
Bayesian Inference

Before discussing Bayesian inference, we recall the fundamental problem of statistics: “The fundamental problem towards which the study of Statistics is addressed is that of inference. Some data are observed and we wish to make statements, inferences, about one or more unknown features of the physical system which gave rise to these data” (O’Hagan, 2010). Upon more careful consideration of the foundations of statistics we find many different schools of thought. Even leaving aside those that are collectively known as classical statistics, this leaves several choices: objective and subjective Bayes, fiducialist inference, likelihood based methods, and more.¹

This diversity is not unexpected! Deriving the desired inference on parameters and models from the data is a problem of induction, which is one of the most controversial problems in philosophy. Each school of thought follows its own principles and methods to lead to statistical inference. Berger (1984) describes this as: “Statistics needs a: ‘foundation’, by which I mean a framework of analysis within which any statistical investigation can theoretically be planned, performed, and meaningfully evaluated. The words ‘any’ and ‘theoretically’ are key, in that the framework should apply to any situation but may only theoretically be implementable. Practical difficulties or time limitations may prevent complete (or even partial) utilisation of such framework, but the direction in which ‘truth’ could be found would at least be known”. The foundations of Bayesian inference are better understood when seen in contrast to those of its mainstream competitor, classical inference.

¹ Subjective Bayes is essentially the subject of this volume. In addition to these schools of thought, there are even half-Bayesians who accept the use of a priori information but believe that probability calculus is inadequate to combine prior information with data, which should instead be replaced by a notion of causal inference.

1.1 The Classical Paradigm

Classical statistics seeks to make inference about a population starting from a sample. Let x (or $x = (x_1, x_2, \dots, x_n)$, where n is a sample size,) denote the data. The set \mathcal{X} of possible samples x is known as the sample space, usually $\mathcal{X} \subseteq \mathbb{R}^n$. Underlying classical inference is the recognition of variability across samples, keeping in mind that the observed data are only one of many – possibly infinitely many – data sets that could have been observed. The interpretation of the data depends not only on the observed data, but also on the assumptions put forward about the process generating the observable data. As a consequence, the data are treated as a realization of a random variable or a random vector X with a distribution F_θ , which of course is not entirely known. However, there is usually some knowledge (theoretical considerations, experimental evidence, etc.) about the nature of the chance experiment under consideration that allow one to conjecture that F_θ is a member of a family of distributions \mathcal{F} . This family of distributions becomes the statistical model for X . The assumption of a model is also known as the *model specification* and is an essential part of developing the desired inference.

Assuming that X is a continuous random variable or random vector, it is common practice to represent the distributions \mathcal{F} by their respective density functions. When the density functions are indexed by a parameter θ in a parameter space Θ , the model can be written as $\mathcal{F} = \{f(x | \theta), x \in \mathcal{X} : \theta \in \Theta\}$. In many cases, the n variables (X_1, X_2, \dots, X_n) are assumed independent conditional on θ and the statistical model can be written in terms of the marginal densities of $X_i, i = 1, 2, \dots, n$:

$$\mathcal{F} = \{f(x | \theta) = \prod_{i=1}^n f_i(x_i | \theta) : \theta \in \Theta\}, x \in \mathcal{X},$$

and $f_i(\cdot | \theta) = f(\cdot | \theta), i = 1, 2, \dots, n$, if additionally the variables X_i are assumed to be identically distributed. The latter is often referred to as random sampling.

Beyond the task of modeling and parametrization, classical inference includes many methods to extract conclusions about the characteristics of the model that best represents the population and tries to answer questions like the following: (1) Are the data x compatible with a family \mathcal{F} ? (2) Assuming that the specification is correct and that the data are generated from a model in the family \mathcal{F} , what conclusions can be drawn about the parameter θ_0 that indexes the distribution F_θ that “appropriately” describes the phenomenon under study?

Classical methods – also known as frequentist methods – are evaluated

under the principle of repeated sampling, that is, with respect to the performance under infinitely many hypothetical repetitions of the experiment carried out under identical conditions. One of the aspects of this principle is the use of frequencies as a measure of uncertainties, that is, a frequentist interpretation of probability. See , Paulino et al. (2018, section 1.2), for a review of this and other interpretations of probability.

In the case of parametric inference, in answer to question (2) above, we need to consider first the question of point estimation, which, *grosso modo*, is: Given a sample $X = (X_1, X_2, \dots, X_n)$, how should one “guess,” estimate, or approximate the true value θ , through an estimator $T(X_1, X_2, \dots, X_n)$. The estimator should have the desired properties such as unbiasedness, consistency, sufficiency, efficiency, etc.

For example, with $X \equiv \mathbb{R}^n$, the estimator $T(X_1, X_2, \dots, X_n)$ based on a random sample is said to be centered or unbiased if

$$E\{T \mid \theta\} = \int_{\mathbb{R}^n} T(x_1, x_2, \dots, x_n) \prod_{i=1}^n f(x_i \mid \theta) dx_1 dx_2 \dots dx_n = \theta, \quad \forall \theta \in \Theta.$$

This is a property related to the principle of repeated sampling, as can be seen by the fact that it includes integration over the sample space (in this case \mathbb{R}^n). Considering this entire space is only relevant if one imagines infinitely many repetitions of the sampling process or observations of the n random variables (X_1, X_2, \dots, X_n) . The same applies when one considers other criteria for evaluation of estimators within the classical paradigm. In other words, implicit in the principle of repeated sampling is a consideration of what might happen in the entire sample space.

Parametric inference often takes the form of confidence intervals. Instead of proposing a single value for θ , one indicates an interval whose endpoints are a function of the sample,

$$(T^*(X_1, X_2, \dots, X_n), T^{**}(X_1, X_2, \dots, X_n)),$$

and which covers the true parameter value with a certain probability, preferably a high probability (typically referred to as the confidence level),

$$P\{T^*(X_1, X_2, \dots, X_n) < \theta < T^{**}(X_1, X_2, \dots, X_n) \mid \theta\} = 1 - \alpha,$$

$0 < \alpha < 1$. This expression pre-experimentally translates a probability of covering the unknown value θ to a random interval (T^*, T^{**}) whose lower and upper limits are functions of (X_1, X_2, \dots, X_n) and, therefore, random variables. However, once a specific sample is observed (i.e., post-experimentally) as n real values, (x_1, x_2, \dots, x_n) , this becomes a specific

interval on the real line (now with real numbers as lower and upper limits).

$$(T^*(x_1, x_2, \dots, x_n), T^{**}(x_1, x_2, \dots, x_n)),$$

and the probability

$$P\{T^*(x_1, x_2, \dots, x_n) < \theta < T^{**}(x_1, x_2, \dots, x_n) \mid \theta\} = 1 - \alpha,$$

$0 < \alpha < 1$, is no longer meaningful. In fact, once θ has an unknown, but fixed, value, this probability can only be 1 or 0, depending upon whether the true value of θ is or is not in the real interval

$$(T^*(x_1, x_2, \dots, x_n), T^{**}(x_1, x_2, \dots, x_n)).$$

Of course, since θ is unknown, the investigator does not know which situation applies. However, a classical statistician accepts the frequentist interpretation of probability and invokes the principle of repeated sampling in the following way: If one imagines a repetition of the sampling and inference process (each sample with n observations) a large number of times, then in $(1 - \alpha)$ 100% of the repetitions the numerical interval will include the value of θ .

Another instance of classical statistical inference is a parametric hypothesis test. In the course of scientific investigation one frequently encounters, in the context of a certain theory, the concept of a hypothesis about the value of one (or multiple) parameter(s), for example in the symbols

$$H_0 : \theta = \theta_0.$$

This raises the following fundamental question: Do the data (x_1, x_2, \dots, x_n) support or not support the proposed hypothesis? This hypothesis is traditionally referred to as the null hypothesis. Also here the classical solution is again based on the principle of repeated sampling if one follows the Neyman–Pearson theory. It aims to find a rejection region W (critical region) defined as a subset of the sample space, $W \subset \mathcal{X}$, such that

$$(X_1, X_2, \dots, X_n) \in W \Rightarrow \text{rejection of } H_0,$$

$$(X_1, X_2, \dots, X_n) \notin W \Rightarrow \text{fail to reject } H_0.$$

The approach aims to control the probability of a type-I error,

$$P\{(X_1, X_2, \dots, X_n) \in W \mid H_0 \text{ is true}\},$$

and minimize the probability of a type-II error,

$$P\{(X_1, X_2, \dots, X_n) \notin W \mid H_0 \text{ is false}\}.$$

What does it mean that the critical region is associated with a type-I error, equal to, for example, 0.05? The investigator can not know whether a false or true hypothesis is being rejected when a particular observation falls into the critical region and the hypothesis is thus rejected. However, being a classical statistician the investigator is convinced that under a large number of repetitions and if the hypothesis were true, then only in 5% of the cases would the observation fall into the rejection region. What does it mean that the critical region is associated with a type-II error equal to, say 0.10? Similarly, when a particular observation is not in the rejection region and thus the hypothesis is not rejected, then the investigator cannot know whether a true or false hypothesis is being accepted. Being a classical statistician, the investigator can affirm that under a large number of repetitions of the entire process and if the hypothesis were in fact false, only in 10% of the cases would the observation not fall into the rejection region.

In the following discussion, it is assumed that the reader is familiar with at least the most elementary aspects of how classical inference approaches estimation and hypothesis testing, which is therefore not discussed here in further detail.

1.2 The Bayesian Paradigm

For Lindley, the substitution of the classical paradigm by the Bayesian paradigm represents a true scientific revolution in the sense of Kuhn (1962) The initial seed for the Bayesian approach to inference problems was planted by Richard Price when, in 1763, he posthumously published the work of Rev. Thomas Bayes titled *An Essay towards Solving a Problem in the Doctrine of Chances*. An interpretation of probability as a degree of belief – fundamental in the Bayesian philosophy – has a long history, including J. Bernoulli, in 1713, with his work *Ars Conjectandi*. One of the first authors to define probabilities as a degree of beliefs in the truth of a given proposition was De Morgan, in *Formal Logic*, in 1847, who stated: (1) probability is identified as a degree of belief; (2) the degrees of belief can be measured; and (3) these degrees of belief can be identified with a certain set of judgments. The idea of coherence of a system of degrees of belief seems to be due to Ramsey, for whom the behavior of an individual when betting on the truth of a given proposition is associated with the degree of belief that the individual attaches to it. If an individual states odds or possibilities (chances) – in favor of the truth or untruth – as $r : s$, then the degree of belief in the proposition is, for this individual, $r/(r + s)$. For Ramsey, no set of bets in given propositions is admissible for a coherent individual if it would

lead to certain loss. The strongest exponent of the concept of personal probabilities is, however, de Finetti. In discussing the Bayesian paradigm and its application to statistics, one must also cite Harold Jeffreys, who, reacting to the predominantly classical position in the middle of the century, besides inviting disapproval, managed to resurrect Bayesianism, giving it a logical basis and putting forward solutions to statistical inference problems in his time. From there the number of Bayesians grew rapidly and it becomes impossible to mention all but the most influential – perhaps Good, Savage, and Lindley.

The well-known Bayes' theorem is a proposition about conditional probabilities. It is simply probability calculus and is thus not subject to any doubts. Only the application to statistical inference problems is subject to some controversy. It obviously plays a central role in Bayesian inference, which is fundamentally different from classical inference. In the classical model, the parameter θ , $\theta \in \Theta$, is an unknown but fixed quantity, i.e., it is a particular value that indexes the sampling model or family of distributions \mathcal{F} that “appropriately” describes the process or physical system that generates the data. In the Bayesian model, the parameter θ , $\theta \in \Theta$, is treated as an unobservable random variable. In the Bayesian view, any unknown quantity – in this case, the parameter θ – is uncertain and all uncertainties are described in terms of a probability model. Related to this view, Bayesians would argue that initial information or *a priori* information – prior or external to the particular experiment, but too important to be ignored – must be translated into a probability model for θ , say $h(\theta)$, and referred to as the prior distribution. The elicitation and interpretation of prior distributions are some of the most controversial aspects of Bayesian theory.

The family \mathcal{F} is also part of the Bayesian model; that is, the sampling model is a common part of the classical and the Bayesian paradigms, except that in the latter the elements $f(x | \theta)$ of \mathcal{F} are in general assumed to also have a subjective interpretation, similar to $h(\theta)$.

The discussion of prior distributions illustrates some aspects of the disagreement between Bayesian and classical statisticians. For the earlier, Berger, for example, the subjective choice of the family \mathcal{F} is often considered a more drastic use of prior information than the use of prior distributions. And some would add: In the process of modeling, a classical statistician uses prior information, albeit in a very informal manner. Such informal use of prior information is seen critically under a Bayesian paradigm, which would require that initial or prior information of an investigator needs to be formally stated as a probability distribution on the random variable θ . Classical statisticians, for example, Lehmann, see an

important difference between the modeling of \mathcal{F} and the specification of $h(\theta)$. In the earlier case one has a data set $x = (x_1, x_2, \dots, x_n)$ that is generated by a member of \mathcal{F} and can be used to test the assumed distribution.

To understand the Bayesian point of view, recall that for a classical statistician *all* problems that involve a binomial random variable X can be reduced to a Bernoulli model with an unknown parameter θ that represents a “success” probability. For Bayesians, each problem is *unique* and has its own real context where θ is an important quantity about which there is, in general, some level of knowledge that might vary from problem to problem and investigator to investigator. Thus, the probability model that captures this variability is based on *a priori* information and is specific to a given problem and a given investigator. In fact, *a priori* information includes personal judgements and experiences of most diverse types, resulting from in general not replicable situations, and can thus only be formalized in subjective terms. This formalism requires that the investigator comply with coherence or consistency conditions that permit the use of probability calculus. However, different investigators can in general use different prior distributions for the same parameter without violating coherence conditions.

Assume that we observe $X = x$ and are given some $f(x | \theta) \in \mathcal{F}$ and a prior distribution $h(\theta)$. Then Bayes’ theorem implies²

$$h(\theta | x) = \frac{f(x | \theta)h(\theta)}{\int_{\Theta} f(x | \theta)h(\theta) d\theta}, \quad \theta \in \Theta, \quad (1.1)$$

where $h(\theta | x)$ is the posterior distribution of θ after observing $X = x$. Here, the initial information of the investigator is characterized by $h(\theta)$, and modified with the observed data by being updated to $h(\theta | x)$. The denominator in (1.1), denoted $f(x)$, is the marginal (or prior predictive) distribution for X ; that is, for an observation of X whatever the value of θ .

The concept of a likelihood function appears in the context of classical inference, and is not less important in the Bayesian context. Regarding its definition, it is convenient to distinguish between the discrete and continuous cases (Kempthorn and Folks, 1971), but both cases lead to the function of θ ,

$$\begin{aligned} L(\theta | x) &= kf(x | \theta), \quad \theta \in \Theta \quad \text{or} \\ L(\theta | x_1, \dots, x_n) &= k\Pi_i f(x_i | \theta), \quad \theta \in \Theta, \end{aligned} \quad (1.2)$$

which expresses for every $\theta \in \Theta$ its likelihood or plausibility when $X = x$ or $(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)$ is observed. The symbol k represents a

² Easily adapted if x were a vector or if the parameter space were discrete.

factor that does not depend on θ . The likelihood function – it is not a probability, and therefore, for example, it is not meaningful to add likelihoods – plays an important role in Bayes' theorem as it is the factor through which the data, x , updates prior knowledge about θ ; that is, the likelihood can be interpreted as quantifying the information about θ that is provided by the data x .

In summary, for a Bayesian the posterior distribution contains, by way of Bayes' theorem, all available information about a parameter:

prior information + information from the sample.

It follows that all Bayesian inference is based on $h(\theta | x)$ [or $h(\theta | x_1, x_2, \dots, x_n)$].

When θ is a parameter vector, that is, $\theta = (\gamma, \phi) \in \Gamma \times \Phi$, it can be the case that the desired inference is restricted to a subvector of θ , say γ . In this case, in contrast to the classical paradigm, the elimination of the nuisance parameter ϕ under the Bayesian paradigm follows always the same principle, namely through the marginalization of the joint posterior distribution,

$$h(\gamma | x) = \int_{\Phi} h(\gamma, \phi | x) d\phi = \int_{\Phi} h(\gamma | \phi, x) h(\phi | x) d\phi. \quad (1.3)$$

Possible difficulties in the analytic evaluation of the marginal disappear when γ and ϕ are *a priori* independent and the likelihood function factors into $L(\theta | x) = L_1(\gamma | x) \times L_2(\phi | x)$, leading to $h(\gamma | x) \propto h(\gamma) L_1(\gamma | x)$.

1.3 Bayesian Inference

In the Bayesian approach, it is convenient to distinguish between two objectives: (1) inference about unknown parameters θ , and (2) inference about future data (prediction).

1.3.1 Parametric Inference

In the case of inference on parameters, we find a certain agreement – at least superficially – between classical and Bayesian objectives, although in the implementation the two approaches differ. On one side, classical inference is based on probabilities associated with different samples, x , that could be observed under some fixed but unknown value of θ . That is, inference is based on sampling distributions that “weigh” probabilistically the values that a variable X or statistic $T(X)$ can assume across the sample

space. On the other hand, Bayesian inference is based on subjective probabilities or *a posteriori* credibilities associated with different values of the parameter θ and conditional on the particular observed x value. The main point is that x is fixed and known and θ is uncertain.

For example, once x is observed, a Bayesian being asked about the hypothesis $\{\theta \leq 0.5\}$ would directly address the question by calculating $P(\theta \leq 0.5 \mid x)$ based on $h(\theta \mid x)$, i.e., without leaving probability calculus. In contrast, a classical statistician would not directly answer the question. Stating, for example, that the hypothesis $H_0 : \theta \leq 0.5$ is rejected at significance level 5% does not mean that its probability is less than 0.05, but that if the hypothesis H_0 were true, (i.e., if in fact $\theta \leq 0.5$), then the probability of X falling into a given rejection region W would be $P(X \in W \mid \theta \leq 0.5) < 0.05$, and if in fact $x \in W$, then the hypothesis is rejected.

In O'Hagan's words (O'Hagan, 2010), while a Bayesian can state probabilities about the parameters, which are considered random variables, this is not possible for a classical statistician, who uses probabilities on data and not on parameters and needs to restate such probabilities such that they seem to say something about the parameter. The question is also related to a different view of the sample space. For a classical statistician, the concept of the sample space is fundamental, as repeated sampling would explore the entire space. A Bayesian would start by objecting to the reliance on repeated sampling and would assert that only the actually observed value x is of interest and not the space that x belongs to, which could be totally arbitrary, and which contains, besides x , observations that could have been observed, but were not.³

In estimation problems a classical statistician has several alternatives for functions of the data – estimators – whose sampling properties are investigated under different perspectives (consistency, unbiasedness, etc.). For a Bayesian there is only *one* estimator, which specifically is the posterior distribution $h(\theta \mid x)$. One can, of course, summarize this distribution in different ways, using mode, mean, median, or variance. But this is unrelated to the problem facing a classical statistician, who has to find a so-called *optimal estimator*. For a Bayesian such a problem only exists in the context of decision theory, an area in which the Bayesian view has a clear advantage over the classical view. Related to this, Savage claims that in past decades

³ The irrelevance of the sample space also leads to the same issue about stopping rules, something which Mayo and Kruse (2001) note, recalling Armitage, could cause problems for Bayesians.

the central problem in the face of uncertainty is shifting from *which inference one should report*, to *which decision should be taken*. As individual decisions have been considered outdated by some philosophers, we have also recently seen a resurgence of the Bayesian approach in the context of group decisions.

Under a Bayesian approach, confidence intervals are replaced by credible intervals (or regions). Given x , and once a posterior distribution is determined, one finds a credible interval for a parameter θ (assume, for the moment, a scalar). The interval is formed by two values in θ , say $[\underline{\theta}(x), \bar{\theta}(x)]$, or simpler, $(\underline{\theta}, \bar{\theta})$, such that

$$P(\underline{\theta} < \theta < \bar{\theta} | x) = \int_{\underline{\theta}}^{\bar{\theta}} h(\theta | x) d\theta = 1 - \alpha, \quad (1.4)$$

where $1 - \alpha$ (usually 0.90, 0.95, or 0.99) is the desired level of credibility. If $\Theta = (-\infty, +\infty)$, then one straightforward way of constructing a (in this case, central) credible interval is based on tails of the posterior distribution such that

$$\int_{-\infty}^{\underline{\theta}} h(\theta | x) d\theta = \int_{\bar{\theta}}^{+\infty} h(\theta | x) d\theta = \frac{\alpha}{2}. \quad (1.5)$$

Equation (1.4) has an awkward implication: The interval $(\underline{\theta}, \bar{\theta})$ is not unique. It could even happen that the values θ in the reported interval have less credibility than values θ outside the same interval. Therefore, to proceed with the choice of an interval that satisfies (1.4) and at the same time is of minimum size, Bayesians prefer to work with HPD (*highest posterior density*) credible sets $A = \{\theta : h(\theta | x_1, x_2, \dots, x_n) \geq k(\alpha)\}$, where $k(\alpha)$ is the largest real number such that $P(A) \geq 1 - \alpha$. For a unimodal posterior, the set becomes a HPD credible interval.

Credible sets have a direct interpretation in terms of probability. The same is not true for confidence intervals, which are based on a probability not related to θ , but rather a probability related to the data; more specifically, they are random intervals based on a generic sample, and which after observing a particular sample become a *confidence* of covering the unknown value θ by the resulting numerical interval. In general, this can not be interpreted as a probability or credibility about θ . Besides other critical aspects of the theory of confidence intervals (or regions), there are the ironical comments of Lindley (1990), who says to know various axioms of probability – for example, those due to Savage, de Finetti, or Kolmogorov – but no axiomatic definition of *confidence*.

For example, when a Bayesian investigates a composite hypothesis H_0 :

$\theta \in \Theta_0$ versus a composite alternative $H_1 : \theta \in \Theta_1$, with $\Theta_0 \cap \Theta_1 = \emptyset$, $\Theta_0 \cup \Theta_1 = \Theta$, she or he uses expressions in terms of probabilities on θ . When the investigator possesses a distribution $h(\theta)$, $\theta \in \Theta$, representing the initial credibility attributed to different parameter values, her or his prior probabilities of the competing hypotheses are determined by

$$P(\Theta_0) = \int_{\Theta_0} h(\theta) d\theta, \quad P(\Theta_1) = \int_{\Theta_1} h(\theta) d\theta.$$

The ratio $P(\Theta_0)/P(\Theta_1)$ is known as the *prior odds* for H_0 versus H_1 . After the experiment resulting in the observations x , and after determining $h(\theta | x)$, a Bayesian statistician calculates the corresponding posterior probabilities

$$P(\Theta_0 | x) = \int_{\Theta_0} h(\theta | x) d\theta, \quad P(\Theta_1 | x) = \int_{\Theta_1} h(\theta | x) d\theta,$$

and usually also the posterior odds for H_0 versus H_1 , that is, $P(\Theta_0 | x)/P(\Theta_1 | x)$. One can therefore say that in the Bayesian framework the inference outcome is not so much the acceptance or rejection of a the hypothesis H_0 – as is the case in the Neyman–Pearson framework – but rather the updating of the plausibility that is attributed to the competing hypotheses. Bayesian inference can be described as a comparison of posterior odds versus prior odds through

$$B(x) = \frac{P(\Theta_0 | x)/P(\Theta_1 | x)}{P(\Theta_0)/P(\Theta_1)}, \quad (1.6)$$

which is known as the Bayes factor in favor of H_0 (or Θ_0). The Bayes factor quantifies the evidence in the data x in favor of H_0 . Of course, the larger the Bayes factor, the larger is the increase of the posterior odds relative to the prior odds and thus the support that the data give to the hypothesis H_0 . In general, the Bayes factor depends on the prior distribution and can be expressed as a ratio of likelihoods weighted by the prior distributions conditional on the respective hypothesis on Θ_0 and Θ_1 (see also Paulino et al., 2003). In this sense one can not say that the Bayes factor is a measure of the support for H_0 based on *only* the data.

When the hypothesis about θ is specific to the point of being defined as $H_0 : \theta = \theta_0$, evaluation of a Bayes factor or of posterior odds requires that the prior distribution be consistent with this conjecture in the sense of avoiding zero probability for H_0 , implying in general that the prior is a mixture model. This implication is considered natural by Bayesians such as Jeffreys, with the argument that a prior distribution needs to integrate

probabilistic judgments that are inherent in the statement of competing hypotheses, which in this case attribute some importance to θ_0 , as opposed to other values of θ .

Other Bayesians such as Lindley and Zellner advocate a different approach, with a certain formal analogy with classical significance tests, in a way in which the statement of point hypotheses does not interfere with the prior distribution. Their approach can be described as a quantification of relative plausibility under the posterior for the value θ_0 , via the evaluation of $P = P(\theta \notin R_0(x) | x)$, where $R_0(x) = \{\theta \in \Theta : h(\theta | x) \geq h(\theta_0 | x)\}$ is the smallest HPD region that contains θ_0 . Large (small) values of the posterior relative plausibility P for H_0 are evidence in favor of (against) this hypothesis.

The fundamental tool of the Bayesian approach and the way the joint model $M = \{f(x | \theta)h(\theta), x \in \mathcal{X}, \theta \in \Theta\}$ is used in the implementation of inference already suggest that the question of evaluating the adequacy of a conjectured model in absolute terms might not have an answer in the sense of Popper (reject/not reject) of the type that is guaranteed by classical goodness-of-fit tests.

Bayes factors can be used if it is possible to extend the model M (or parts of it) to a larger family that contains the true model as an unknown quantity, and that allows comparing models within it. Otherwise, one can only define various measures of model adequacy for a relative analysis of a reference model in the context of a class of suitably defined competing models (see Chapter 5). The unsatisfactory nature of such options has led some statisticians to defend a Bayesian approach only when the underlying model is not put into question, a condition which Gillies (2001) refers to as *fixity of the theoretical framework*.

1.3.2 Predictive Inference

Many Bayesians believe that inference should not be restricted to statements about unobservable parameters. They note that parametric inference is awkward in that actual values are rarely known for parameters and therefore such inference can rarely be compared with reality. For Bayesians like Lindley, the most fundamental problem is to start with a set of observations (x_1, x_2, \dots, x_n) (yesterday) and infer conclusions, in terms of (subjective) probability, about a set of future observations $(x_{n+1}, x_{n+2}, \dots, x_{n+m})$ (tomorrow).

For easier exposition we assume $m = 1$ and that the $n + 1$ random variables $X_1, X_2, \dots, X_n, X_{n+1}$ are independent and identically distributed, given

θ with probability density function $f(x | \theta)$. The problem is to predict the random variable X_{n+1} after observing $(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)$. Trying to predict X_{n+1} with sampling model $f(x | \theta)$ we face two sources of randomness: (1) uncertainty that has to do with X_{n+1} being a random variable; (2) the impact of the uncertainty on θ . For example, if we estimate θ with the maximum likelihood estimator $\hat{\theta} = \hat{\theta}(x_1, x_2, \dots, x_n)$ and write $P(a < X_{n+1} < b | x_1, x_2, \dots, x_n) \cong \int_a^b f(x | \hat{\theta}) dx$, as estimate of the probability of the event $a < X_{n+1} < b$, then this expression ignores the randomness in the substitution of the parameter by its estimate. However, both types of randomness need to enter the prediction. The method of substituting an estimate for an unknown parameter in the sampling model (*plug-in procedure*) should thus be seen with some caveat.

Although the classical solution of the prediction problem involves much more than this (Amaral Turkman, 1980), it still could be said that the Bayesian solution is much cleaner. If one has only prior information, formalized as a prior distribution $h(\theta)$, the natural tool to use is the already discussed marginal or prior predictive distribution $f(x)$. The more interesting case is when one observes $x = (X_1 = x_1, X_2 = x_2, \dots, X_n = x_n)$ and wishes to predict X_{n+1} , assuming that conditional on θ the latter is independent of the previous observations [the problem of predicting $(X_{n+1}, X_{n+2}, \dots, X_{n+m})$ is not very different]. Using a completely probabilistic argument we have $f(x_{n+1} | x) = \int_{\theta} f(x_{n+1} | \theta) h(\theta | x) d\theta$, where the posterior distribution takes the place of the prior, as representing the information given the sample. One can then report summaries of this predictive distribution, including probabilities of any region in the sample space for X_{n+1} or values, $a = a(x)$ and $b = b(x)$ for any pre-specified probability $P(a < X_{n+1} < b | x) = \int_a^b f(x_{n+1} | x) dx_{n+1}$, which then determine a prediction interval (of HPD type if desired).

1.4 Conclusion

In summary, from a Bayesian point of view:

- The classical approach to statistical inference proceeds by arguments of an inductive type, such as, the notion of confidence intervals, which do not have a direct interpretation as probabilities. The difficulty or impossibility of making inference with a direct probabilistic interpretation – as the parameter θ itself is not even considered as a random variable – is strongly criticized by Jaynes (2003).
- Under a Bayesian approach, all inference can be derived from a logical

application of probability calculus. Bayesian statistical inference does not rely on any results that could not be derived from the rules of probability calculus, in particular Bayes' theorem. As O'Hagan (2010) puts it: "Probability theory is a completely self-consistent system. Any question of probabilities has one and only one answer, although there may be many ways to derive it."

Short of taking an extremist position, it is convenient to recall statements like that of Dawid (1985), who, besides confessing a clear preference for Bayesian theory comments that no statistical theory, Bayesian or not, can ever be entirely satisfactory. A position that some statisticians would argue for today is not the *exclusively Bayesian* option of authors like Savage, but rather an eclectic position shared by Wasserman (2004) when he argues that, in summary, combining prior judgments with data is naturally done by Bayesian methods, but to construct methods that guarantee good results in the long run under repeated observations, one needs to resort to frequentist methods.

Problems

- 1.1** Suppose there are N cable cars in San Francisco, numbered 1 to N . You don't know the value of N , so this is the unknown parameter. Your prior distribution on N is a geometric distribution with mean 100; that is

$$h(N) = \frac{1}{100} \left(\frac{99}{100} \right)^{N-1},$$

$N = 1, 2, \dots$ You see a cable car at random. It is numbered $x = 203$. Assume that x = number on a randomly picked car and has the probability distribution $f(x | N) = 1/N$ for $x = 1, \dots, N$, and $f(x | N) = 0$ for $x > N$.

- a.** Find the posterior distribution $h(N | x)$. Find the Bayes estimate of N , i.e., the posterior mean of N , and the posterior standard deviation of N (use results for a geometric series $\sum_{x=k}^{\infty} a^x$; or use a numerical approximation).
 - b.** Find a 95% HPD credible interval for N (you will not be able to exactly match the 95% – get as close as possible).
- 1.2** Recording the number of bacteria (y_i) found in $n = 6$ water samples (of the same volume), we find $y_i = 2, 3, 8, 6, 4$, and 1 (you might need $S = \sum y_i = 24$). It is thought reasonable to assume that y_i follows a Poisson distribution with mean θ , $i = 1, \dots, n$.
Also suppose that the prior $h(\theta) \propto 1/\sqrt{\theta}$ is used for the parameter θ (this is a so-called improper prior – see Chapter 2 for more discussion).

- a. Find the posterior distribution $h(\theta | y)$, and show how you will obtain a 95% credible interval for θ .
- b. Suppose that you are now told that only non-zero outcomes were recorded in the above experiment, and therefore the correct distribution for y_i , $i = 1, \dots, 6$, is the truncated Poisson given by

$$f(y | \theta) = \frac{e^{-\theta}\theta^y}{y!(1 - e^{-\theta})}, \quad y = 1, 2, \dots$$

- (i) Write down the likelihood function, and the posterior (using the same prior as before) up to a constant.
- (ii) Find a 95% credible interval using numerical integration (the previous posterior is no longer in a simple form that allows analytic evaluation of credible intervals).
- 1.3** Assume that the waiting time for a bus follows an exponential distribution with parameter θ . We have a single observation, $x = 3$. Assume that θ can only take one of 5 values, $\Theta = \{1, 2, 3, 4, 5\}$, with prior probabilities $h(\theta) \propto 1/\theta$, $\theta \in \Theta$.
- a. Find the posterior mode (MAP), the posterior standard deviation $SD(\theta | x)$, and the (frequentist) standard error of the estimator, i.e., $SD(\text{MAP} | \theta)$.
- b. Find a 60% HPD credible interval $A = (\theta_0, \theta_1)$, i.e., find the shortest interval A with $P(A | x) \geq 0.6$.
- c. Find $h(\theta | x, \theta \geq 2)$, i.e., find the posterior distribution for θ when we know that $\theta \geq 2$.
- d. Find the posterior mode MAP_0 conditional on $\theta \geq 2$, the conditional posterior standard deviation $SD(\theta | x, \theta \geq 2)$, and the (frequentist) standard error of the estimator, $SD(\text{MAP}_0 | \theta)$.
Compare with the answers under item (a) and comment.
- e. How would you justify the choice of the prior distribution $h(\theta) \propto 1/\theta$?
- 1.4** Your friend always uses a certain coin to bet “heads or tails”⁴ and you have doubts whether the coin is unbiased or not. Let θ denote the probability of a head. You want to test $H_1 : \theta < 0.5$ versus $H_2 : \theta = 0.5$ versus $H_3 : \theta > 0.5$. Assign a prior probability 1/2 that the coin is unbiased and equal probability to the other two hypotheses. That is, $p(H_2) = 0.5$ and $p(H_1) = p(H_3) = 0.25$. The prior distribution for θ under H_1 and H_3 is uniform. That is, $h(\theta | H_1) = U(0, 0.5)$ and $h(\theta | H_3) = U(0.5, 1)$. Assume you observe $n = 1$ coin toss. Let $x_1 \in \{0, 1\}$ denote an indicator of “head.”
- a. Find the predictive distribution (i.e., the marginal distribution) for the first toss, under each of the three hypotheses, i.e., find $p(x_1 = 1 | H_1)$, $p(x_1 = 1 | H_2)$, and $p(x_1 = 1 | H_3)$.
- b. Find the predictive distribution for the first toss, $p(x_1 = 1)$.

⁴ The labels “head” and “tail” refer to the two sides of a 25 cent coin.

- c. Still using the data x_1 from one coin toss, given $x_1 = 1$,
- (i) find the Bayes factors for H_1 versus H_2 and for H_3 versus H_2 ;
 - (ii) find the posterior probability that the coin is unbiased.
 - (iii) In general it is not meaningful to compute a Bayes factor with a non-informative uniform prior $h(\theta) \propto c$ (because the choice of c is arbitrary). Why is it okay here?