gamma(n(1 + \bar{y}) + 1/2)} \frac{1}{\Gamma(y_{(1)} + 1 + \frac{1}{2})}, \quad y_{(1)} = \min_{i=1, \dots, n} \{y_i\}$$

And the empirical IBF lower bound is:

$$\underline{B}_{01}^{I*}(\mathbf{y}) = \frac{\Gamma(n(1 + \bar{y}) + 1/2) \Gamma(y_{(n)} + 1 + \frac{1}{2})}{\prod_{i=1}^n y_i! (n-1)! n^{(n\bar{y} + \frac{1}{2})}}, \quad y_{(n)} = \max_{i=1, \dots, n} \{y_i\}$$

We aim to ascertain the reliability and consistency of our obtained results by implementing an **R** script specifically designed to calculate the Bayes factor in two distinct scenarios: when the null hypothesis (H_0) holds true and when the alternative hypothesis (H_1) is true. By conducting this analysis, we will demonstrate that when the samples are drawn from the Poisson distribution, the computed Bayes factor approaches zero, indicating stronger support for H_0 . Conversely, when the samples originate from the Geometric distribution, the Bayes factor exhibits values significantly greater than one, favoring H_1 . Furthermore, we will calculate the empirical IBF bounds, providing numerical bounds for both $B_{01}^N(\mathbf{y})$ and $B_{10}^N(\mathbf{y})$. The **R** script below generates one hundred simulations of thirty independent samples from the Poisson distribution and from the Geometric distribution. By employing this approach, we calculated the strength of evidence in favor of the null hypothesis in both scenarios. As anticipated, the obtained values are consistent with our prior knowledge about the data that we simulated, indicating the expected outcome of growing indefinitely when the null is false and going to zero under the null. We smoothed the results by calculating the average value for each sample size over all of the simulations.

R Program

```

1
2 #100 simulations of 30 samples each
3
4 #Null hypothesis true, lambda =1
5 lambda = 1
6 n_sim = 100; n_data = 30; p=0.5
7 bf10ss_pois = c(); data_pois = c(); bfss10_pois = c();
8 bf10ss_geom = c(); data_geom = c(); bfss10_geom = c();
9 for (i in 1:n_sim) {
10   for (j in 1:n_data) {
11
12     data_geom = c(data_geom, rgeom(n=1, prob=p));
13     data_pois = c(data_pois, rpois(n=j, lambda=lambda));
14     s_geom = sum(data_geom)
15     s_pois = sum(data_pois)
16     y_fact_geom = factorial(data_geom)
17     y_fact_pois = factorial(data_pois)
18
19     #Marginals
20     m1_geom = (j^(s_geom + 0.5)*factorial(j-1)*prod(y_fact_geom)*2)
21     m0_geom = (gamma(j + s_geom + 0.5)*sqrt(pi))
22     bfss10_geom = c(bfss10_geom, m1_geom/m0_geom)
23     m1_pois = (j^(s_pois + 0.5)*factorial(j-1)*prod(y_fact_pois)*2)
24     m0_pois = (gamma(j + s_pois + 0.5)*sqrt(pi))
25     bfss10_pois = c(bfss10_pois, m1_pois/m0_pois)
26   }
27   bf10ss_pois = cbind(bf10ss_pois, bfss10_pois)
28   bf10ss_geom = cbind(bf10ss_geom, bfss10_geom)
29 }
30 bf10ss_avg_geom = rep(0, n_data)
31 bf10ss_avg_pois = rep(0, n_data)
32 for (i in 1:n_data) {
33   bf10ss_avg_geom[i] = mean(bf10ss_geom[i,])
34   bf10ss_avg_pois[i] = mean(bf10ss_pois[i,])
35 }
36 par(mfrow=c(1,2))
37 plot(1:n_data, bf10ss_avg_pois, type="l", col="red", main="Geometric vs Poisson Distribution Bayes Factor Bound",
38      ylab="Average SS B10 Bound", xlab="Sample size")
39 mtext(text="Data generated from the Poisson distribution", side=3)
40 plot(1:n_data, bf10ss_avg_geom, type="l", col="red", main="Geometric vs Poisson Distribution Bayes Factor Bound",
41      ylab="Average SS B10 Bound", xlab="Sample size")
42 mtext(text="Data generated from the Geometric distribution", side=3)

```




Figure 8–1: Poisson vs. Geometric IBF 10 bound as the sample size increases when H_0 is true (left) and when H_0 is false (right).

Once we obtain the vector of samples \mathbf{y} , we can proceed to compute the empirical bound for $B_{10}^{AI}(\mathbf{y})$. The code below calculates the value of the Bayes Factor for each sample in \mathbf{y} and then identifies the maximum value among them, representing the empirical maximum. By employing this code, we obtain a numerical measure that reflects the strength of evidence in favor of the alternative hypothesis.

R Program

```

1 bfss01_mts = function(y, print=TRUE) {
2   n = length(y); bfss_01=c()
3   for (i in 1:n) { m1 = gamma(y[i] + 1.5); m0 = 1; bfss_01 = c(bfss_01, m0/m1) }
4   y_max = y[which.max(bfss_01)]; bfss_max = max(bfss_01)
5   if (print) {
6     print("Samples"); print(y); print("BFSS01 for each MTS"); print(bfss_01)
7     print("The value that maximize the BFSS 01 is"); print(y_max)
8     print("The maximal value of the BFSS 01 is"); print(bfss_max)
9   }
10  else { bfss = c(); bfss$y_max= y_max; bfss$bfss_max = bfss_max; return(bfss) }
11 }

```

Upon implementing the aforementioned algorithm, leveraging a vector of samples drawn from the Poisson distribution, we acquired the empirical supremum of the Bayes Factor. Notably, this empirical outcome corresponded precisely with the theoretical supremum. This concordance is attributable to the inclusion of the value '0' within the sample set, as is evident in the code output below

R Program

```

1 > bfss01_mts(y)
2 [1] "Samples"
3 [1] 3 1 2 0 5 0 2 1 0 2
4 [1] "BFSS01 for each MTS"
5 [1] 0.085971746 0.752252778 0.300901111 1.128379167 0.003473606
6 [1] 1.128379167 0.300901111 0.752252778 1.128379167 0.300901111
7 [1] "The value that maximizes the BFSS 01 is"
8 [1] 0
9 [1] "The maximal value of BFSS 01 is"
10 [1] 1.128379
11 [1] "The theoretical sup of BFSS 01 is"
12 [1] 2/sqrt(pi)
13 [1] 1.128379

```

A valuable approach involves comparing the theoretical supremum with the empirical supremum as the sample size (n) expands. This can be accomplished utilizing specific computational methodologies, particularly when the underlying data are derived from Poisson or Geometric distributions. The implementation of this comparison can be conducted via the following code

R Program

```

1 y_pois = rpois(n=20,lambda=0.5); y_geom = rgeom(n=20,prob=0.8)
2 bfss10_comparison = function(y,c_theoretical = 2/sqrt(pi)) {
3   bfss10_empirical = c(); bfss10_theoretical = c(); s = 0
4   for (i in 1:n) {
5     s = s + y[i]; y_fact = factorial(y[1:i]);
6     m1 = (i^(s + 0.5)*factorial(i-1)*prod(y_fact)); m0 = (gamma(i + s + 0.5))
7     bfss01= bfss01_mts(y[1:i],print=FALSE); c_empirical = bfss01$bfss_max
8     bfss10_empirical = c(bfss10_empirical,c_empirical*m1/m0); bfss10_theoretical= c(bfss10_theoretical,c_
9       theoretical*m1/m0)
10  }
11 plot(1:n,bfss10_empirical,type="l",xlab="Number of samples",ylab="BFSS 10 Empirical")
12 plot(1:n,bfss10_theoretical,type="l",xlab="Number of samples",ylab="BFSS 10 Theoretical")
13 }

```


The following charts were generated using the code above

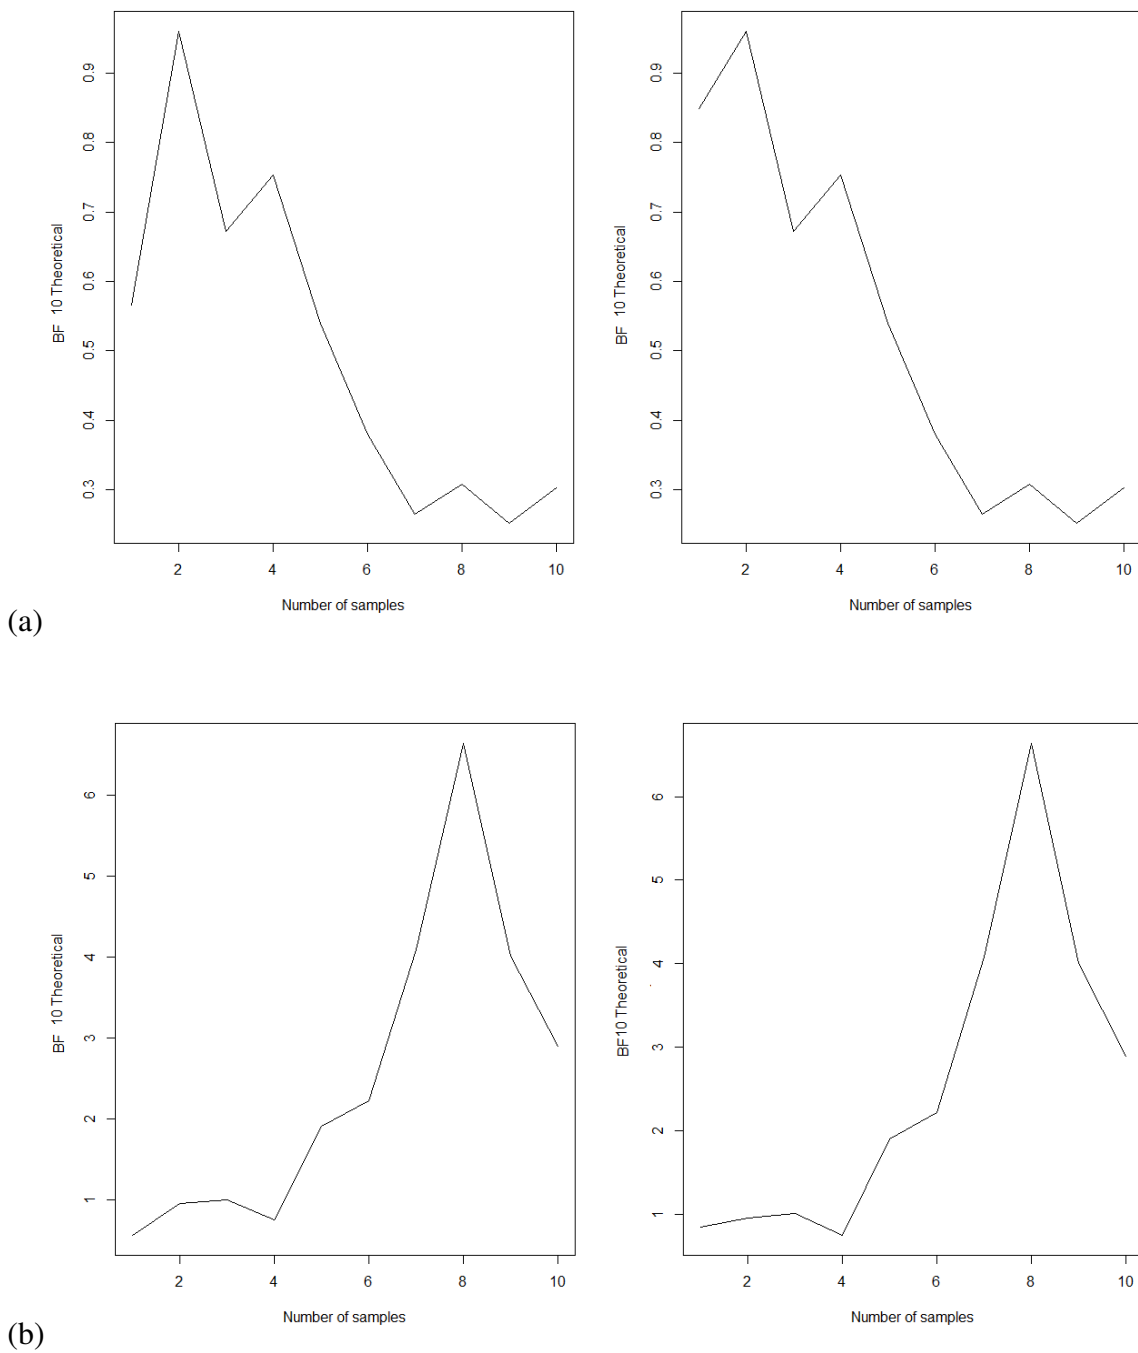


Figure 8-2: (a) Empirical and Theoretical \bar{B}_{10}^J when H_0 is true (b) Empirical and Theoretical \bar{B}_{01}^J when H_0 is false

Our methodological approach involves the generation of one hundred simulations, each encompassing thirty samples. These simulations will be conducted under both the null hypothesis and the alternative hypothesis. For each sample size, we aim to compute the average Bayes factors derived from all the simulations. Further, our study entails the graphical representation of these averaged Bayes factors against their respective sample sizes. Both the averages under the null and the alternative hypotheses will be depicted within the same coordinate system, facilitating a clear comparison and analysis. To automate the generation of such illustrative graphics, we utilize **R** programming language. The corresponding script to generate these charts is

R Program

```

1 Geometric_Poisson_SSBF = function(n,prob=0.5,lambda=1,simulations=100) {
2   geometric_data = c(); poisson_data = c(); bfss01_geometric = c(); bfss01_poisson = c(); bf01_geo_total = c();
   bf01_pois_total = c(); ibf01_geometric = c(); ibf01_poisson = c();ibf01_geo_total = c(); ibf01_pois_
   total = c()
3   for (k in 1:simulations) {
4     for (i in 1:n) {
5       geometric_data = c(geometric_data,rgeom(n=1,prob=prob)); s = sum(geometric_data); y_fact = factorial(
       geometric_data)
6       m1 = (i^(s + 0.5)*factorial(i-1)*prod(y_fact)*2); m0 = (gamma(i + s + 0.5)*sqrt(pi))
7       m1_ibf = (i^(s + 0.5)*factorial(i-1)*prod(y_fact)); m0_ibf = (gamma(i + s + 0.5))
8       ibf10_correction = mean(gamma(geometric_data + 1.5))
9       bfss01_geometric = c(bfss01_geometric,m0/m1); ibf01_geometric = c(ibf01_geometric,(m0_ibf/m1_ibf)*ibf10_
       correction)
10      poisson_data = c(poisson_data,rpois(n=1,lambda=lambda))
11      s = sum(poisson_data); y_fact = factorial(poisson_data)
12      m1 = (i^(s + 0.5)*factorial(i-1)*prod(y_fact)*2); s = sum(log(y_fact)); m0 = (gamma(i + s + 0.5)*sqrt(pi)
        ); m1_ibf = (i^(s + 0.5)*factorial(i-1)*prod(y_fact)); m0_ibf = (gamma(i + s + 0.5))
13      ibf10_correction = mean(gamma(poisson_data + 1.5))
14      bfss01_poisson = c(bfss01_poisson,m0/m1); ibf01_poisson = c(ibf01_poisson,(m0_ibf/m1_ibf)*ibf10_
        correction)
15    }
16    bf01_geo_total = cbind(bf01_geo_total,bfss01_geometric); bf01_pois_total = cbind(bf01_pois_total,bfss01_
        poisson); ibf01_pois_total = cbind(ibf01_pois_total,ibf01_poisson); ibf01_geo_total = cbind(ibf01_geo_
        total,ibf01_geometric)
17  }
18  bf01_geo_avg = c(); bf01_pois_avg = c(); ibf01_geo_avg = c(); ibf01_pois_avg = c(); n = nrow(bf01_geo_total)
19  for (i in 1:n) {
20    bf01_geo_avg = c(bf01_geo_avg,mean(bf01_geo_total[i,])); bf01_pois_avg = c(bf01_pois_avg,mean(bf01_pois_
        total[i,])); ibf01_pois_avg = c(ibf01_pois_avg,mean(ibf01_pois_total[i,])); ibf01_geo_avg = c(ibf01_
        geo_avg,mean(ibf01_geo_total[i,]));
21  }
22  par(mfrow=c(1,2))
23  plot(1:n,bf01_pois_avg,main="BFSS01 (Black) vs IBF10 (Red) with Poisson samples ",xlab="Number of samples",
        ylab="BFSS01",type="l",ylim = c(min(bf01_pois_avg,ibf01_pois_avg),max(bf01_pois_avg,ibf01_pois_avg)))
24  mtext("Simulations: 100 / Samples: 30 / H0: TRUE",side=3)
25  lines(1:n,ibf01_pois_avg, type="l",col="red")
26  plot(1:n,bf01_geo_avg,main="BFSS01 (Black) vs IBF10 (Red) with Geometric samples",xlab="Number of samples",
        ylab="BFSS01",type="l", ylim=c(min(bf01_geo_avg,ibf01_geo_avg),max(bf01_geo_avg,ibf01_geo_avg)))
27  mtext("Simulations: 100 / Samples: 30 / H0: FALSE",side=3)
28  lines(1:n,ibf01_geo_avg, type="l",col="red")
29 }
30 Geometric_Poisson_SSBF(n=30,prob=0.5,lambda=1,simulations=100)

```




Figure 8–3: $AIBF_{10}$ and The IBF_{10} Lower Bound Comparison for the Poisson vs Geometric Test

Examining the presented figure above elucidates several significant observations. Most prominently, it is discernible that \underline{B}_{01}^I is a lower bound of the Average Intrinsic Bayes Factor (AIBF). Further, we note that the IBF lower bound is a more stable and robust form of evidence compared to the Arithmetic Intrinsic Bayes Factor (AIBF). In the scenario where data is generated under the alternative hypothesis, the AIBF presents lower values. However, these values do not furnish sufficient evidence to confidently reject the null hypothesis, H_0 . On the other hand, the lower bound presents a more compelling case. It succeeds in providing an accurate decision - one favoring the rejection of the null hypothesis - when we deal with a sample size exceeding ten. This finding underscores the valuable role that the IBF lower bound can play in hypothesis testing, particularly when dealing with larger sample sizes.

Moving forward, we now turn our attention to an empirical data-set that was previously studied by Cox (1962, p. 414) [6]. This data-set comprises a sample of 30 observations, which were drawn from a Poisson distribution. Notably, this Poisson distribution is characterized by a mean value of 0.8. We aim to delve into a comprehensive analysis of this data-set within the purview of our ongoing investigation.

TABLE 2
Example illustrating test for Poisson and geometric distributions

<i>Variate value</i>	<i>Observed frequency</i>	<i>Estimated Poisson frequency</i>	<i>Estimated geometric frequency</i>
0	12	12.61	16.07
1	11	10.93	7.46
2	6	4.74	3.47
3	1	1.37	1.61
≥ 4	0	0.35	1.40

Figure 8–4: Cox (1962) p. 414

The data utilized for this study was managed within the **R** programming environment. To ensure the robustness of our findings and to minimize the risk of any inherent bias or patterns influencing the results, we employed a randomization process on the data. This was achieved using the *sample()* function, a commonly used **R** command for random sampling. Subsequent to this randomization, we proceeded to compute \bar{B}'_{10} , specifically for the comparison between Poisson and Geometric distributions. The detailed steps followed in this computation are described next.

R program

```

1 bfss10_comparison = function(y,c_theoretical = 2/sqrt(pi),type="empirical",chart=FALSE) {
2   bfss10_empirical = c();
3   bfss10_theoretical = c();
4   s = 0;
5   n = length(y)
6   for (i in 1:n) {
7     s = s + y[i]
8     y_fact = factorial(y[1:i])
9     m1 = (i^(s + 0.5)*factorial(i-1)*prod(y_fact))
10    m0 = (gamma(i + s + 0.5))
11    bfss01= bfss01_mts(y[1:i],print=FALSE)
12    c_empirical = bfss01$bfss_max
13    bfss10_empirical = c(bfss10_empirical,c_empirical*m1/m0)
14    bfss10_theoretical= c(bfss10_theoretical,c_theoretical*m1/m0)
15  }
16  if (chart) {
17    plot(1:n,bfss10_empirical,type="l",
18        xlab="Number of samples",ylab="BFSS 10",
19        main="Poisson data with parameter 0.8 (Cox 1962) Table 2")
20  }
21  else {
22    if (type=="empirical") { return(bfss10_empirical) }
23    else if (type=="theoretical") { return(bfss10_theoretical) }
24  }
25 }
26 data = c(rep(0,12),rep(1,11),rep(2,6),3)
27 data = sample(data)
28 bfss10_comparison(y=data,type="theoretical",chart=TRUE)

```

Using the code above we obtained the following chart:




Figure 8-5: We used the randomized Cox data considered a growing vector with those samples and calculated the Bayes Factor as we added more data.

It is evident from the chart above that the IBF bound is consistent with the results obtained by Cox (1962)[6] and the IBF bound converges to the correct decision in favor of the adequate model which is the Poisson distribution. The graphical representation depicted above yields insightful observations about the performance of the IBF bound. Moreover, as sample sizes increase, the IBF bound shows convergence toward an optimal decision. It correctly leans in favor of the model that accurately reflects the underlying distribution of the data in this case, the Poisson. This empirical evidence reinforces the strength and consistency of the IBF as a tool for model selection.

8.2 The Negative Binomial vs. Poisson Separated Hypothesis Test

In this section, we'll focus on a separated hypothesis test comparing the Negative Binomial and Poisson distributions. The Negative Binomial distribution is a discrete probability distribution that models the number of failures in a sequence of independent and identically distributed Bernoulli trials, before a specified (non-random) number of successes (denoted r) occurs. An alternative formulation is to model the number of total trials (instead of the number of failures). In fact, for a specified (non-random) number of successes (r), the number of failures ($n - r$) is random because the total trials (n) are random. For example, we could use the negative binomial distribution to model the number of days n (random) a certain machine works (specified by r) before it breaks down. Imagine a sequence of independent Bernoulli trials: each trial has two potential outcomes called "success" and "failure". In each trial, the probability of success is θ and of failure is $(1 - \theta)$. We observe this sequence until a predefined number r of successes occurs. Then the random number of observed failures, Y follows the negative binomial distribution: $Y \sim NB(r, \theta)$. The pmf of Y is given by

$$P(Y = y|r, \theta) = \binom{y+r-1}{r-1} \theta^r (1-\theta)^y$$

where r is the number of successes, y is the number of failures, and θ is the probability of success on each trial. The likelihood of the Negative Binomial is given by

$$\prod_{i=1}^n \binom{y_i+r-1}{r-1} \theta^{rn} (1-\theta)^{\sum y_i}.$$

The maximum likelihood estimator for θ , the Fisher Information Matrix and the Jeffreys's prior for the Negative Binomial are:

$$\hat{\theta} = \frac{r}{\bar{y}+r}, \quad I(\theta) = \frac{r}{\theta(1-\theta)^2}, \quad \sqrt{|I(\theta)|} \propto \frac{r^{1/2}}{\theta^{1/2}(1-\theta)}$$

The marginal distribution in this case is:

$$\begin{aligned} m_1(\mathbf{y}) &\propto \int_0^1 \left[\prod_{i=1}^n \binom{y_i + r - 1}{r - 1} \right] \theta^{rn} (1 - \theta)^{\sum y_i} \frac{r^{1/2}}{\theta^{1/2} (1 - \theta)} d\theta \\ &= \left[\prod_{i=1}^n \binom{y_i + r - 1}{r - 1} \right] r^{1/2} \int_0^1 \theta^{rn+1-1/2-1} (1 - \theta)^{\sum y_i - 1} d\theta \end{aligned}$$

The change of variable $\beta = n\bar{y}$, $\alpha = rn + 1/2$ yields:

$$\int_0^1 \theta^{\alpha-1} (1 - \theta)^{\beta-1} d\theta = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}$$

The marginal above becomes:

$$m_1(\mathbf{y}) \propto \left[\prod_{i=1}^n \binom{y_i + r - 1}{r - 1} \right] r^{1/2} \frac{\Gamma(rn + 1/2)\Gamma(n\bar{y})}{\Gamma(n(r + \bar{y}) + 1/2)}$$

In this case, our null hypothesis is that M_0 is *Poisson*(λ), thus

$$m_0(\mathbf{y}) = \frac{c_0 \Gamma(\sum y_i + \frac{1}{2})}{n} (1/n)^{(\sum y_i + \frac{1}{2})} \prod_{i=1}^n y_i!$$

The Bayes factor can be expressed as:

$$B_{10}(\mathbf{y}) = \frac{c_1}{c_0} \left[\prod_{i=1}^n \binom{y_i + r - 1}{r - 1} y_i! \right] r^{1/2} \frac{\Gamma(rn + 1/2)\Gamma(n\bar{y})}{\Gamma(n(r + \bar{y}) + 1/2)\Gamma(n\bar{y} + 1/2)} n^{(n\bar{y} + 1/2)}$$

By definition, we have that

$$\binom{y_i + r - 1}{r - 1} y_i! = \frac{(y_i + r - 1)! y_i!}{(r - 1)! y_i!} = \frac{(y_i + r - 1)!}{(r - 1)!}$$

Hence, the Bayes factor can be expressed as:

$$B_{10}(\mathbf{y}) = \frac{c_1}{c_0} \left[\prod_{i=1}^n \frac{(y_i + r - 1)!}{(r - 1)!} \right] r^{1/2} \frac{\Gamma(rn + 1/2)\Gamma(n\bar{y})}{\Gamma(n(r + \bar{y}) + 1/2)\Gamma(n\bar{y} + 1/2)} n^{(n\bar{y} + 1/2)}$$

Similarly, the Bayes factor for a minimal training sample $y(\ell)$ is:

$$B_{10}(y(\ell)) = \frac{c_0}{c_1} \left[\frac{(r - 1)!}{(y(\ell) + r - 1)!} \right] r^{-1/2} \frac{\Gamma(r + y(\ell) + 1/2)\Gamma(y(\ell) + 1/2)}{\Gamma(r + 1/2)\Gamma(y(\ell))}$$

The Intrinsic Bayes Factor in this case is:

$$B_{10}(y(-\ell)|y(\ell)) = \frac{\left[\prod_{i=1}^n (y_i + r - 1)! \right] n^{(n\bar{y}+1/2)} \Gamma(r + y(\ell) + 1/2) \Gamma(y(\ell) + 1/2) \Gamma(rn + 1/2) \Gamma(n\bar{y})}{(y(\ell) + r - 1)! [(r - 1)!]^{n-1} \Gamma(n(r + \bar{y}) + 1/2) \Gamma(n\bar{y} + 1/2) \Gamma(r + 1/2) \Gamma(y(\ell))}$$

By using the identity $(y(\ell) + r - 1)! = \Gamma(y(\ell) + r)$, we found the Empirical SS Bayes factor upper bound for this test:

$$\bar{B}_{10}^I(\mathbf{y}) = \frac{\left[\prod_{i=1}^n (y_i + r - 1)! \right] n^{(n\bar{y}+1/2)} \Gamma(rn + 1/2) \Gamma(n\bar{y})}{[(r - 1)!]^{n-1} \Gamma(n(r + \bar{y}) + 1/2) \Gamma(n\bar{y} + 1/2) \Gamma(r + 1/2)} \sup_{\ell=1, \dots, L} \frac{\Gamma(r + y(\ell) + 1/2) \Gamma(y(\ell) + 1/2)}{\Gamma(y(\ell) + r) \Gamma(y(\ell))}$$

To obtain the supremum above, we must note that:

$$\Gamma(r + y(\ell) + 1/2) \Gamma(y(\ell) + 1/2) \geq \Gamma(y(\ell) + r) \Gamma(y(\ell)), \quad \forall y(\ell) = 0, 1, 2, \dots$$

Thus,

$$\sup_{\ell=1, \dots, L} \frac{\Gamma(r + y(\ell) + 1/2) \Gamma(y(\ell) + 1/2)}{\Gamma(y(\ell) + r) \Gamma(y(\ell))} = \frac{\Gamma(r + y_{(n)} + 1/2) \Gamma(y_{(n)} + 1/2)}{\Gamma(y_{(n)} + r) \Gamma(y_{(n)})}$$

The result above, allowed us to express the Empirical IBF upper bound as:

$$\bar{B}_{10}^I(\mathbf{y}) = \frac{\left[\prod_{i=1}^n (y_i + r - 1)! \right] n^{(n\bar{y}+1/2)} \Gamma(rn + 1/2) \Gamma(n\bar{y})}{[(r - 1)!]^{n-1} \Gamma(n(r + \bar{y}) + 1/2) \Gamma(n\bar{y} + 1/2) \Gamma(r + 1/2)} \frac{\Gamma(r + y_{(n)} + 1/2) \Gamma(y_{(n)} + 1/2)}{\Gamma(y_{(n)} + r) \Gamma(y_{(n)})}$$

Alternatively, we can express the Empirical IBF lower bound as:

$$\underline{B}_{01}^I(\mathbf{y}) = \frac{[(r - 1)!]^{n-1} \Gamma(n(r + \bar{y}) + 1/2) \Gamma(n\bar{y} + 1/2) \Gamma(r + 1/2)}{\left[\prod_{i=1}^n (y_i + r - 1)! \right] n^{(n\bar{y}+1/2)} \Gamma(rn + 1/2) \Gamma(n\bar{y})} \frac{\Gamma(y_{(1)} + r) \Gamma(y_{(1)})}{\Gamma(r + y_{(1)} + 1/2) \Gamma(y_{(1)} + 1/2)}$$

Similarly, the Theoretical IBF Lower bound is:

$$\underline{B}_{01}^I(\mathbf{y}) = \frac{[(r - 1)!]^{n-1} \Gamma(n(r + \bar{y}) + 1/2) \Gamma(n\bar{y} + 1/2)}{\left[\prod_{i=1}^n (y_i + r - 1)! \right] n^{(n\bar{y}+1/2)} \Gamma(rn + 1/2) \Gamma(n\bar{y})} \frac{\Gamma(r)}{\Gamma(1/2)}$$




Figure 8-6: Poisson vs. Negative Binomial IBF lower bound as the sample size increases with samples from H_0 .

```

1 bfss01 = function(y,m0="poisson",m1="nbinom",r=1) {
2   ybar = mean(y)
3   n = length(y)
4   lf1 = (n-1)*log(factorial(r-1)) + log(gamma( n*(r+ybar)+0.5 ))
5   + log(gamma(n*ybar + 0.5)) + log(gamma(r+0.5)) + log(gamma(r))
6   summ = 0
7   for (i in 1:n) { summ = log(factorial(y[i] + r - 1)) + summ }
8   lf2 = (n*ybar + 0.5)*log(n) + log(gamma(r*n+0.5)) + log(gamma(n*ybar)) + summ
9   exp(lf1)/exp(lf2)
10 }
11
12 n=30; y = c(); samples = c()
13 for (i in 10:n) {
14   samples = c(samples,rpois(n=1,lambda=1))
15   y = c(y,bfss01(samples,r=1))
16 }
17
18 plot(10:n,y,type="l",col="red",xlab="Sample size",
19 main="Bayes factor of Poisson vs Negative Binomial with Poisson data"
20 ,ylab="Bayes Factor")

```

8.3 Conclusions

In conclusion, our exploration of separated hypothesis testing has provided valuable insights. Upon comprehensive analysis involving Bayes factor bounds, numerical simulations, and real-world data, our findings suggest a promising application of these bounds in general hypothesis testing. However, a crucial consideration lies in the careful selection of distributions for M_0 and M_1 . It is imperative to choose these models in a manner that ensures $\inf B_{10}(\mathbf{y}(\ell)) \neq 0$, or equivalently, $\sup B_{01}(\mathbf{y}(\ell)) < \infty$. This selection criterion is pivotal as the effectiveness of the bound is contingent upon the appropriate identification of M_0 and M_1 and we used it in our examples, when we considered the Poisson as M_0 and the Geometric or Negative Binomial as M_1 . Without this careful choice, the bound's utility in hypothesis testing will be affected.

CHAPTER 9

GHOSH AND SAMANTA ALTERNATIVE APPROACH

This chapter explores a crucial 2001 study by J.K. Ghosh and T. Samanta [4], presenting two methodologies that, when specific conditions are met, align with the SS approach. One such method involves setting $y_0 = \arg \sup_{y(\ell)} m_0(y(\ell))$ and subsequently resolving the equation $cB_{01}(y_0) = 1$. The Bayes factor that arises from this method is:

$$\boxed{B_{10}^{GS}(\mathbf{y}) = cB_{10}(\mathbf{y}) = B_{01}(y_0)B_{10}(\mathbf{y})}.$$

9.1 The Normal Mean Hypothesis Test Example

Consider the following simple hypothesis test

$$H_0 : N(0, 1) \text{ vs. } H_1 : N(\mu, 1)$$

The marginal distributions for H_0 and H_1 are:

$$m_0(\mathbf{y}) = \left(\frac{1}{\sqrt{2\pi}}\right)^n \exp\left\{-\frac{\sum y_i^2}{2}\right\} \quad m_1(\mathbf{y}) = 1$$

To calculate the GS Bayes Factor, we calculated the following:

$$y_0 = \arg \sup_{y(\ell) \in D} m_0(y(\ell)) = \arg \sup_{y(\ell) \in D} \left(\frac{1}{\sqrt{2\pi}}\right) \exp\left\{-\frac{y(\ell)^2}{2}\right\} = 0$$

Additionally, we must solve this equation:

$$cB_{10}(y_0) = 1 \implies c = B_{01}(y_0) = \frac{1}{\sqrt{2\pi}}$$

The Bayes Factor in this case is:

$$B_{10}(\mathbf{y}) = (\sqrt{2\pi})^n \exp\left\{\frac{\sum y_i^2}{2}\right\}$$

Similarly, the GS Bayes Factor is:

$$B_{10}^{GS}(\mathbf{y}) = \frac{1}{\sqrt{2\pi}} (\sqrt{2\pi})^n \exp\left\{\frac{\sum y_i^2}{2}\right\} = (\sqrt{2\pi})^{(n-1)} \exp\left\{\frac{\sum y_i^2}{2}\right\}.$$

In contrast, the IBF upper bound for this test is:

$$\bar{B}_{10}^I(\mathbf{y}) = B_{10}^N(\mathbf{y}) \frac{1}{\sqrt{2\pi}}$$

Therefore,

$$\bar{B}_{10}^I(\mathbf{y}) = B_{10}^{GS}(\mathbf{y}).$$

9.2 The Exponential Hypothesis Test Example

Consider the following simple exponential hypothesis test

$$H_0 : \lambda = \lambda_0 \text{ vs. } H_1 : \lambda \neq \lambda_0$$

In the exponential case, for a sample \mathbf{y} of size n , the marginal distribution is

$$m(\mathbf{x}) \propto \int_0^\infty \lambda^{(n-1)} \exp(-\lambda \sum_{i=1}^n y_i) d\lambda = \frac{(\sum_{i=1}^n y_i)^n}{\Gamma(n)}$$

The marginal distributions for this hypothesis test are:

$$m_0(\mathbf{y}) = \lambda_0^n \exp(-\lambda_0 \sum_{i=1}^n y_i) \quad m_1(\mathbf{y}) = \frac{(\sum_{i=1}^n y_i)^n}{\Gamma(n)}$$

To calculate we GS Bayes factor, we found that:

$$y_0 = \arg \sup_{y(\ell) \in D} m_0(y(\ell)) = \arg \sup_{y(\ell) \in D} \lambda_0 \exp(-\lambda_0 y(\ell)) = 0$$

The next step is to solve the following equation:

$$cB_{10}(y_0) = 1 \implies c = B_{01}(y_0) = \lambda_0 \exp(-\lambda_0 y_0) \frac{1}{y_0} = \infty$$

Evidently, the GS Bayes Factor does not exist for this case.

9.2.1 Conclusions

The resemblance between the Ghost and Samanta (GS) approach and the IBF upper bound is evident. It is clear that if the marginal distribution $m_1(y(\ell))$ is independent of $y(\ell)$, then the two methods yield identical results. In other words, $B_{10}^{GS} = \bar{B}_{10}^I$. This equality arises due to the following reasoning:

$$\begin{aligned}
 B_{10}^{GS} &= B_{01}(y_0)B_{10}(\mathbf{y}) \\
 &= B_{01}(\arg \sup_{y(\ell) \in D} m_0(y(\ell)))B_{10}(\mathbf{y}) \\
 &= B_{01}(\arg \sup_{y(\ell) \in D} [m_0(y(\ell))/m_1(y(\ell))])B_{10}(\mathbf{y}) \\
 &= \sup_{y(\ell) \in D} B_{01}(y(\ell))B_{10}(\mathbf{y}) = \bar{B}_{10}^I
 \end{aligned}$$

However, it is important to note that in certain cases, such as the exponential distribution, the IBF bound provides accurate results while the GS Bayes factor does not exist. Further research can be conducted to compare the suitability of each approach for different hypothesis tests in specific scenarios. Additionally, efforts can be made to identify general conditions that dictate the existence of the IBF bound and GS Bayes factors. Such investigations will contribute to a deeper understanding of the applicability and limitations of these alternative methodologies.

CHAPTER 10

RESEARCH CONTRIBUTIONS

My doctoral thesis makes significant contributions to the field of Bayes factors theory. In this research, I introduce a new universal robust bound for Bayes factors that is an alternative to previous methods and is applicable to a variety of models. The bound is based on the SS Bayes Factor, which is a measure of evidence in favor of one hypothesis compared to another. This research significantly advances our understanding of Intrinsic Bayes Factors (IBF), presenting a noteworthy contribution to the field. Specifically, we have established stringent lower/upper bounds on all types of IBF - arithmetic, geometric, and median.

Our proposed bound surpasses the existing benchmark of $-ep \log(p)$ in two crucial aspects. Firstly, it shares identical asymptotic behavior as a Bayes Factor, adhering more closely to the foundational tenets of Bayesian analysis. Secondly, and more importantly, our lower bound corresponds to an actual prior distribution. This prior can be interpreted as a form of least favorable prior for the null hypothesis, embodying the spirit of conservatism intrinsic to statistical hypotheses testing.

In contrast to other IBFs, our derived variant exhibits computational efficiency. The ease of computation is particularly noteworthy given the generally arduous computational demands of other IBFs. Therefore, our lower bound provides not only a stronger mathematical formulation but also a more practical tool for researchers employing Bayesian methods. This approach bridged the gap between the Intrinsic Bayes Factors Theory with the SS Bayes Factors and compared the arithmetic Intrinsic Bayes Factor with the IBF

bounds. Moreover, I developed the "Empirical IBF Bounds" methodology, resulting in a novel Empirical IBF lower bound that is even better than the theoretical bound.

To develop the Empirical Intrinsic Bayes Factor Bound, I constructed a mathematical formula and an algorithm to determine the set of all possible empirical minimal training samples and compared the empirical and theoretical Bayes Factor bounds. The thesis also contains a comparison of the lower bound with the Sellke et. al (1990) [13] $-ep \log(p)$ universal lower bound and the new Empirical and Theoretical SP Bayes Factors that are defined for the first time.

The thesis includes several theorems that demonstrate the effectiveness of these new methods. For example, **Theorem 3.1** shows that the IBF lower bound gets better than the $-ep \log(p)$ universal lower bound as the sample size increases under the null hypothesis, **Theorem 3.2** states and shows that using our novel Empirical SP Prior, we can obtain our Empirical IBF Bound, and **Theorem 3.3** states and shows that our novel Empirical SP Prior and SP Bayes factor, converges to the Theoretical SP Prior and SP Bayes Factor. This work also includes several examples and simulations, such as the Normal and Exponential distribution, ANOVA, and Linear models. In section 3.5, I generalized the technique used to compare the AIBF with the IBF Bound, and I showed that the same can be said for any other kind of Bayes factor that uses any measure of central tendency.

I also derived and calculated various Bayes factors and priors for linear models and specifically for ANOVA. I developed the necessary calculations for the Theoretical and Empirical Generalized IBF Bounds in General linear models by introducing a generalized prior to obtain both empirical and theoretical IBF bounds for two nested linear models and ANOVA. Using this generalization, we derived the ANOVA Bayes factors for different priors such as the Full Jeffreys prior and the Modified Jeffrey's prior. Furthermore, through these developments, I was able to demonstrate that the Generalized SS Bayes factor for ANOVA does not exist for the Reference Prior.

In this work, we also calculated the EP-Priors and EP-Bayes factor for the exponential distribution and the hypothesis test $\lambda = \lambda_0$ vs $\lambda \neq \lambda_0$ and compared their results with the between the SP Priors and the SP Bayes factor through numerical simulations and charts. We found that the results are similar but the prior obtained in the SP-Prior approach does not depend on the samples, as in the EP-Prior and it's much easier to calculate while their difference is small. Additionally, for the exponential case, the GS Bayes Factor was calculated in section 4.2 and compared to the SS Bayes factor and we've concluded that the GS Bayes factor does not exist while the IBF Bound both the theoretical and the empirical are well-defined and consistent. In the Geometric vs Poisson Separated models example, we calculated the AIBF and we showed the usefulness consistency, and stability of our lower bound of the AIBF.

These contributions have the potential to improve the accuracy and efficiency of statistical modeling methods, making them more accessible to researchers and practitioners in various fields. Ultimately, this research provides a valuable contribution to the field of Bayes factors theory and statistical modeling, with the potential to advance scientific knowledge and improve decision-making in a variety of domains.

REFERENCE LIST

- [1] Atkinson, A. C. (1978). Posterior probabilities for choosing a regression model. *Biometrika* 65(1), 39–48. Accessed 20 Dec. 2023.
- [2] Berger, J. O. and J. M. Bernardo (1992). Ordered group reference priors with applications to a multinomial problem. *Biometrika* 79, 25–37.
- [3] Berger, J. O. and L. R. Pericchi (1996, March). The intrinsic bayes factor for model selection and prediction. *Journal of the American Statistical Association* 91(433), 109–122.
- [4] Berger, J. O., L. R. Pericchi, J. K. Ghosh, T. Samanta, and F. De Santis (2001). Objective bayesian methods for model selection: Introduction and comparison. *Lecture Notes-Monograph Series* 38, 135–207.
- [5] Berger, J. O., L. R. Pericchi, and J. A. Varshavsky (1998). Bayes factors and marginal distributions in invariant situations. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)* 60(3), 307–321.
- [6] Cox, D. R. (1962). Further results on tests of separate families of hypotheses. *Journal of the Royal Statistical Society, Series B* 24, 406–424.
- [7] Jeffreys, H. (1961). *Theory of Probability* (3rd ed.). Oxford: Oxford University Press.
- [8] Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association* 90(430), 773–795.
- [9] Lempers, F. B. (1971). *Posterior probabilities of alternative linear models. Some theoretical considerations and empirical experiments*. Rotterdam: Rotterdam University Press.
- [10] Pérez, J. M. and J. O. Berger (2002). Expected-posterior prior distributions for model selection. *Biometrika* 89(3), 491–511.

- [11] Pericchi, L. (2005). Model selection and hypothesis testing based on objective probabilities and bayes factors. In *Handbook of Statistics*, Volume 25, pp. 115–149. Elsevier.
- [12] Self, S. G. and K. Y. Liang (1987). Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *Journal of the American Statistical Association* 82(398), 605–610.
- [13] Sellke, T., M. J. Bayarri, and J. O. Berger (2001). Calibration of p values for testing precise null hypotheses. *The American Statistician* 55, 62–71.
- [14] Smith, A. and D. Spiegelhalter (1980). Bayes factors and choice criteria for linear models. *Journal of the Royal Statistical Society. Series B (Methodological)* 42(2), 213–220.
- [15] Spiegelhalter, D. and A. Smith (1982). Bayes factors for linear and log-linear models with vague prior information. *Journal of the Royal Statistical Society. Series B (Methodological)* 44(3), 377–387.
- [16] Wasserstein, R., A. Schirm, and N. Lazar (2019). Moving to a world beyond $p < 0.05$. *American Statistician* 73, 1–19.
- [17] Wasserstein, R. L. and N. A. Lazar (2016). The asa statement on p-values: Context, process, and purpose. *The American Statistician* 70(2), 129–133.

A UNIVERSAL AND ROBUST BOUND FOR THE INTRINSIC BAYES FACTORS

Richard Anthony Clare Morales

richard.clare@upr.edu

Department of Mathematics

Chair: Luis Raúl Pericchi Guerra Ph.D

María E. Pérez Hernández Ph.D

Grade: Doctor of Philosophy in Mathematics

Date of Graduation: Feb 2024