We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Why the Decision-Theoretic Perspective Misrepresents Frequentist Inference: Revisiting Stein's Paradox and Admissibility

Aris Spanos

Additional information is available at the end of the chapter

<http://dx.doi.org/10.5772/65720>

Abstract

The primary objective of this paper is to make a case that R.A. Fisher's objections to the decision-theoretic framing of frequentist inference are not without merit. It is argued that this framing is congruent with the Bayesian but incongruent with the frequentist approach; it provides the former with a theory of optimal inference but misrepresents the optimality theory of the latter. Decision-theoretic and Bayesian rules are considered optimal when they minimize the expected loss "for all possible values of θ in Θ " [$\forall \theta \in \Theta$], irrespective of what the true value θ^* [state of Nature] happens to be; the value that gave rise to the data. In contrast, the theory of optimal frequentist inference is framed entirely in terms of the capacity of the procedure to pinpoint θ^* . The inappropriateness of the quantifier $\forall \theta \in \Theta$ calls into question the relevance of admissibility as a minimal property for frequentist estimators. As a result, the pertinence of Stein's paradox, as it relates to the capacity of frequentist estimators to pinpoint θ^* , needs to be reassessed. The paper also contrasts loss-based errors with traditional frequentist errors, arguing that the former are attached to θ , but the latter to the inference procedure itself.

Keywords: decision theoretic inference, Bayesian vs. frequentist inference, Stein's paradox, James-Stein estimator, loss functions, admissibility, error probabilities, loss functions, risk functions, complete class theorem

1. Introduction

Wald's [1] decision-theoretic framework is widely viewed as providing a broad enough perspective to accommodate and compare the frequentist and Bayesian approaches to inference, despite their well-known differences. It is perceived as offering a neutral framing of inference that brings into focus their common features and tones down their differences; see Refs. [2–4].

Historically, Wald [5] proposed the original variant of the decision-theoretic framework with a view to unify Neyman's [6] rendering of frequentist interval estimation and testing:

"The problem in this formulation is very general. It contains the problems of testing hypotheses and of statistical estimation treated in the literature." (p. 340)

Among the frequentist pioneers, Jerzy Neyman accepted enthusiastically this broader perspective, primarily because the concepts of decision rules and action spaces seemed to provide a better framing for his *behavioristic interpretation* of Neyman-Pearson (N-P) testing based on the accept/reject rules; see Refs. [7, 8]. Neyman's attitude towards Wald's [1] framing was also adopted wholeheartedly by some of his most influential students/colleagues at Berkeley, including [9, 10]. In a foreword of a collection of Neyman's early papers, his students/editors described the Wald's framing as ([11], p. vii):

"A natural but far reaching extension of their [N-P formulation] scope can be found in Abraham Wald's theory of statistical decision functions."

At the other end of the argument, Fisher [12] rejected Wald's framing on the grounds that it seriously distorts his rendering of frequentist statistics:

"The attempt to reinterpret the common tests of significance used in scientific research as though they constituted some kind of acceptance procedure and led to "decisions" in Wald's sense, originated in several misapprehensions and has led, apparently, to several more." (p. 69)

With a few exceptions, such as Refs. [13–15], Fisher's [12] viewpoint has been inadequately discussed and evaluated by the subsequent statistics literature. The primary aim of this paper is to revisit Fisher's minority view by taking a closer look at the decision-theoretic framework with a view to reevaluate the claim that it provides a neutral framework for comparing the frequentist and Bayesian approaches. It is argued that Fisher's view that the decision theoretic framing is germane to "acceptance sampling," but misrepresents frequentist inference, is not without merit. The key argument of the discussion that follows is that the decision-theoretic notions of loss function and admissibility are congruent with the Bayesian approach, but incongruent with both the primary objective and the underlying reasoning of the frequentist approach.

Section 2 introduces the basic elements of the decision theoretic set-up with a view to bring out its links to the Bayesian and frequentist approaches, calling into question the conventional wisdom concerning its neutrality. Section 3 takes a closer look at the Bayesian approach and argues that had the decision-theoretic apparatus not exist, Bayesians would have been forced to invent it in order to establish a theory of optimal Bayesian inference. Section 4 discusses critically the notions of loss functions and admissibility, focusing primarily on their role in giving rise to Stein's paradox and their incompatibility with the frequentist approach. It is argued that the frequentist dimension of the notions of a loss function and admissibility is more apparent than real. Section 5 makes a case that the decision-theoretic framework misrepresents both the primary objective and the underlying reasoning of the frequentist approach. Section 6 revisits the notion of a loss function and its dependence on "information other than the data." It is argued that loss-based errors are both different and incompatible with the traditional frequentist errors because they are attached to the unknown parameters instead of the inference procedures themselves, as the traditional frequentist errors (Type I, II and coverage).

2. The decision theoretic set-up

2.1. Basic elements of the decision-theoretic framing

The current decision-theoretic set-up has three basic elements:

1. A prespecified (parametric) statistical model $\mathcal{M}_\theta(\mathbf{x})$, generically specified by

$$\mathcal{M}_\theta(\mathbf{x}) = \{f(\mathbf{x}; \theta), \theta \in \Theta\}, \mathbf{x} \in \mathbb{R}_X^n, \text{ for } \theta \in \Theta \subset \mathbb{R}^m, m \ll n, \quad (1)$$

where $f(\mathbf{x}; \theta)$ denotes the (joint) *distribution of the sample* $\mathbf{X} := (X_1, \dots, X_n)$, $\mathbb{R}_X^n$ denotes the *sample space* and Θ the *parameter space*. This model represents the stochastic mechanism assumed to have given rise to data $\mathbf{x}_0 := (x_1, \dots, x_n)$.

2. A decision space D containing all mappings $d(\cdot) : \mathbb{R}_X^n \rightarrow A$, where A denotes the set of all actions available to the statistician.
3. A loss function $L(\cdot, \cdot) : [D \times \Theta] \rightarrow R$, representing the numerical loss if the statistician takes action $a \in A$ when the state of Nature is $\theta \in \Theta$; see Refs. [2, 16–18].

The basic idea is that when the decision-maker selects action a , he/she does not know the “true” state of Nature, represented by θ^* . However, contingent on each action $a \in A$, the decision maker “knows” the losses (gains and utilities) resulting from different choices $(d, \theta) \in [D \times \Theta]$. The decision maker observes data $\mathbf{x}_0$, which provides some information about θ^* and then maps each $\mathbf{x} \in \mathbb{R}_X^n$ to a certain action $a \in A$ guided solely by $L(d, \theta)$.

2.2. The original Wald framing

It is important to bring out the fact that the original Wald [5] framing was much narrower than the above basic elements 2 and 3, due to its original objective to formalize the Neyman-Pearson (N-P) approach; see [19]. What were the key differences?

- i. The decision (action) space D was defined exclusively in terms of subsets of the parameter space Θ . For estimation purposes $D := \{\theta : \theta \in \Theta\}$ is the set of all singleton points of Θ and for testing $D := (\Theta_0, \Theta_1)$, the null and alternative regions, respectively.
- ii. The original loss (weight) was a zero-positive function, with zero loss at:

$$L_{0-c}(\theta, \hat{\theta}(\mathbf{X})) = \begin{cases} 0 & \text{if } \hat{\theta}(\mathbf{X}) = \theta^* \\ c_\theta > 0 & \text{if } \hat{\theta}(\mathbf{X}) = \theta \neq \theta^*, \theta \in \Theta, \end{cases} \quad (2)$$

where θ^* is the true value of θ in Θ . For the discussion that follows, it is important to note that Eq. (2) is nonoperational in practice because θ^* is unknown.

The more general framing, introduced by Wald ([1, 20]) and broadened by Le Cam [21], extended the scope of the original set-up by generalizing the notions of loss functions and decision spaces. In what follows it is argued that these extensions created serious incompatibilities with both the objective and the underlying reasoning of frequentist inference.

In addition, it is both of historical and methodological interest to note that Wald [5] introduced the notion of a prior distribution, $\pi(\theta)$, $\forall \theta \in \Theta$, into the original decision-theoretic machinery reluctantly, and justified it on being a useful tool for proving certain technical results:

“The situation regarding the introduction of an a priori probability distribution of θ is entirely different. First, the objection can be made against it, as Neyman has pointed out, that θ is merely an unknown constant and not a variate, hence it makes no sense to speak of the probability distribution of θ . Second, even if we may assume that θ is a variate, we have in general no possibility of determining the distribution of θ and any assumptions regarding this distribution are of hypothetical character. The reason why we introduce here a hypothetical probability distribution of θ is simply that it proves to be useful in deducing certain theorems and in the calculation of the best system of regions of acceptance.” (p. 302)

2.3. A shared neutral framework?

The frequentist, Bayesian, and the decision-theoretic approaches share the notion of a statistical model by viewing data $\mathbf{x}_0 := (x_1, \dots, x_n)$ as a realization of a sample $\mathbf{X} := (X_1, \dots, X_n)$ from Eq. (1).

The key differences between the three approaches are as follows:

- a. The frequentist approach relies exclusively on $\mathcal{M}_\theta(\mathbf{x})$
- b. The Bayesian approach adds a *prior* distribution, $\pi(\theta)$, $\forall \theta \in \Theta$ (for all $\theta \in \Theta$)
- c. The decision-theoretic framing revolves around a *loss* (gain or utility) function:

$$L(d(\mathbf{x}), \theta), \forall \theta \in \Theta, \forall \mathbf{x} \in \mathbb{R}_X^n. \quad (3)$$

The loss function is often assumed to be an even, differentiable and convex function of $(d(\mathbf{x}) - \theta)$ and can take numerous functional forms; see Refs. [17, 18] inter alia.

The claim that the decision-theoretic perspective provides a neutral ground is often justified [3] on account of the loss function being a function of the sample and parameter spaces through the two *universal quantifiers*:

- (i) “ $\forall \mathbf{x} \in \mathbb{R}_X^n$,” associated with the *distribution of the sample*:

$$\text{frequentist : } f(\mathbf{x}; \theta), \forall \mathbf{x} \in \mathbb{R}_X^n, \quad (4)$$

- (ii) “ $\forall \theta \in \Theta$ ” associated with the *posterior distribution*:

$$\text{Bayesian : } \pi(\theta|\mathbf{x}_0) = \frac{\pi(\theta) \cdot f(\mathbf{x}_0|\theta)}{\int_{\theta \in \Theta} \pi(\theta) \cdot f(\mathbf{x}_0|\theta) d\theta}, \forall \theta \in \Theta. \quad (5)$$

The idea is that allowing for all values of $\mathbf{x}$ in $\mathbb{R}_X^n$ goes beyond the Bayesian perspective, which relies exclusively on a single point $\mathbf{x}_0$. What is not obvious is whether that is sufficient to do justice to the frequentist approach. A closer scrutiny suggests that frequentist inference is misrepresented by the way both quantifiers are employed in the decision-theoretic framing of inference.

First, the quantifier $\forall \mathbf{x} \in \mathbb{R}_X^n$ plays only a minor role in transforming a loss function, say $L(\theta, \hat{\theta}(\mathbf{x}))$, into a risk function:

$$R(\theta, \hat{\theta}) = E_X[L(\theta, \hat{\theta}(\mathbf{X}))] = \int_{\mathbf{x} \in \mathbb{R}_X^n} L(\theta, \hat{\theta}(\mathbf{x})) f(\mathbf{x}; \theta) d\mathbf{x}, \forall \theta \in \Theta. \quad (6)$$

This is the only place where the distribution of the sample, $f(\mathbf{x}; \theta)$, $\forall \mathbf{x} \in \mathbb{R}_X^n$ enters the decision-theoretic framing, and the only relevant part of the behavior of $\hat{\theta}(\mathbf{X})$ is how it affects the risk function for different values of θ in Θ . In frequentist inference, however, the distribution of the sample takes center stage for the theory of optimal frequentist inference. It determines the sampling distribution of any statistic $Y_n = g(\mathbf{X})$ (estimator, test, and predictor) through:

$$F(y; \theta) := \mathbb{P}(Y_n \leq y; \theta) = \underbrace{\int \cdots \int}_{\{\mathbf{x}: g(\mathbf{x}) \leq y; \mathbf{x} \in \mathbb{R}_X^n\}} f(\mathbf{x}; \theta) d\mathbf{x}, \quad (7)$$

and that, in turn, yields the relevant error probabilities that determine optimal inference procedures.

Second, the decision-theoretic notion of optimality revolves around the universal quantifier " $\forall \theta \in \Theta$," rendering it congruent with the Bayesian but incongruent with the frequentist approach. To be more specific, since different risk functions often intersect over Θ , an optimal rule is usually selected after the risk function is reduced to a scalar. Two such choices of risk are:

$$\begin{aligned} \textbf{Maximum risk : } R_{\max}(\hat{\theta}) &= \sup_{\theta \in \Theta} R(\theta, \hat{\theta}), \\ \textbf{Bayes risk : } R_B(\hat{\theta}) &= \int_{\theta \in \Theta} R(\theta, \hat{\theta}) \pi(\theta) d\theta. \end{aligned} \quad (8)$$

Hence, an obvious way to choose among different rules is to find the one that minimizes the relevant risk with respect to all possible estimates $\tilde{\theta}(\mathbf{x})$. In the case of Eq. (8), this gives rise to two corresponding decision rules:

$$\begin{aligned} \textbf{Minimax rule : } \inf_{\tilde{\theta}(\mathbf{x})} R_{\max}(\hat{\theta}) &= \inf_{\tilde{\theta}(\mathbf{x})} [\sup_{\theta \in \Theta} R(\theta, \hat{\theta})], \\ \textbf{Bayes rule : } \inf_{\tilde{\theta}(\mathbf{x})} R_B(\hat{\theta}) &= \inf_{\tilde{\theta}(\mathbf{x})} \int_{\theta \in \Theta} R(\theta, \hat{\theta}) \pi(\theta) d\theta. \end{aligned} \quad (9)$$

In this sense, a decision or a Bayes rule $\tilde{\theta}(\mathbf{x})$ will be considered optimal when it minimizes the relevant risk, no matter what the true state of Nature θ^* happens to be. The last clause, "irrespective of θ^* " constitutes a crucial caveat that is often ignored in discussions of these approaches. When viewed as a game against Nature, the decision maker selects action a from A , irrespective of what value θ^* Nature has chosen. That is, θ^* plays no role in selecting the optimal rules since the latter have nothing to do with the true value θ^* of θ . To avoid any misreading of this line of reasoning, it is important to emphasize that "the true value θ^* " is shorthand for saying that "data $\mathbf{x}_0$ constitute a typical realization of the sample $\mathbf{X}$ with distribution $f(\mathbf{x}; \theta^*)$ "; see Ref. [22].

This should be contrasted with the notion of optimality in frequentist inference that gives θ^* center stage, in the sense that it evaluates the capacity of the inference procedure to inform the modeler about θ^* ; no other value is relevant. According to Reid [23]:

“A statistical model is a family of probability distributions $[\mathcal{M}_\theta(\mathbf{x})]$, the central problem of statistical inference being to identify which member of the family $[\theta^*]$ generated the data of interest.” (p. 418)

3. The Bayesian approach

To shed further light on the affinity between the decision-theoretic framework and the Bayesian approach, let us take a closer look at the latter.

3.1. Bayesian inference and its primary objective

A key argument in favor of the Bayesian approach is often its simplicity in the sense that all forms of inference revolve around a single function, the posterior distribution: $\pi(\theta|\mathbf{x}_0) \propto \pi(\theta) \cdot f(\mathbf{x}_0|\theta)$, $\forall \theta \in \Theta$. Hence, an outsider looking at Bayesian approach might naturally surmise that its primary objective is to yield “a probabilistic ranking” (ordering) of all values of θ in Θ . According to O’Hagan [4]:

“Having obtained the posterior density $\pi(\theta|\mathbf{x}_0)$, the final step of the Bayesian method is to derive from it suitable inference statements. The most usual inference question is this: After seeing the data $\mathbf{x}_0$, what do we now know about the parameter θ . The only answer to this question is to present the entire posterior distribution.” (p. 6)

The idea is that the modeling begins with an a priori probabilistic ranking based on $\pi(\theta)$, $\forall \theta \in \Theta$, which is revised after observing $\mathbf{x}_0$ to derive $\pi(\theta|\mathbf{x}_0)$, $\forall \theta \in \Theta$; hence the key role of the quantifier $\forall \theta \in \Theta$. O’Hagan [4], echoing earlier views in [24, 25], contrast the frequentist (classical) inferences with the Bayesian inference arguing:

“Classical inference theory is very concerned with constructing *good* inference rules. The primary concern of Bayesian inference, ..., is entirely different. The objective is to extract information concerning θ from the posterior distribution, and to present it helpfully via effective summaries. There are two criteria in this process. The first is to identify interesting features of the posterior distribution. ... The second criterion is good communication. Summaries should be chosen to convey clearly and succinctly all the features of interest. ... In Bayesian terms, therefore, a good inference is one which contributes effectively to appropriating the information about θ which is conveyed by the posterior distribution.” (p. 14)

Clearly, O’Hagan’s [4] attempt to define what is a “good” Bayesian inference begs the question: what does constitute “effective appropriation of information about θ ” mean, beyond the probabilistic ranking? That is, the issue of optimality is inextricably bound up with what the primary objective of Bayesian inference is. If the primary objective of Bayesian inference is not the revised probabilistic ranking, what is it? The answer is that the ranking is only half the story. The other half is concerned with the optimality for Bayesian inference which cannot be

framed exclusively in terms of the posterior distribution. The decision-theoretic perspective provides the Bayesian approach with a theory of optimal inference as well as a primary objective: minimize expected losses for all values of θ in Θ .

In his attempt to defend his stance that the entire posterior distribution *is* the inference, O'Hagan [4] argues that criteria for "optimal" Bayesian inferences are only *parasitical* on the Bayesian approach and enter the picture through the decision theoretic perspective:

"... a study of decision theory has two potential benefits. First, it provides a link to classical inference. It thereby shows to what extent classical estimators, confidence intervals and hypotheses tests can be given a Bayesian interpretation or motivation. Second, it helps identify suitable summaries to give Bayesian answers to stylized inference questions which classical theory addresses." (p. 14)

Both of the above mentioned potential benefits to the Bayesian approach, are questionable for two reasons. First, the link between the decision-theoretic and the classical (frequentist) inference is more apparent than real because it is fraught with misleading definitions and unclari- ties pertaining to the reasoning and objectives of the latter. As argued in the sequel, the quantifier " $\forall \theta \in \Theta$ " used to define "optimal" decision-theoretic or Bayes rules is at odds with and misrepresents frequentist inference. Second, the claim concerning Bayesian answers to frequentist questions of interest is misplaced because the former provides no real answers to the frequentist primary question of interest which pertains to learning about θ^* . An optimal Bayes rule offers very little, if anything, relevant for learning about the value θ^* that gave rise to $\mathbf{x}_0$. Let us unpack this answer in some more detail.

3.2. Optimality for Bayesian inference

What does minimizing the Bayes risk amount to? Substituting the risk function in Eq. (6) into the Bayes risk in Eq. (8), one can show that:

$$\begin{aligned} R_B(\hat{\theta}) &= \int_{\theta \in \Theta} \left(\int_{\mathbf{x} \in \mathbb{R}_X^n} L(\theta, \hat{\theta}(\mathbf{x})) f(\mathbf{x}; \theta) d\mathbf{x} \right) \pi(\theta) d\theta \\ &= \int_{\mathbf{x} \in \mathbb{R}_X^n} \int_{\theta \in \Theta} L(\theta, \hat{\theta}(\mathbf{x})) f(\mathbf{x}; \theta) \pi(\theta) d\theta d\mathbf{x} \\ &= \int_{\mathbf{x} \in \mathbb{R}_X^n} \left\{ \int_{\theta \in \Theta} L(\theta, \hat{\theta}(\mathbf{x})) \pi(\theta | \mathbf{x}) d\theta \right\} m(\mathbf{x}) d\mathbf{x}, \end{aligned} \quad (10)$$

where $m(\mathbf{x}) = \int_{\theta \in \Theta} f(\mathbf{x}; \theta) d\theta$; see Ref. [18]. The second and third equalities presume that one can reverse the order of integration (a technical issue), and treat $f(\mathbf{x}; \theta)$ as the joint distribution of $\mathbf{X}$ and θ so that the following equalities hold:

$$f(\mathbf{x}; \theta) = f(\mathbf{x} | \theta) \pi(\theta) = \pi(\theta | \mathbf{x}) m(\mathbf{x}). \quad (11)$$

In this case, these equalities are questionable due to the blurring of the distinction between $\mathbf{x}$, a generic value of $\mathbb{R}_X^n$, and the particular value $\mathbf{x}_0$; see Ref. [26].

In light of Eq. (10), a Bayesian estimate is “optimal” relative to a particular loss function $L(\hat{\theta}(\mathbf{X}), \theta)$, when it minimizes $R_B(\hat{\theta})$, or equivalently $\int_{\theta \in \Theta} L(\theta, \hat{\theta}(\mathbf{x}))\pi(\theta|\mathbf{x})d\theta$. This makes it clear that what constitutes an “optimal” Bayesian estimate is primarily determined by $L(\hat{\theta}(\mathbf{X}), \theta)$ [27]:

- i. When $L_2(\hat{\theta}, \theta) = (\hat{\theta} - \theta)^2$, the Bayes estimate $\hat{\theta}$ is the *mean* of $\pi(\theta|\mathbf{x}_0)$.
- ii. When $L_1(\tilde{\theta}, \theta) = |\tilde{\theta} - \theta|$, the Bayes estimate $\tilde{\theta}$ is the *median* of $\pi(\theta|\mathbf{x}_0)$.
- iii. When $L_{0-1}(\bar{\theta}, \theta) = \delta(\bar{\theta}, \theta) = \begin{cases} 0 & \text{for } |\bar{\theta} - \theta| < \varepsilon \\ 1 & \text{for } |\bar{\theta} - \theta| \geq \varepsilon \end{cases}$, for $\varepsilon > 0$, the Bayes estimate $\bar{\theta}$ is the *mode* of $\pi(\theta|\mathbf{x}_0)$.

In practice, the most widely used loss function is the square:

$$L_2(\hat{\theta}(\mathbf{X}); \theta) = (\hat{\theta}(\mathbf{X}) - \theta)^2, \forall \theta \in \Theta, \quad (12)$$

whose risk function is the decision-theoretic *Mean Square Error* (MSE_1):

$$R(\theta, \hat{\theta}) = E(\hat{\theta}(\mathbf{X}) - \theta)^2 = MSE_1(\hat{\theta}(\mathbf{X}); \theta), \forall \theta \in \Theta. \quad (13)$$

Surprising, however, this definition of the MSE, denoted by MSE_1 , is different from the frequentist MSE, which is defined by:

$$MSE(\hat{\theta}_n(\mathbf{X}); \theta^*) = E(\hat{\theta}_n(\mathbf{X}) - \theta^*)^2. \quad (14)$$

The key difference is that Eq. (14) is defined at the point $\theta = \theta^*$, as opposed to $\forall \theta \in \Theta$. Unfortunately, statistics textbooks adopt one of the two definitions of the MSE— either at $\theta = \theta^*$ or $\forall \theta \in \Theta$ —and ignore (or seem unaware) of the other. At first sight, this difference might appear pedantic, but it turns out that it has very serious implications for the relevant theory of optimality for the frequentist vs. Bayesian inference procedures. Indeed, reliance on $\forall \theta \in \Theta$ undermines completely the relevance of admissibility as a minimal property for estimators in frequentist inference.

Admissibility. An estimator $\tilde{\theta}(\mathbf{X})$ is *inadmissible* if there exists another estimator $\hat{\theta}(\mathbf{X})$ such that:

$$R(\theta, \hat{\theta}) \leq R(\theta, \tilde{\theta}), \forall \theta \in \Theta, \quad (15)$$

and the strict inequality ($<$) holds for at least one value of θ . Otherwise, $\tilde{\theta}(\mathbf{X})$ is said to be *admissible* with respect to the loss function $L(\theta, \hat{\theta})$.

The objective of minimizing losses weighted by $\pi(\theta|\mathbf{x}_0)$ for all value of θ in Θ , is in direct contrast to the frequentist primary objective, which is to learn from data about the true value θ^* underlying the generation of $\mathbf{x}_0$. Hence, the question that naturally arises is: what does an optimal Bayes rule, stemming from Eq. (17) convey about the underlying data generating mechanism in Eq. (1)? It is not obvious why the highest ranked value $\tilde{\theta}(\mathbf{x}_0)$ (mode), or some

other feature of the posterior distribution, has any value in pinpointing θ^* knowing that $\tilde{\theta}(\mathbf{x}_0)$ is selected irrespective of θ^* the true state of Nature.

3.3. The duality between loss functions and priors

The derivation in Eq. (10) brings out the built-in affinity between the decision-theoretic framing of inference and the Bayesian approach. As shown above, minimizing the Bayes risk:

$$R_B(\hat{\theta}) = \int_{\theta \in \Theta} R(\hat{\theta}, \theta) \pi(\theta) d\theta, \quad (16)$$

is equivalent to minimizing the integral:

$$\int_{\theta \in \Theta} L(\hat{\theta}(\mathbf{X}), \theta) \pi(\theta | \mathbf{x}) d\theta. \quad (17)$$

This result brings out two important features of optimal Bayesian inference.

First, it confirms the minor role played by the quantifier $\mathbf{x} \in \mathbb{R}_X^n$ in both the Bayesian and decision-theoretic optimality theory of inference.

Second, it indicates that $L(\theta, \hat{\theta})$ and $\pi(\theta)$ are perfect substitutes with respect to any weight function $w(\theta) > 0, \forall \theta \in \Theta$, in the derivation of Bayes rules. Modifying the loss function or the prior yields the same result:

“... the problem of estimating θ with a modified (weighted) loss function is identical to the problem with a simple loss but with modified hyperparameters of the prior distribution while the form of the prior distribution does not change.” ([28], p. 522)

This implies that in practice a Bayesian could derive a particular Bayes rule by attaching the weight to the loss function or to the prior distribution depending on which derivation is easier; see Refs. [18, 28].

3.4. Revisiting the complete class theorem

The issue of contrasting objectives highlights the key built-in tension between the frequentist and Bayesian approaches to optimality, which in turn undermines several important results, including the complete class theorem, first proved in Ref. [20]:

“Wald showed that under fairly general conditions the class of Bayes decision functions forms an essentially complete class; in other words, for any decision function that is not Bayesian, there exists one that is Bayes and is at least as good no matter what the true state of Nature may be.” ([19], p. 341)

As argued in the sequel, it should come as no surprise to learn that Bayes rules dominate all other rules when admissibility is given center stage. The key result is that a Bayes rule $\hat{\theta}_B(\mathbf{x})$ with respect to a prior distribution $\pi(\theta)$ is:

- i. *Admissible*, under certain regularity conditions, including when $\hat{\theta}_B(\mathbf{x})$ is unique up to equivalence relative to the same risk function $R(\theta, \hat{\theta}_B)$.
- ii. *Minimax* when $R(\theta, \hat{\theta}_B) = c < \infty$.
- iii. An admissible, relative to a risk function $R(\theta, \hat{\theta}_B)$, estimate $\hat{\theta}(\mathbf{x})$ is either Bayes $\hat{\theta}_B(\mathbf{x})$ or the limit of a sequence of Bayes rules; see Refs. [2, 17, 28].

Ignoring the contrasting objectives, these results have been interpreted as evidence for the superiority of the Bayesian perspective, and led to the intimation that an effective way to generate optimal frequentist procedures is to find the Bayes solution using a reasonable prior and then examine their frequentist properties to see whether it is satisfactory from the latter viewpoint; see Refs. [29, 30].

As argued next, even if one were to agree that Bayes rules and admissible estimators largely coincide, the importance of such a result hinges on the relevance of admissibility as a key property for frequentist estimators.

4. Loss functions and admissibility revisited

The claim to be discussed in this section is that the notions of a “loss function” and “admissibility” are incompatible with the optimal theory of frequentist estimation as framed by Fisher; see Ref. [31].

4.1. Admissibility as a minimal property

The following example brings out the inappropriateness of admissibility as a minimal property for optimal frequentist estimators.

Example. In the context of the simple Normal model:

$$X_k \sim \text{NIID}(\theta, 1), k=1, 2, \dots, n, \text{ for } n > 2, \quad (18)$$

consider the decision-theoretic notion of MSE_1 in Eq. (13) to compare two estimators of θ :

- i. The maximum likelihood estimator (MLE): $\bar{X}_n = \frac{1}{n} \sum_{k=1}^n X_k$
- ii. The “crystalball” estimator: $\theta_{cb} = 7405926, \forall \mathbf{x} \in \mathbb{R}_X^n$

When compared on admissibility grounds, both estimators are admissible and thus equally acceptable. Common sense, however, suggests that if a particular criterion of optimality cannot distinguish between $\bar{X}_n$ [a strongly consistent, unbiased, fully efficient and sufficient estimator] and θ_{cb} , an arbitrarily chosen real number that ignores the data altogether, is not much of a minimal property.

A moment’s reflection suggests that the inappropriateness of admissibility stems from its reliance on the quantifier “ $\forall \theta \in \Theta$.” The admissibility of θ_{cb} arises from the fact that for certain

values of θ close enough to θ_{cb} , say $\theta \in (\theta_{cb} \pm \frac{\lambda}{\sqrt{n}})$, for $0 < \lambda < 1$, θ_{cb} is “better” than $\bar{X}_n$ on MSE_1 grounds:

$$MSE_1(\bar{X}_n; \theta) = \frac{1}{n} > MSE_1(\theta_{cb}; \theta) \leq \frac{\lambda^2}{n} \text{ for } \theta \in \left(\theta_{cb} \pm \frac{\lambda}{\sqrt{n}}\right). \quad (19)$$

Given that the primary objective of a frequentist estimator is to pin-point θ^* , the result in Eq. (19) seems totally irrelevant as a gauge of its capacity to achieve that!

This example indicates that admissibility is totally ineffective as a *minimal* property because it does not filter out θ_{cb} , the worst possible estimator! Instead, it excludes potentially good estimators like the *sample median*; see Ref. [32]. This highlights the “extreme relativism” of admissibility to the particular loss function, $L_2(\hat{\theta}(\mathbf{X}); \theta)$, in this case. For the absolute loss function $L_1(\hat{\theta}(\mathbf{X}); \theta) = |\hat{\theta}(\mathbf{X}) - \theta|$, however, the sample median would have been the optimal estimator. Despite his wholehearted embrace of the decision-theoretic framing, Lehmann [33] warned statisticians about the perils of arbitrary loss functions:

“It is argued that the choice of a loss function, while less crucial than that of the model, exerts an important influence on the nature of the solution of a statistical decision problem, and that an arbitrary choice such as squared error may be baldly misleading as to the relative desirability of the competing procedures.” (p. 425)

A strong case can be made that *the* key minimal property (necessary but not sufficient) for frequentist estimation is *consistency*, an extension of the Law of Large Numbers (LLN) to estimators, more generally. For instance, consistency would have eliminated θ_{cb} from consideration because it is inconsistent. This makes intuitive sense because if an estimator $\hat{\theta}(\mathbf{X})$ cannot pinpoint θ^* with an infinite data information, it should be considered irrelevant for learning about θ^* . Indeed, there is nothing in the notion of admissibility that advances learning from data about θ^* .

Further to relative (to particular loss functions) efficiency being a dubious property for frequentist estimators, the pertinent measure of finite sample precision for frequentist estimators is full efficiency, which is defined relative to the assumed statistical model (1).

4.2. Stein's paradox and admissibility

The quintessential example that has bolstered the appeal of the Bayesian claims concerning admissibility is the James-Stein estimator [34], which gave rise to an extensive literature on *shrinkage estimators*, see Ref. [35].

Let $\mathbf{X} := (X_1, X_2, \dots, X_m)$ be independent sample from a Normal distribution:

$$X_k \sim \text{NI}(\theta_k, \sigma^2), k=1, 2, \dots, m, \quad (20)$$

where σ^2 is known. Using the notation $\boldsymbol{\theta} := (\theta_1, \theta_2, \dots, \theta_m)$ and $\mathbf{I}_m := \text{diag}(1, 1, \dots, 1)$, this can be denoted by:

$$\mathbf{X} \sim N(\theta, \sigma^2 \mathbf{I}_m).$$

Find an optimal estimator $\tilde{\theta}(\mathbf{X})$ of θ with respect to the square “overall” loss function:

$$L_2(\theta, \hat{\theta}(\mathbf{X})) = (\|\hat{\theta}(\mathbf{X}) - \theta\|^2) = \sum_{k=1}^m (\hat{\theta}_k(\mathbf{X}) - \theta_k)^2. \quad (21)$$

Stein [36] astounded the statistical world by showing that for $m=2$ the least-squares (LS) estimator $\hat{\theta}_{LS}(\mathbf{X}) = \mathbf{X}$ is admissible, but for $m > 2$ $\hat{\theta}_{LS}(\mathbf{X})$ is inadmissible. Indeed, James and Stein [37] were able to come up with a nonlinear estimator:

$$\hat{\theta}_{JS}(\mathbf{X}) = \left(1 - \frac{(m-2)\sigma^2}{\|\mathbf{X}\|^2}\right) \mathbf{X}, \quad (22)$$

that became known as the James-Stein estimator, which dominates $\hat{\theta}_{LS}(\mathbf{X}) = \mathbf{X}$ in MSE_1 terms by demonstrating that:

$$\text{MSE}_1(\hat{\theta}_{JS}(\mathbf{X}); \theta) < \text{MSE}_1(\hat{\theta}_{LS}(\mathbf{X}); \theta), \forall \theta \in \mathbb{R}^m. \quad (23)$$

It turns out that $\hat{\theta}_{JM}(\mathbf{X})$ is also inadmissible for $m > 2$ and dominated by the modified James-Stein estimator that is *admissible*:

$$\hat{\theta}_{JS}^+(\mathbf{X}) = \left(1 - \frac{(m-2)\sigma^2}{\|\mathbf{X}\|^2}\right)^+ \mathbf{X}, \quad (24)$$

where $(z)^+ = \max(0, z)$; see Ref. [17].

The traditional interpretation of this result is that for the Normal, Independent model in Eq. (20), the James-Stein estimator (15) of $\theta := (\theta_1, \theta_2, \dots, \theta_m)$, for $m > 2$, reduces the *overall* MSE_1 in Eq. (21). This result seems to imply that one will “do better” (in overall MSE_1 terms) by using a combined nonlinear (shrinkage) estimator, instead of estimating these means separately. What is surprising about this result is that there is no statistical reason (due to independence) to connect the inferences pertaining to the different individual means, and yet the obvious estimator (LS) is inadmissible.

As argued next, this result calls into question the appropriateness of the notion of admissibility with respect to a particular loss function, and not the judiciousness of frequentist estimation.

5. Frequentist inference and learning from data

The objectives and underlying reasoning of frequentist inference are inadequately discussed in the statistics literature. As a result, some of its key differences with Bayesian inference remain beclouded.

5.1. Frequentist approach: primary objective and reasoning

All forms of parametric frequentist inference begin with a prespecified statistical model $\mathcal{M}_\theta(\mathbf{x}) = \{f(\mathbf{x}; \theta), \theta \in \Theta\}$, $\mathbf{x} \in \mathbb{R}_X^n$. This model is chosen from the set of all possible models that could have given rise to data $\mathbf{x}_0 := (x_1, \dots, x_n)$, by selecting the probabilistic structure for the underlying stochastic process $\{X_t, t \in \mathbb{N} := (1, 2, \dots, n, \dots)\}$ in such a way so as to render the observed data $\mathbf{x}_0$ a “typical” realization thereof. In light of the fact that each value of $\theta \in \Theta$ represents a different element of the family of models represented by $\mathcal{M}_\theta(\mathbf{x})$, the primary objective of frequentist inference is to learn from data about the “true” model:

$$\mathcal{M}^*(\mathbf{x}) = \{f(\mathbf{x}; \theta^*)\}, \mathbf{x} \in \mathbb{R}_X^n, \quad (25)$$

where θ^* denotes the true value of θ in Θ . The “typicality” is testable vis-a-vis the data $\mathbf{x}_0$ using misspecification testing; see Ref. [38].

The frequentist approach relies on two modes of reasoning for inference purposes:

$$\begin{aligned} \text{Factual (estimation, prediction)} : & \quad f(\mathbf{x}; \theta^*), \forall \mathbf{x} \in \mathbb{R}_X^n, \\ \text{Hypothetical (hypothesis testing)} : & \quad f(\mathbf{x}; \theta_0), f(\mathbf{x}; \theta_1), \forall \mathbf{x} \in \mathbb{R}_X^n, \end{aligned} \quad (26)$$

where θ^* denotes the true value of θ in Θ , and $\theta_i, i = 0, 1$ denote hypothesized values of θ associated with the hypotheses, $H_0: \theta_0 \in \Theta_0, H_1: \theta_1 \in \Theta_1$, where Θ_0 and Θ_1 constitute a partition of Θ .

A frequentist estimator $\hat{\theta}$ aims to pinpoint θ^* , and its optimality is evaluated by how effectively it achieves that. Similarly, a test statistic usually compares a good estimator $\hat{\theta}$ of θ with a prespecified value θ_0 , but behind $\hat{\theta}$ is the value θ^* assumed to have generated data $\mathbf{x}_0$. Hence, the hypothetical reasoning is used in testing to learn about θ^* , and has nothing to do with all possible values of θ in Θ .

This contradicts misleading claims by Bayesian textbooks ([3], p. 61):

“The frequentist paradigm relies on this criterion [risk function] to compare estimators and, if possible, to select the best estimator, the reasoning being that estimators are evaluated on their long-run performance for all possible values of the parameter θ .”

Contrary to this claim, the only relevant value of θ in evaluating the “optimality” of $\hat{\theta}$ is θ^* . Such misleading claims stem from an apparent confusion between the existential and universal quantifiers in framing certain inferential assertions.

The existence of θ^* can be formally defined using the *existential* quantifier:

$$\exists \theta^* \in \Theta : \quad \text{there exists a } \theta^* \in \Theta \text{ such that.} \quad (27)$$

This introduces a potential conflict between the *existential* and the *universal* quantifier “ $\forall \theta \in \Theta$ ” because neither the decision theoretic nor the Bayesian approach explicitly invoke θ^* . Decision-theoretic and Bayesian rules are considered optimal when they minimize the expected loss $\forall \theta \in \Theta$, no matter what θ^* happens to be.

Any attempt to explain away the crucial differences between the two quantifiers can be easily scotched using elementary logic. The two quantifiers could not be more different since, using the logical connective for negation ($\neg$), the equivalence between the two involves double negations:

$$(i) \exists \theta^* \in \Theta \Leftrightarrow \neg \forall \theta \notin \Theta, \quad (ii) \forall \theta \in \Theta \Leftrightarrow \neg \exists \theta^* \notin \Theta. \quad (28)$$

Similarly, invoking intuition to justify the quantifier $\forall \theta \in \Theta$ as innocuous and natural on the grounds that one should care about the behavior of an estimator $\hat{\theta}$ for all possible values of θ , is highly misleading. The behavior of $\hat{\theta}$, for all $\theta \in \Theta$, although relevant, is not what determines how effective a frequentist estimator is at pinpointing θ^* ; what matters is its sampling behavior around θ^* . Assessing its effectiveness calls for evaluating (deductively) the sampling distribution of $\hat{\theta}$ under factual $\theta = \theta^*$, or hypothetical values θ_0 and θ_1 , and not for all possible values of θ in Θ . Let's unpack the details of this claim.

5.2. Frequentist estimation

The underlying reasoning for frequentist estimation is *factual*, in the sense the optimality of an estimator is appraised in terms of its generic capacity of $\hat{\theta}_n(\mathbf{X})$ to *zero-in* on (pinpoint) the true value θ^* , whatever the sample realization $\mathbf{X} = \mathbf{x}_0$. Optimal properties like consistency, unbiasedness, full efficiency, sufficiency, etc., calibrate this generic capacity using its sampling distribution of $\hat{\theta}_n(\mathbf{X})$ evaluated under $\theta = \theta^*$ i.e., in terms of $f(\hat{\theta}_n(\mathbf{x}); \theta^*)$, for $\mathbf{x} \in \mathbb{R}_X^n$. For instance, *strong consistency* asserts that as $n \rightarrow \infty$, $\hat{\theta}_n(\mathbf{X})$ will zero-in on θ^* almost surely:

$$\mathbb{P}(\lim_{n \rightarrow \infty} \hat{\theta}_n(\mathbf{X}) = \theta^*) = 1. \quad (29)$$

Similarly, *unbiasedness* asserts that the mean of $\hat{\theta}_n(\mathbf{X})$ is the true value θ^* :

$$E(\hat{\theta}_n(\mathbf{X})) = \theta^*. \quad (30)$$

In this sense, both of these optimal properties are defined at the point $\theta = \theta^*$. This is achieved by using *factual reasoning*, i.e., evaluating the sampling distribution of $\hat{\theta}_n(\mathbf{X})$ under the true state of Nature ($\theta = \theta^*$), without having to know θ^* . This is in contrast to using loss functions, such as Eq. (2), which are defined in terms of θ^* but are rendered nonoperational without knowing θ^* .

Example. In the case of the simple Normal model in Eq. (18), the point estimator, $\bar{X}_n$, is consistent, unbiased, fully efficient, sufficient, with a sampling distribution:

$$\bar{X}_n \sim N\left(\theta, \frac{1}{n}\right). \quad (31)$$

What is not usually appreciated sufficiently is that the evaluation of that distribution is *factual*, i.e., $\theta = \theta^*$, and should formally be denoted by:

$$\bar{X}_n \stackrel{\theta=\theta^*}{\sim} N\left(\theta^*, \frac{1}{n}\right). \quad (32)$$

When $\bar{X}_n$ is standardized, it yields the pivotal function:

$$d(\mathbf{X}; \theta) := \sqrt{n}(\bar{X}_n - \theta^*) \stackrel{\theta=\theta^*}{\sim} N(0, 1), \quad (33)$$

whose distribution only holds for the true θ^* , and no other value. This provides the basis for constructing a $(1 - \alpha)$ confidence interval (CI):

$$\mathbb{P}\left(\bar{X}_n - c_{\frac{\alpha}{2}}\left(\frac{1}{\sqrt{n}}\right) \leq \theta \leq \bar{X}_n + c_{\frac{\alpha}{2}}\left(\frac{1}{\sqrt{n}}\right); \theta=\theta^*\right) = 1 - \alpha, \quad (34)$$

which asserts that the random interval $\left[\bar{X}_n - c_{\frac{\alpha}{2}}\left(\frac{1}{\sqrt{n}}\right), \bar{X}_n + c_{\frac{\alpha}{2}}\left(\frac{1}{\sqrt{n}}\right)\right]$ will cover (overlay) the true mean θ^* , *whatever that happens to be*, with probability $(1 - \alpha)$, or equivalently, the error of coverage is α . Hence, frequentist evaluation of the coverage error probability depends only on the sampling distribution of $\bar{X}_n$ and is attached to random interval for all values $\theta \neq \theta^*$ without requiring one to know θ^* .

The evaluation at $\theta=\theta^*$ calls into question the decision-theoretic definition of unbiasedness:

$$E_1(\hat{\theta}_n(\mathbf{X})) = \theta, \forall \theta \in \Theta, \quad (35)$$

in the context of frequentist estimation since this assertion makes sense only when defined at $\theta=\theta^*$. Similarly, the appropriate frequentist definition of the MSE for an estimator, initially proposed by Fisher [39], is defined at the point $\theta=\theta^*$:

$$MSE(\hat{\theta}_n(\mathbf{X}); \theta^*) = E(\hat{\theta}_n(\mathbf{X}) - \theta^*)^2, \text{ for } \theta^* \text{ in } \Theta. \quad (36)$$

Indeed, the well-known decomposition:

$$MSE(\hat{\theta}_n(\mathbf{X}); \theta^*) = Var(\hat{\theta}_n(\mathbf{X})) + [E(\hat{\theta}_n(\mathbf{X})) - \theta^*]^2, \text{ for } \theta^* \text{ in } \Theta, \quad (37)$$

is meaningful only when defined at the point $\theta=\theta^*$ (true mean) since by definition:

$$\begin{aligned} Var(\hat{\theta}_n(\mathbf{X})) &= E[\hat{\theta}_n(\mathbf{X}) - \theta_m]^2, \theta_m = E(\hat{\theta}_n(\mathbf{X})) \\ Bias(\hat{\theta}_n(\mathbf{X}); \theta^*) &= E(\hat{\theta}_n(\mathbf{X})) - \theta^*, \end{aligned} \quad (38)$$

and thus, the variance and the bias involve only two values of θ in Θ , θ_m and θ^* , and when $\theta_m = \theta^*$ the estimator is unbiased. This implies that the apparent affinity between the MSE_1 defined in Eq. (13) and the variance of an estimator is more apparent than real because the latter makes frequentist sense only when $\theta_m = E(\hat{\theta}_n(\mathbf{X}))$ is a single point.

5.3. James-Stein estimator from a frequentist perspective

For a proper frequentist evaluation of the above James-Stein result, it is important to bring out the conflict between the *overall MSE* (14) and the factual reasoning underlying frequentist estimation. From the latter perspective, the James-Stein estimator raises several issues of concern.

First, both the least-squares $\hat{\theta}_{LS}(\mathbf{X})$ and the James-Stein $\hat{\theta}_{JS}(\mathbf{X})$ estimators are *inconsistent* estimators of θ since the underlying model suffers from the incidental parameter problem: there is essentially one observation (X_k) for each unknown parameter (θ_k), and as $m \rightarrow \infty$ the number of unknown parameters increases at the same rate. To bring out the futility of comparing these two estimators more clearly, consider the following simpler example.

Example. Let $\mathbf{X} := (X_1, X_2, \dots, X_n)$ be a sample from the simple Normal model in Eq. (18). Comparing the two estimators $\hat{\theta}_1 = X_n$ and $\hat{\theta}_2 = \frac{1}{2}(X_1 + X_n)$ and inferring that $\hat{\theta}_2$ is relatively more efficient than $\hat{\theta}_1$ relative to a square loss function, i.e.,

$$\text{MSE}(\hat{\theta}_2(\mathbf{X}); \theta) = 1 < \text{MSE}(\hat{\theta}_1(\mathbf{X}); \theta) = \frac{1}{2}, \forall \theta \in \mathbb{R}, \quad (39)$$

is totally uninteresting because both estimators are inconsistent!

Second, to be able to discuss the role of admissibility in the Stein [37] result, we need to consider a consistent James-Stein estimator, by extending the original data to a panel (longitudinal) data where the sample is:

$\mathbf{X}_t := (X_{1t}, X_{2t}, \dots, X_{mt})$, $t = 1, 2, \dots, n$. In this case, the consistent least-squares and James-Stein estimators are:

$$\begin{aligned} \hat{\theta}_{LS}(\mathbf{X}) &= (\bar{X}_1, \bar{X}_2, \dots, \bar{X}_m), \text{ where } \bar{X}_k = \frac{1}{n} \sum_{t=1}^n X_{kt}, k = 1, 2, \dots, m, \\ \hat{\theta}_{JS}^+(\mathbf{X}) &= \left(1 - \frac{(m-2)\sigma^2}{\|\bar{\mathbf{X}}\|^2}\right)^+ \bar{\mathbf{X}}, \text{ where } \bar{\mathbf{X}} := (\bar{X}_1, \bar{X}_2, \dots, \bar{X}_m). \end{aligned} \quad (40)$$

This enables us to evaluate the notion of “relatively better” more objectively.

Admissibility relative to the overall loss function in Eq. (21) introduces a trade-off between the accuracy of the estimators for individual parameters $\theta := (\theta_1, \theta_2, \dots, \theta_m)$ and the “overall” expected loss. The question is: “In what sense the overall MSE among a group of mean estimates provides a better measure of “error” in learning about the true values $\theta^* := (\theta_1^*, \theta_2^*, \dots, \theta_m^*)$?” The short answer is: it does not. Indeed, the overall MSE will be irrelevant when the primary objective of estimation is to learn from data about θ^* . This is because the particular loss function penalizes the estimator’s capacity to pin-point θ^* by trading an increase in *bias* for a decrease in the overall MSE in Eq. (21), when the latter is misleadingly evaluated over all θ in $\Theta := \mathbb{R}^m$. That is, the James-Stein estimator flouts the primary objective of pinpointing θ^* in favor of reducing the overall MSE $\forall \theta \in \Theta$.

In summary, the above discussion suggests that there is nothing paradoxical about Stein’s [37] original result. What is problematic is not the least-squares estimator, but the choice of “better”

in terms of admissibility relative to an overall MSE in evaluating the accuracy of the estimators of θ .

5.4. Frequentist hypothesis testing

Another frequentist inference procedure one can employ to learn from data about θ^* is hypothesis testing, where the question posed is whether θ^* is close enough to some prespecified value θ_0 . In contrast to estimation, the reasoning underlying frequentist testing is *hypothetical* in nature.

5.4.1. Legitimate frequentist error probabilities

For testing the hypotheses:

$H_0: \theta \leq \theta_0$ vs. $H_1: \theta > \theta_0$, where θ_0 is a prespecified value,

one utilizes the same sampling distribution $\bar{X}_n \sim N(\theta, \frac{1}{n})$, but transforms the pivot $d(\mathbf{X}; \theta) := \sqrt{n}(\bar{X}_n - \theta^*)$ into the test statistic by replacing θ^* with the prespecified value θ_0 , yielding $d(\mathbf{X}) := \sqrt{n}(\bar{X}_n - \theta_0)$. However, instead of evaluating it under the factual $\theta = \theta^*$, it is now evaluated under various *hypothetical scenarios* associated with H_0 and H_1 to yield two types of (hypothetical) sampling distributions:

- (I) $d(\mathbf{X}) := \sqrt{n}(\bar{X}_n - \theta_0) \sim^{\theta=\theta^*} N(0, 1)$,
- (II) $d(\mathbf{X}) := \sqrt{n}(\bar{X}_n - \theta_0) \sim^{\theta=\theta^*} N(\delta_1, 1)$, $\delta_1 = \sqrt{n}(\theta_1 - \theta_0)$ for $\theta_1 > \theta_0$.

In both cases, (I) and (II), the underlying reasoning is hypothetical in the sense that the factual in Eq. (33) is replaced by hypothesized values of θ , and the test statistic $d(\mathbf{X})$ provides a standardized distance between the hypothesized values (θ_0 or θ_1) and θ^* the true θ , assumed to underlie the generation of the data $\mathbf{x}_0$, yielding $d(\mathbf{x}_0)$. Using the sampling distribution in (I), one can define the following legitimate error probabilities:

$$\begin{aligned} \text{significance level : } & \mathbb{P}(d(\mathbf{X}) > c_\alpha; H_0) = \alpha, \\ \text{p-value : } & \mathbb{P}(d(\mathbf{X}) > d(\mathbf{x}_0); H_0) = p(\mathbf{x}_0). \end{aligned} \quad (41)$$

Using the sampling distribution in (II), one can define:

$$\begin{aligned} \text{type II error prob. : } & \mathbb{P}(d(\mathbf{X}) \leq c_\alpha; \theta = \theta_1) = \beta(\theta_1), \text{ for } \theta_1 > \theta_0, \\ \text{power : } & \mathbb{P}(d(\mathbf{X}) > c_\alpha; \theta = \theta_1) = \rho(\theta_1), \text{ for } \theta_1 > \theta_0. \end{aligned} \quad (42)$$

It can be shown that the test T_α , defined by the test statistic $d(\mathbf{X})$ and the rejection region $C_1(\alpha) = \{\mathbf{x} : d(\mathbf{x}) > c_\alpha\}$, constitutes a uniformly most powerful (UMP) test for significance level α ; see Ref. [9]. The type I [II] error probability is associated with test T_α erroneously rejecting [accepting] H_0 . The type I and II error probabilities evaluate the generic capacity [whatever the sample realization $\mathbf{x} \in \mathbb{R}^n$] of a test to reach correct inferences. Contrary to Bayesian claims, these error probabilities have nothing to do with the temporal or the physical dimension of the long-run metaphor associated with repeated samples. The relevant feature of the long-run

metaphor is the repeatability (in principle) of the DGM represented by $\mathcal{M}_\theta(\mathbf{x})$; this feature can be easily operationalized using computer simulation; see Ref. [40].

The key difference between the significance level α and the p -value is that the former is a *pre-data* and the latter a *post-data* error probability. Indeed, the p -value can be viewed as the smallest significance level α at which H_0 would have been rejected with data $\mathbf{x}_0$. The legitimacy of postdata error probabilities underlying the hypothetical reasoning can be used to go beyond the N-P accept/reject rules and provide an evidential interpretation pertaining to the discrepancy γ from the null warranted by data $\mathbf{x}_0$; see Ref. [41].

Despite the fact that frequentist testing uses hypothetical reasoning, its main objective is also to learn from data about the true model $\mathcal{M}^*(\mathbf{x})=\{f(\mathbf{x};\theta^*)\}$, $\mathbf{x} \in \mathbb{R}_X^n$. This is because a test statistic like $d(\mathbf{X}):=\sqrt{n}(\bar{X}_n - \theta_0)$ constitutes nothing more than a scaled distance between θ^* [the value behind the generation of $\bar{x}_n$], and a hypothesized value θ_0 , with θ^* being replaced by its “best” estimator $\bar{X}_n$.

6. Revisiting loss and risk functions

The above discussion raises serious doubts about the role of loss functions and admissibility in evaluating learning from data $\mathbf{x}_0$ about θ^* . To understand why the decision-theoretic framing misrepresents the frequentist approach, one needs to consider the role of loss functions in statistical inference more generally.

6.1. Where do loss functions come from?

A closer scrutiny of the decision-theoretic set up reveals that the loss function needs to invoke “information from sources other than the data,” which is usually not readily available. Indeed, such information is available in very restrictive situations, such as acceptance sampling in quality control. In light of that, a proper understanding of the intended scope of statistical inference calls for distinguishing the special cases where the loss function is part and parcel of the available substantive information from those that no such information is either relevant or available.

Tiao and Box [25], p. 624, reiterated Fisher’s [42] distinction:

“Now it is undoubtedly true that on the one hand that situations exist where the loss function is at least approximately known (for example, certain problems in business) and sampling inspection are of this sort. ... On the other hand, a vast number of inferential problems occur, particularly in the analysis of scientific data, where there is no way of knowing in advance to what use the results of research will subsequently be put.”

Cox [43] went further and questioned this framing even in cases where the inference might involve a decision:

“The reasons that the detailed techniques [decision-theoretic] seem of fairly limited applicability, even when a fairly clear cut decision element is involved, may be (i) that, except in

such fields as control theory and acceptance sampling, a major contribution of statistical technique is in presenting the evidence in incisive form for discussion, rather than in providing mechanical presentation for the final decision. This is especially the case when a single major decision is involved. (ii) The central difficulty may be in formulating the elements required for the quantitative analysis, rather than in combining these elements via a decision rule." (p. 45)

Another important aspect of using loss functions in inference is that in practice they seem to be an add-on to the inference itself since they bring to the problem the information other than the data. In particular, the same statistical inference problem can give rise to very different decisions/actions depending on one's loss function. To illustrate that consider an example from [44]:

"... consider the case of a new drug whose effects are studied by a research scientist attached to the laboratory of a pharmaceutical company. The conclusion of the study may have different bearings on the action to be taken by (a) the scientist whose line of further investigation would depend on it, (b) the company whose business decisions would be determined by it, and (c) the Government whose policies as to health care, drug control, etc., would take shape on that basis." (p. 72)

In practice, each one of these different agents is likely to have a very different loss function, but their inferences should have a common denominator: the scientific evidence pertaining to θ^* , the true θ , that stems solely from the observed data.

6.2. Decisions vs. inferences

The above discussion brings out the crucial distinction between a "decision" and an "inference" stemming from data x_0 . Even before Wald [5] introduced the decision-theoretic perspective, Fisher [42] perceptively argued:

"In the field of pure research no assessment of the cost of wrong conclusions, or of delay in arriving at more correct conclusions can conceivably be more than a pretence, and in any case such an assessment would be inadmissible and irrelevant in judging the state of the scientific evidence." (pp. 25–26)

Tukey (1960) echoed Fisher's view by contrasting decisions vs. inferences:

"Like any other human endeavor, science involves many decisions, but it progresses by the building up of a fairly well established body of knowledge. This body grows by the reaching of conclusions — by acts whose essential characteristics differ widely from the making of decisions. Conclusions are established with careful regard to evidence, but without regard to consequences of specific actions in specific circumstances." (p. 425)

Hacking [45] brought out the key difference between an "inference pertaining to evidence" for or against a hypothesis, and a "decision to do something" as a result of an inference:

"... to conclude that an hypothesis is best supported is, apparently, to decide that the hypothesis in question is best supported. Hence it is a decision like any other. But this inference is

fallacious. Deciding that something is the case differs from deciding to do something. ... Hence deciding to do something falls squarely in the province of decision theory, but deciding that something is the case does not.” (p. 31)

This issue was elaborated upon by Birnbaum [15], p. 19:

“Two contrasting interpretations of the decision concept are formulated: *behavioral*, applicable to “decisions” in a concrete literal sense as in acceptance sampling; and *evidential*, applicable to “decisions” such as “reject H_0 ” in a research context, where the pattern and strength of statistical evidence concerning statistical hypotheses is of central interest.”

6.3. Loss functions vs. inherent distance functions

The notion of a loss function stemming from “information other than the data” raises another source of potential conflict. This stems from the fact that within each statistical model $\mathcal{M}_\theta(\mathbf{x})$ there exists an *inherent* statistic distance function, often relating to the log-likelihood and the score function, which constitutes information contained in the data; see Ref. [46].

It is well known that when the distribution underlying $\mathcal{M}_\theta(\mathbf{x})$ is normal, the *inherent distance function* for comparing estimators of the mean (θ) is the square:

$$ND(\hat{\theta}_n(\mathbf{X}); \theta^*) = (\hat{\theta}_n(\mathbf{X}) - \theta^*)^2. \quad (43)$$

On the other hand, when the distribution is *Laplace*, the relevant statistical distance function is the *absolute distance* (see Ref. [47]):

$$AD(\hat{\theta}_n(\mathbf{X}); \theta^*) = |\hat{\theta}_n(\mathbf{X}) - \theta^*|. \quad (44)$$

Similarly, when the distribution underlying $\mathcal{M}_\theta(\mathbf{x})$ is *uniform*, the inherent distance function is:

$$SUP(\hat{\theta}_n(\mathbf{X}); \theta^*) = \sup_{\mathbf{x} \in \mathbb{R}_x^n} |\hat{\theta}_n(\mathbf{x}) - \theta^*|. \quad (45)$$

Note that these distance functions are defined at the point $\theta = \theta^*$ and not for all θ in Θ , as traditional loss functions.

The dilemma facing a Bayesian or a decision-theoretic statistician is to decide when it makes sense to override the MLE and select the optimal rule stemming from an externally given loss function. The dilemma is not as trivial as it might seem at first sight for two reasons. First, the key difference between the two is that the assumptions of the likelihood function $L(\theta)$ are testable vis-a-vis the data, but those underlying the loss function are not. Second, the likelihood function renders the notion of efficiency “global,” *full efficiency*, in terms of Fisher’s information:

$$CR(\theta^*) = I_n^{-1}(\theta^*), I_n(\theta^*) := E \left(- \frac{\partial^2 \ln L(\theta)}{\partial \theta \partial \theta^T} \right). \quad (46)$$

Hence, the optimality of an estimator can be affirmed using testable information comprising the statistical model $\mathcal{M}_\theta(\mathbf{x})$. This is in direct contrast with admissibility, which is a property

defined in terms of “local” efficiency—*relative* to a loss function—based on external (nontestable) information.

6.4. Acceptance sampling vs. learning from data

Let us bring out the key features of a situation where the above decision-theoretic set up makes perfectly good sense. This is the situation Fisher [12] called *acceptance sampling*, such as an industrial production process where the objective is quality control, i.e., to make a decision pertaining to shipping sub-standard products (e.g., nuts and bolts) to a buyer using the expected loss/gain as the ultimate criterion.

In an acceptance sampling context, the $MSE(\hat{\theta}(\mathbf{X}); \theta)$, or some other risk function, are relevant because they evaluate genuine losses associated with a decision related to the choice of an estimate $\hat{\theta}(\mathbf{x}_0)$, say the cost of the observed percentage of defective products, but that has nothing to do with type I and II error probabilities.

Acceptance sampling differs from a scientific enquiry in two crucial respects:

- a. The primary aim is to use statistical rules to minimize the expected loss associated with “a decision.”
- b. The sagacity of all actions is determined by the respective “losses” stemming from “*relevant information other than the data*” ([32], p. 251)."
- c. The trade-off between the two types of error probabilities is determined by the risk function itself and not by any endeavor to learn from data about θ^* . Indeed, the learning is deliberately undermined by certain loss function such as the *overall* MSE (14) that favor biased estimators of the James-Stein type.

The key difference between acceptance sampling and a scientific inquiry is that the primary objective of the latter is *not* to minimize expected loss (costs and utility) associated with different values of $\theta \in \Theta$, but to use data $\mathbf{x}_0$ to learn about the “true” model (17). The two situations are drastically different mainly because the key notion of a “true θ ” calls into question the above acceptance sampling set up. Indeed, the loss function being defined “ $\forall \theta \in \Theta$,” will penalize θ^* , since there is no reason to expect that the highest ranked θ would coincide with θ^* , unless by accident.

The extreme relativism of loss function optimality renders decision-theoretic and Bayes rules highly vulnerable to abuse. In practice, one can justify any estimator as optimal, however lame in terms of other criteria, by selecting an “appropriate” loss function.

Example 1. Consider a manufacturer of high precision bolts and nuts who has information that the buyer only checks the first and last box for quality control when accepting an order. This suggests that to minimize losses, stemming from the return of its products as defective, an appropriate loss function might be:

$$L(\mathbf{X}; \theta) = \left([(X_1 + X_n)/2] - \theta \right)^2, \theta \in (0, 1). \quad (47)$$

From the *acceptance sampling* perspective, the “optimal” estimator $\tilde{\theta} = (X_1 + X_n)/2$ is excellent because it minimizes the expected losses, but it is a terrible estimator for pinpointing θ^* because it is inconsistent!

Consider a more general case where acceptance sampling resembles hypothesis testing in so far as final products are randomly selected for inspection during the production process. In such a situation the main objective can be viewed as operationalizing the probabilities of false acceptance/rejection with a view to minimize the expected losses. The conventional wisdom has been that this situation is similar enough to Neyman-Pearson (N-P) testing to render the latter as the appropriate framing for the decision to ship this particular batch or not. However, a closer look at some of the examples used to illustrate such a situation [48], reveals that the decisions are driven exclusively by the risk function and not by any quest to learn from data about the true θ^* . For instance, N-P way of addressing the trade-off between the two types of error probabilities, fixing α to a small value and seek a test that minimizes the type II error probability, seems utterly irrelevant in such a context. One can easily think of a loss function where the “optimal” trade-off calls for a much larger type I than type II error probability. As argued in Ref. [14]:

“Wald’s decision theory ... has given up fixed probability of errors of the first kind, and has focused on gains, losses or regrets.” (p. 433)

Indeed, Wald [5] was the first to highlight that the decision-theoretic notion of “optimality” revolves around a particular loss function:

“The “best” system of regions of acceptance ... will depend only on the weight function of the errors.” ([5], p. 302)

Given the crucial differences in [a]–[c], one can make a strong case that the objectives and the underlying reasoning of acceptance sampling are drastically different from those pertaining to learning from data in a scientific context.

6.5. Is expected loss a legitimate frequentist error?

The key question is: what do expected losses and traditional frequentist errors, such as bias, MSE and the type I–II errors, have in common, if anything?

First, they stem directly from the statistical model $\mathcal{M}_\theta(\mathbf{x})$ since the underlying sampling distributions of estimators, test statistics, and predictors are derived *exclusively from the distribution of the sample* $f(\mathbf{x}; \theta)$ through Eq. (7). In this sense, the relevant error probabilities are directly related to statistical information pertaining to the data as summarized by the statistical model $\mathcal{M}_\theta(\mathbf{x})$ itself.

Second, they are attached to a particular frequentist inference procedure as they relate to a relevant inferential claim. These error probabilities calibrate the effectiveness of inference procedures in learning from data about the true statistical model $\mathcal{M}^*(\mathbf{x}) = \{f(\mathbf{x}; \theta^*)\}$, $\mathbf{x} \in \mathbb{R}_X^n$.

In light of these features, the question is: “in what sense a risk function could potentially represent relevant frequentist errors?” That argument that the risk function represents legitimate frequentist errors because it is derived by taking expectations with respect to $f(\mathbf{x}; \theta)$, $\mathbf{x} \in \mathbb{R}_X^n$ [3], is misguided for two reasons.

- a. The relevant errors in estimation, including the bias $E(\hat{\theta}_n(\mathbf{X})) - \theta^*$ and MSE $E(\hat{\theta}_n(\mathbf{X}) - \theta^*)^2$, are evaluated with respect to $f(\mathbf{x}; \theta^*)$, $\mathbf{x} \in \mathbb{R}_X^n$, by invoking factual reasoning; θ^* denotes the state of Nature. Wald's [5] original loss function in Eq. (2) represents an interesting case because it is defined in terms of θ^* , which renders it nonoperational when evaluated for all θ in Θ , since θ^* is unknown in practice. In contrast, the errors associated with the bias and MSE are rendered operational by the factual reasoning fashioned to forgo knowing θ^* .
- b. The expected losses stemming from the risk function $R(\theta, \hat{\theta})$ are attached to particular values of θ in Θ . Such an assignment is in direct conflict with all the above legitimate error probabilities that are attached to the inference procedure itself, and never to the particular values of θ in Θ . The expected loss assigned to each value of θ in Θ has nothing to do with learning from data about θ^* . Indeed, the risk function will penalize a procedure for pinpointing θ^* since the latter is unknown in practice. This is in direct conflict with the main objective of frequentist estimation but in sync with "acceptance sampling," where the objective of the inference has everything to do with expected losses.

7. Summary and conclusions

The paper makes a case for Fisher's [12, 42] assertions concerning the appropriateness of the decision-theoretic framing for "acceptance sampling" and its inappropriateness for frequentist inference. A closer look at this framing reveals that it is congruent with the Bayesian approach because it supplements the posterior distribution with a theory of optimal inference. Decision-theoretic and Bayesian rules are considered optimal when they minimize the expected loss for all possible values of θ [$\forall \theta \in \Theta$], irrespective of what the true value θ^* happens to be. In contrast, the theory of optimal frequentist inference revolves around the true value θ^* , since it depends entirely on the capacity of the procedure to pinpoint θ^* . The frequentist approach relies on factual (estimation and prediction), as well as hypothetical (testing) reasoning, both of which revolve around the existential quantifier $\exists \theta^* \in \Theta$. The inappropriateness of the quantifier $\forall \theta \in \Theta$ calls into question the relevance of *admissibility* as a minimal property for frequentist estimators. A strong case can be made that the relevant minimal property for frequentist estimators is *consistency*. In addition, full efficiency provides the relevant measure of an estimator's finite sample efficiency (accuracy) in pinpointing θ^* . Both of these properties stem from the underlying statistical model $\mathcal{M}_\theta(\mathbf{x})$, in contrast to admissibility which relies on loss functions based on information other than the data.

It is argued that Stein's [36] result stems from the fact that admissibility introduces a trade-off between the accuracy of the estimator in pinpointing θ^* and the "overall" expected loss. That is, the James-Stein estimator achieves a higher overall MSE by blunting the capacity of a frequentist estimator to pinpoint θ^* . Why would a frequentist care about the overall MSE defined for all θ in Θ ? After all, expected losses are not legitimate errors similar to bias and MSE (when properly defined), as well as coverage, type I and II errors. The latter are attached to the frequentist procedures themselves to calibrate their capacity to achieve learning from data about θ^* . In contrast, expected losses are assigned to different values of θ in Θ , using information other than the data.

Author details

Aris Spanos

Address all correspondence to: aris@vt.edu

Department of Economics, Virginia Tech, Blacksburg, VA, USA

References

- [1] Wald, A. Statistical Decision Functions. NY: Wiley; 1950.
- [2] Berger, J.O. Statistical Decision Theory and Bayesian Analysis, 2nd ed., NY: Springer; 1985.
- [3] Robert, C.P. The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation, 2nd ed. NY: Springer; 2001.
- [4] O'Hagan, A. Bayesian Inference, London: Edward Arnold; 1994.
- [5] Wald, A. Contributions to the theory of statistical estimation and testing hypotheses. *Annals of Mathematical Statistics*. 1939; **10**: 299–326.
- [6] Neyman, J. Outline of a theory of statistical estimation based on the classical theory of probability. *Philosophical Transactions of the Royal Society of London, Series A*, 1937; **236**: 333–380.
- [7] Neyman, J. Lectures and Conferences on Mathematical Statistics and Probability, 2nd ed. Washington: U.S. Department of Agriculture; 1952.
- [8] Neyman, J. Foundations of Behavioristic Statistics. Godambe, V. and Sprott, D. eds., *Foundations of Statistical Inference*. Toronto: Holt, Rinehart and Winston of Canada; 1971: pp. 1–13.
- [9] Lehmann, E.L. Testing Statistical Hypotheses. NY: Wiley; 1959.
- [10] LeCam, L. Asymptotic Methods in Statistical Decision Theory. NY: Springer; 1986.
- [11] Neyman, J.A. Selection of Early Statistical Papers by J. Neyman. Moss Landing, CA: University of California Press; 1967.
- [12] Fisher, R.A. Statistical methods and scientific induction. *Journal of the Royal Statistical Society, B*, 1955; **17**: 69–78.
- [13] Cox, D.R. Some problems connected with statistical inference. *The Annals of Mathematical Statistics*. 1958; **29**: 357–372.
- [14] Tukey, J.W. Conclusions vs Decisions. *Technometrics*. 1960; **2**: 423–433.
- [15] Birnbaum, A. The Neyman-Pearson theory as decision theory, and as inference theory; with a criticism of the Lindley-Savage argument for Bayesian Theory. *Synthese*. 1977; **36**: 19–49.

- [16] Ferguson, T.S. *Mathematical Statistics: A Decision Theoretic Approach*. London: Academic Press. 1967.
- [17] Wasserman, L. *All of Statistics*. NY: Springer; 2004.
- [18] Bansal, A.K. *Bayesian Parametric Inference*. Oxford: Alpha Science. 2007.
- [19] Ferguson, T.S. Development of the Decision Model. In: *On the History of Statistics and Probability*, edited by D. B. Owen. NY: Marcel Dekker. 1976; ch. 16.
- [20] Wald, A. An essentially complete family class of admissible decision functions. *Annals of Mathematical Statistics*. 1947; **18**: 549–555.
- [21] LeCam, L. An extension of Wald's theory of statistical decision functions. *Annals of Mathematical Statistics*. 1955; **26**: 69–81.
- [22] Spanos, A. and Mayo, D.G. Error statistical modeling and inference: where methodology meets ontology. *Synthese*. 2015; **192**: 3533–3555.
- [23] Reid, N. Statistical Sufficiency. *International Encyclopedia of the Social & Behavioral Sciences*, edited by Wright J.D., 2nd edition, Vol 23. Oxford: Elsevier. 2015; pp. 418–422.
- [24] Lindley, D.V. *Introduction to Probability and Statistics from a Bayesian Viewpoint, Part 2: Inference*. Cambridge: Cambridge University Press; 1965.
- [25] Tiao, G.C. and Box, G.E.P. Some comments on “Bayes” estimators. *Studies in Bayesian Econometrics and Statistics, In Honor of Leonard Savage J.*, edited by Fienberg S.E. and Zellner A. Amsterdam: North-Holland. 1975; pp. 619–626.
- [26] Spanos, A. *Revisiting Bayes' Rule and Evidence: Notional Events vs. Real Data*. Virginia Tech, working paper; Blacksburg, VA. 2015.
- [27] Schervish, M.J. *Theory of Statistics*. New York, NY: Springer-Verlag; 1995.
- [28] Srivastava, M.K., Khan A.H. and Srivastava, N. *Statistical Inference: Theory of Estimation*, Delhi, India: PHI Learning; 2014.
- [29] Rubin, D.B. Bayesianly justifiable and relevant frequency calculation for the applied statistician. *Annals of Statistics*. 1984; **12**: 1151–1172.
- [30] Gelman, A., Carlin J.B. and Rubin, D.B., *Bayesian Data Analysis*, 2nd edition, London: Chapman & Hall; 2004.
- [31] Savage, L.J. On rereading Fisher, R.A. *The Annals of Statistics*. (1976); **4**(3): 441–500.
- [32] Cox, D.R. and Hinkley, D.V. *Theoretical Statistics*. London: Chapman & Hall; 1974.
- [33] Lehmann, E.L. Specification Problems in the Neyman-Pearson-Wald Theory. *Statistics: An Appraisal*, edited by David, H.A. and David, H.T., IA: The Iowa State University Press; 1984; 425–436.
- [34] Efron, B. and Morris, C.N. Stein's estimation rule and its competitors—an empirical Bayes approach. *Journal of the American Statistical Association*. 1973; **68**: 117–130.

- [35] Saleh, A.K. Md. E. Theory of Preliminary Test and Stein-Type Estimation with Applications. NY: Wiley; 2006.
- [36] Stein, C. Inadmissibility of the usual estimator for the mean of a multivariate distribution. Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability. 1956; 1: 197–206.
- [37] James, W. and Stein, C. Estimation with quadratic loss. Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability. 1961; 1: 361–379.
- [38] Spanos, A. Where Do Statistical Models Come From? Revisiting the Problem of Specification. The Second Erich L. Lehmann Symposium, edited by Rojo J., Lecture Notes-Monograph Series, vol. 49, Hayward CA: Institute of Mathematical Statistics. 2006; pp. 98–119.
- [39] Fisher, R.A. A Mathematical examination of the methods of determining the accuracy of an observation by the mean error, and by the mean square error. Monthly Notices of the Royal Astronomical Society. 1920; 80: 758–770.
- [40] Spanos, A. A frequentist interpretation of probability for model-based inductive inference. Synthese. 2013; 190: 1555–1585.
- [41] Mayo, D.G. and Spanos, A. Severe testing as a basic concept in a Neyman-Pearson philosophy of induction. The British Journal for the Philosophy of Science; 2006; 57: 323–357.
- [42] Fisher, R.A. The Design of Experiments. Edinburgh: Oliver and Boyd; 1935.
- [43] Cox, D.R. Foundations of statistical inference: the case for eclecticism. Australian Journal of Statistics. 1978; 20: 43–59.
- [44] Chatterjee, S.K. Statistical Thought: A Perspective and History. Oxford: Oxford University Press. 2002.
- [45] Hacking, I. Logic of Statistical Inference. Cambridge: Cambridge University Press; 1965.
- [46] Casella, G. and Berger, R.L. Statistical Inference, 2nd ed. CA: Duxbury. 2002.
- [47] Shao, J. Mathematical Statistics. 2nd ed. NY: Springer. 2003.
- [48] Silvey, S.D. Statistical Inference, London: Chapman & Hall. 1975.