

# A Statistical Decision-Theoretical Perspective on the Two-Stage Approach to Parameter Estimation

Braghadeesh Lakshminarayanan and Cristian R. Rojas

**Abstract**—One of the most important problems in system identification and statistics is how to estimate the unknown parameters of a given model. Optimization methods and specialized procedures, such as Empirical Minimization (EM) can be used in case the likelihood function can be computed. For situations where one can only simulate from a parametric model, but the likelihood is difficult or impossible to evaluate, a technique known as the Two-Stage (TS) Approach can be applied to obtain reliable parametric estimates. Unfortunately, there is currently a lack of theoretical justification for TS. In this paper, we propose a statistical decision-theoretical derivation of TS, which leads to *Bayesian* and *Minimax* estimators. We also show how to apply TS on models for independent and identically distributed samples, by computing quantiles of the data as a first step, and using a linear function as the second stage. The proposed method is illustrated via numerical simulations.

**Index Terms**—Two-Stage approach, estimation theory, statistical decision theory.

## I. INTRODUCTION

Engineering disciplines rely on mathematical models in order to understand, modify and control physical systems. To derive such models, practitioners utilize techniques from several research fields, such as statistics [1], chemometrics [2], econometrics [3], machine learning [4] and system identification [5, 6], which build models based on data gathered from a physical system.

Many models built from data are *parametric statistical models*, *i.e.*, they define (explicitly or implicitly) a probability distribution that describes the observed data in terms of a set of unknown parameters, which remain to be estimated. Many techniques have been devised for estimating such parameters, including Maximum Likelihood [1, 7], Prediction Error Methods [5], Instrumental Variables [8], Methods of Moments [3], and so on. While some of these techniques achieve statistical efficiency [1], their implementations have several drawbacks: they may require having an explicit expression for the likelihood function, solving non-convex optimization problems, or that the distribution of the samples has a very specific structure. To address some of these issues, novel techniques have been introduced in econometrics and identification, such as Indirect Inference and the Method of Simulated Moments [9], and the Two-Stage (TS) Approach [10, 11].

This work has been supported by the Swedish Research Council under contract number 2016-06079 (NewLEADS) and by the Digital Futures project EXTREMUM. The authors are with the Division of Decision and Control Systems, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden (e-mails: blak@kth.se; crro@kth.se).

Out of these new methods, TS is an attractive alternative to traditional parametric techniques, since it only requires one to be able to simulate data for a fixed choice of the parameters, and, by properly choosing its second stage, it requires only solving convex optimization problems (thus avoiding getting trapped into local optima). The TS approach works by seeing the estimation problem as an inverse *supervised learning* setup [4], where one simulates a large number of samples for different values of the parameters, and uses any supervised learning method to predict the values of the parameters for a given data sample. The supervised learning problem is decomposed into two stages, where the first stage involves “compressing” the data into a small set of values, and the second one uses such compressed samples to “predict” the values of the parameters. In spite of its advantages with respect to other approaches, however, TS relies on several user choices, whose impact on the performance of the method is not yet clear.

In this paper we provide a formalization of TS, by deriving it within a statistical decision theory setup. This allows us to note that the standard TS is actually a Bayes estimator (with respect to a given prior on the parameters), and to extend the approach to minimax decision rules. In addition, we focus on the case of independent and identically distributed (*i.i.d.*) data, for which we suggest a specific choice for its first (or “compression”) and second stages, which leads to convex optimization problems that can be efficiently solved using off-the-shelf solvers.

In summary, the main contributions of this paper are:

- We provide a statistical decision-theoretical derivation of TS, which leads to Bayes and minimax formulations, and give a reason for decomposing the problem into two stages;
- we suggest a specific structure for the second stage of TS, which leads to simple convex programs for both Bayes and minimax formulations;
- for *i.i.d.* data, we propose a specific choice for the first stage of TS, based on sample quantiles of the data; and
- we illustrate the performance of the novel Bayes and minimax TS formulations with a numerical example based on a Weibull distribution.

The rest of this work is organized as follows. In Section II, the statistical decision theory setup for parameter estimation is presented. Section III briefly describes the procedure of the TS approach. Section IV describes a computational way to approximate Bayes and minimax estimators, and explains their connection to TS. In Section V, the Bayes

and minimax formulations of TS are specialized for *i.i.d.* data. A simulation study assessing the performance of the above formulations is presented in Section VI, and finally conclusions are drawn in Section VII.

## II. PROBLEM FORMULATION

Consider an environment, or data generating mechanism, that produces a quantity  $\mathbf{y} \in \mathcal{Y}$  according to a probability distribution  $\mathbb{P}(\mathbf{y}|\theta)$  determined by an unknown parameter  $\theta \in \Theta \subseteq \mathbb{R}^d$ , where  $d \geq 1$ . Based on the value of  $\mathbf{y}$ , the goal is to design an *estimator* (or *decision rule*)  $\delta: \mathcal{Y} \rightarrow \Theta$  of the true unknown parameter  $\theta$  such that the *risk*

$$R(\theta, \delta) := \mathbb{E}_{\mathbf{y} \sim \mathbb{P}(\mathbf{y}|\theta)} \{L(\theta, \delta(\mathbf{y}))\} \quad (1)$$

is minimized, where  $L: \Theta \times \Theta \rightarrow \mathbb{R}_0^+$  is a *loss* function. In other words,  $L(\theta, \delta(\mathbf{y}))$  is the loss of estimating the parameter as  $\delta(\mathbf{y})$  when its true value is  $\theta$ . Since the risk depends on  $\theta$ , to evaluate the performance of estimators  $\delta(\mathbf{y})$ , it is required to reduce the risk to a function that depends only on  $\delta$ . The notion of an “optimal” decision rule is then defined with respect to this function. In this paper, we consider the following types of optimal rules:

- 1) **Bayes rules:** Given a probability distribution  $\mathbb{P}(\theta)$  on  $\Theta$ , a decision rule  $\delta_{\text{Bayes}}^*$  is said to be a *Bayes rule* (with respect to  $\mathbb{P}(\theta)$ ) if it minimizes  $\mathbb{E}_{\theta \sim \mathbb{P}(\theta)} [R(\theta, \delta)]$ .
- 2) **Minimax rules:** A decision rule  $\delta_{\text{minimax}}^*$  is called a *minimax rule* if it minimizes  $\max_{\theta \in \Theta} R(\theta, \delta)$ .

In some of the subsequent sections, we will consider the case where  $\mathbf{y} = (y^1, \dots, y^N)$ , *i.e.*,  $\mathbf{y}$  is a vector consisting of multiple observations.

## III. TS APPROACH

In this section, we summarise the working principle of the TS approach; please refer to [10, 11] for details. The TS approach works by posing the estimation problem as an inverse *supervised learning* setup [4], where different values of the unknown parameter, given by the set  $\{\theta_i\}_{i=1}^{M_\theta}$ , are drawn from a fixed prior, and for each  $\theta_i$ , large number of samples, represented by vector  $\mathbf{y}_i = (y_i^1, \dots, y_i^N)$  are drawn. Thus,  $\{(\mathbf{y}_i, \theta_i)\}_{i=1}^{M_\theta}$  constitutes a training set. The supervised learning problem is then decomposed into two stages, where the first stage involves “compressing” the data  $\mathbf{y}_i = (y_i^1, \dots, y_i^N)$  into a small set of values  $\alpha_i = \{\alpha_i^1, \dots, \alpha_i^n\}$ , where  $n \ll N$ , and the second one uses such compressed samples on a supervised learning algorithm (such as Kernel ridge regression, XGBoost, or Deep Neural Networks) to train a model that can “predict” the values of the parameters by minimizing the training error. This trained model (including the compression stage) constitutes the sought TS estimator.

## IV. A STATISTICAL DECISION-THEORETICAL FORMULATION OF TS

In this section, we will show how to derive, in a computational manner, Bayesian and minimax decision rules, and describe their connection to the Two-Stage Approach.

### A. Bayes estimator

Expanding the expression for  $R(\theta, \delta)$  in (1), we have

$$R(\theta, \delta) = \mathbb{E}_{\mathbf{y} \sim \mathbb{P}(\mathbf{y}|\theta)} [L(\theta, \delta(\mathbf{y}))] = \int_{\mathcal{Y}} L(\theta, \delta(\mathbf{y})) \mathbb{P}(\mathbf{y}|\theta) d\mathbf{y}.$$

Now, we define *weighted risk*

$$R_{\text{weighted}}(\delta) = \int_{\Theta} R(\theta, \delta) \pi(\theta) d\theta, \quad (2)$$

where  $\pi: \Theta \rightarrow \mathbb{R}_0^+$  is a weighting function over  $\Theta$ . Note that  $\pi$  need not be a probability density function (*pdf*) over  $\Theta$ , that is,  $\int_{\Theta} \pi(\theta) d\theta$  does not need to be equal to 1 (or even finite). However, if  $\pi(\theta) = \mathbb{P}(\theta)$  happens to be a pdf over  $\mathbb{R}^d$ , then we obtain the definition of *Bayes risk*,

$$R_{\text{Bayes}}(\delta) = \int_{\Theta} R(\theta, \delta) \mathbb{P}(\theta) d\theta, \quad (3)$$

where  $\mathbb{P}(\theta)$  is called a *prior* on  $\theta$ . The minimizer of  $R_{\text{Bayes}}$  is called the *Bayes rule* associated with  $\mathbb{P}(\theta)$ .

The standard approach to derive the Bayes rule is as follows. We first expand (3):

$$R_{\text{Bayes}}(\delta) = \int_{\Theta} \int_{\mathcal{Y}} L(\theta, \delta(\mathbf{y})) \mathbb{P}(\mathbf{y}|\theta) \mathbb{P}(\theta) d\mathbf{y} d\theta \quad (4)$$

It is clear from (4) that, in order to minimize  $R_{\text{Bayes}}$  we have to minimize the inner integral  $\int_{\Theta} L(\theta, \delta(\mathbf{y})) \mathbb{P}(\theta|\mathbf{y}) d\theta$  over  $\delta$ , since  $\mathbb{P}(\mathbf{y})$  is non-negative. Therefore, the Bayes decision rule is given by  $\delta_{\text{Bayes}}^* = \arg \min_{\delta} \int_{\Theta} L(\theta, \delta(\mathbf{y})) \mathbb{P}(\theta|\mathbf{y}) d\theta$ .

We now propose an alternative computational approach to derive the Bayes rule associated with a prior  $\mathbb{P}(\theta)$ , which essentially corresponds to the TS method [10, 11]. To this end,

$$R_{\text{Bayes}}(\delta) \approx \int_{\Theta} \left[ \frac{1}{M_{\mathbf{y}}} \sum_{j=1}^{M_{\mathbf{y}}} L(\theta, \delta(\mathbf{y}_j)) \right] \mathbb{P}(\theta) d\theta. \quad (5)$$

where  $\mathbf{y}_1, \dots, \mathbf{y}_{M_{\mathbf{y}}}$  are independent random samples from  $\mathbb{P}(\mathbf{y}|\theta)$ . We can further approximate the outer integral of (5) by Monte-Carlo sampling, by taking  $M_\theta$  independent samples  $\theta_1, \dots, \theta_{M_\theta}$  from  $\mathbb{P}(\theta)$ :

$$R_{\text{Bayes}}(\delta) \approx \frac{1}{M_{\mathbf{y}}} \frac{1}{M_\theta} \sum_{i=1}^{M_\theta} \sum_{j=1}^{M_{\mathbf{y}}} L(\theta_i, \delta(\mathbf{y}_{ij})), \quad (6)$$

where  $\{\mathbf{y}_{ij}\}$  *i.i.d.*  $\mathbb{P}(\mathbf{y}|\theta_i)$  and  $\{\theta_i\}$  *i.i.d.*  $\mathbb{P}(\theta)$ , for  $i = 1, \dots, M_\theta$  and  $j = 1, \dots, M_{\mathbf{y}}$ .

From (6), it follows that one can approximate the Bayes rule by solving the optimization problem

$$\tilde{\delta}_{\text{Bayes}} := \arg \min_{\delta \in \Delta} \frac{1}{M_{\mathbf{y}}} \frac{1}{M_\theta} \sum_{i=1}^{M_\theta} \sum_{j=1}^{M_{\mathbf{y}}} L(\theta_i, \delta(\mathbf{y}_{ij})), \quad (7)$$

where  $\Delta$  is a suitable function space. If  $M_{\mathbf{y}}$ ,  $M_\theta$  are sufficiently large, and  $\Delta$  is sufficiently large, one should expect that  $\tilde{\delta}_{\text{Bayes}} \approx \delta_{\text{Bayes}}^*$  in a suitable sense.

### Choosing the function space

The next step is to choose a function space  $\Delta$ . To this end, notice that typically  $\mathcal{Y}$  is a product space, *i.e.*,  $\mathcal{Y} = \tilde{\mathcal{Y}}^N$ , so that  $\mathbf{y} \in \mathcal{Y}$  is of the form  $\mathbf{y} = (y^1, \dots, y^N)^T$ , where  $y^i \in \tilde{\mathcal{Y}}$  ( $i = 1, \dots, N$ ) are the individual *samples* in the data, and  $N$  is the *sample size*. Furthermore,  $\mathbb{P}(\mathbf{y}|\theta)$  may also have a product structure, *i.e.*, if  $(y^1, \dots, y^N)^T \sim \mathbb{P}(\mathbf{y}|\theta)$ , then  $y^i \stackrel{i.i.d.}{\sim} \mathbb{P}(\tilde{y}|\theta)$ , or  $\mathbb{P}(\mathbf{y}|\theta)$  may correspond to a segment of a stationary distribution; in those cases, the measure  $\mathbb{P}(\mathbf{y}|\theta)$  may be subject to a measure concentration phenomenon [12], which means that  $\mathbb{P}(\mathbf{y}|\theta)$  may be highly concentrated on a subset of  $\mathcal{Y}$ . This implies that the decision rule  $\delta_{\text{Bayes}}^*$  may be difficult to approximate using standard functional approximation methods. To overcome this issue, we will add the subscript  $N$  to  $\delta$ , to emphasize the dependence of the number of samples  $N$ , and express  $\delta_N$  as the composition of two functions,

$$\delta_N = g_N \circ h_N, \quad (8)$$

where  $h_N: \mathcal{Y} \rightarrow \mathbb{R}^n$  and  $g_N: \mathbb{R}^n \rightarrow \Theta$ , with  $n \ll N$ . For identifiability purposes [7], we require that  $N$  is larger than  $n$ . Function  $h$  can be chosen as a fixed function (*i.e.*, not to be optimized with respect to) with the property that if  $\mathbf{y} \sim \mathbb{P}(\mathbf{y}|\theta)$ , then  $h_N(\mathbf{y}) \rightarrow \bar{h}(\mathbf{y})$  in probability as  $N \rightarrow \infty$ , where  $\bar{h}: \mathcal{Y} \rightarrow \mathbb{R}^n$  is differentiable and does not depend on  $N$ . This means that the measure concentration of  $\mathbb{P}(\mathbf{y}|\theta)$  as  $N$  grows can be taken care of by  $h_N$ , and if  $N$  is sufficiently large,  $g_N$  can be designed independently of  $N$ , *i.e.*,  $g_N \equiv g$  for a fixed function  $g: \mathbb{R}^n \rightarrow \Theta$ .

A possibility for  $h_N$  is to use an estimator for a *misspecified model* of  $\mathbf{y}$ , *i.e.*, a model which is not related to  $\mathbb{P}(\mathbf{y}|\theta)$ , but which is nevertheless easy to estimate (using, for instance, least squares [1]). The value of  $h_N(\mathbf{y})$  should serve as a *approximate sufficient statistic* [13] for  $\theta$ , in the sense of containing, approximately, the same information about  $\theta$  as  $\mathbf{y}$ . Notice also that this resembles the use of misspecified models in techniques such as Indirect Inference [14].

Once  $h_N$  is chosen, we are left with the problem of determining  $g$  according to

$$g^* = \arg \min_{g \in G} \frac{1}{M_{\mathbf{y}}} \frac{1}{M_{\theta}} \sum_{i=1}^{M_{\theta}} \sum_{j=1}^{M_{\mathbf{y}}} L(\theta_i, g(h_N(\mathbf{y}_{ij}))). \quad (9)$$

Here,  $G$  is a function space that can be chosen according to general functional approximation methods, such as deep neural networks, boosted decision trees, radial basis functions, or finite dimensional function spaces (*e.g.*, polynomials, piecewise affine functions, ...).

*Remark 4.1:* Under *i.i.d.* or stationarity assumptions, if  $N$  is sufficiently large, each sample will give a good estimate of the loss (due to a concentration of measure), so we can take  $M_{\mathbf{y}} = 1$  without significant loss in performance.

The approach given here to approximate  $\delta_{\text{Bayes}}^*$  is essentially identical to the TS approach given in [10], where the first stage consists in the computation of  $h_N(\mathbf{y}_{ij})$  (called the “compression stage”), and in the second stage  $g^*$  is determined. However, our formulation clarifies that the way

how the values of  $\theta$  are sampled (via  $\mathbb{P}(\theta)$ ) affects the resulting estimate, which is in fact an approximate Bayes estimator; for this reason, the estimator  $\delta_N^*$  will be called the *Bayes TS estimator* of  $\theta$  under the prior  $\mathbb{P}(\theta)$ .

### B. Minimax estimator

In this subsection we modify the previous formulation in order to arrive at a minimax estimator, given by  $\delta_{\text{minimax}}^* = \arg \min_{\delta \in \Delta} [\max_{\theta \in \Theta} R(\theta, \delta)]$ .

Let us rewrite the inner optimization problem  $\max_{\theta \in \Theta} R(\theta, \delta)$  as

$$\max_{\theta \in \Theta} R(\theta, \delta) = \max_{\pi \in \mathcal{P}(\Theta)} \int_{\Theta} R(\theta, \delta) \pi(\theta) d\theta, \quad (10)$$

where  $\mathcal{P}(\Theta)$  is the space of probability distributions in  $\Theta$ , and  $\pi(\cdot)$  can be interpreted as a prior distribution on  $\theta$ .

$$\delta_{\text{minimax}}^* = \min_{\delta \in \Delta} \max_{\pi \in \mathcal{P}(\Theta)} \int_{\Theta} \int_{\mathcal{Y}} L(\theta, \delta(\mathbf{y})) \mathbb{P}(\mathbf{y}|\theta) d\mathbf{y} \pi(\theta) d\theta, \quad (11)$$

assuming that we can interchange the integrals in (11).

Using Monte-Carlo sampling, we approximate (11) to obtain the following optimization problem

$$\min_{\delta \in \Delta} \max_{\pi \in \mathcal{P}(\Theta)} \int_{\Theta} \left[ \frac{1}{M_{\mathbf{y}}} \sum_{j=1}^{M_{\mathbf{y}}} L(\theta, \delta(\mathbf{y}_{j,\theta})) \right] \pi(\theta) d\theta, \quad (12)$$

where  $\{\mathbf{y}_{j,\theta}\} \stackrel{i.i.d.}{\sim} \mathbb{P}(\mathbf{y}|\theta)$  for a fixed  $\theta$ , and  $M_{\mathbf{y}}$  is the number of samples generated according to  $\mathbb{P}(\mathbf{y}|\theta)$ .

In contrast to Bayes rules, we will now apply *importance sampling* [15] to further approximate the integral in (12). Notice that the optimization problem in (12) involves a maximization over  $\pi \in \mathcal{P}(\Theta)$ . If we used Monte-Carlo sampling instead, then we would need to fix  $\pi$ , and for each choice of  $\pi$  we would need to generate samples of  $\theta$  according to  $\pi$ ; finally, we would have to find the value of  $\pi$  for which the approximate Monte-Carlo sum is a maximum for a fixed choice of  $\delta$ . This is computationally tedious. An alternative approach, through Importance Sampling, is to fix a distribution  $s(\theta)$ , called a *proposal distribution* (to be decided by a user) from which we can sample values of  $\theta$ . We use these samples to approximate the integral in (12) by employing *self-normalized importance sampling* [15, Chapter 9, page 8] as

$$\begin{aligned} \min_{\delta \in \Delta} \max_{w_1, \dots, w_{M_{\theta}} \in \mathbb{R}} & \frac{1}{M_{\mathbf{y}}} \frac{1}{M_{\theta}} \sum_{j=1}^{M_{\mathbf{y}}} \sum_{i=1}^{M_{\theta}} L(\theta_i, \delta(\mathbf{y}_{ij})) w_i \\ \text{s.t.} & w_i \geq 0, \quad i = 1, \dots, M_{\theta} \\ & \sum_{i=1}^{M_{\theta}} w_i = 1, \end{aligned} \quad (13)$$

where  $\mathbf{y}_{ij} = \mathbf{y}_{j,\theta_i}$ ,  $w_i = \frac{\pi(\theta_i)/s(\theta_i)}{\sum_{k=1}^{M_{\theta}} \pi(\theta_k)/s(\theta_k)}$ , and  $\{\theta_i\} \stackrel{i.i.d.}{\sim} s(\theta)$ .

Furthermore, we can notice that the inner maximization problem takes place over a probability simplex, so the maximum is obtained by assigning probability mass 1 to  $\max_{i=1, \dots, M_\theta} L_i(\delta)$ . Therefore, by using the epigraph formulation of the optimization problem [16, page 134], one can show that (13) is equivalent to

$$\begin{aligned} \min_{\delta \in \Delta, t \in \mathbb{R}} \quad & t \\ \text{s.t.} \quad & L_i(\delta) \leq t, \quad i = 1, \dots, M_\theta. \end{aligned} \quad (14)$$

The form of the decision rule  $\delta$  can be taken, as in the Bayesian case, as  $\delta = g \circ h_N$  (cf. (8)). Due to its similarity to TS, we will call it the *minimax TS estimator*.

Finally, we will consider a very special case of the minimax TS estimator, where  $d = 1$ , the loss function is the *mean square error*  $L(\theta, \hat{\theta}) = (\theta - \hat{\theta})^2$ , where  $\hat{\theta}$  is an estimate of  $\theta$  and  $g$  is a linear function<sup>1</sup>; furthermore, we assume that  $\mathbf{y}$  is a vector of *i.i.d.* or stationarity entries, so by Remark 4.1 we take  $M_{\mathbf{y}} = 1$ . These restrictions will allow us to derive an explicit computation procedure for the second stage. In this case,  $L_i(\delta) = L(\theta_i, \delta(\mathbf{y}_{i1}))$ . Furthermore,  $\delta(\mathbf{y}_{i1}) = \beta^T \phi(\alpha_i)$ , where the vector  $\alpha_i \in \mathbb{R}^n$  contains the compressed data obtained in the first stage of TS,  $\phi(\cdot)$  is a fixed non-linear function  $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^m$  with  $n < m < N$  and  $\beta \in \mathbb{R}^n$  is the linear regression coefficient vector that describes the linear function in the second stage of TS. Then, we obtain the following optimization problem

$$\begin{aligned} \min_{\beta \in \mathbb{R}^n, t \in \mathbb{R}} \quad & t \\ \text{s.t.} \quad & \beta^T (\phi(\alpha_i) \phi(\alpha_i)^T) \beta - 2\theta_i \phi(\alpha_i)^T \beta + \theta_i^2 \leq t, \\ & i = 1, \dots, M_\theta. \end{aligned} \quad (15)$$

This optimization problem can be solved using a standard convex programming solver, such as CVXPY [17].

## V. TWO-STAGE APPROACH FOR I.I.D. DATA

The formulations of TS given in Section IV make no special assumptions on the distribution of the data  $\mathbf{y}$ . In this section, we assume that  $\mathbf{y}$  is a vector whose entries,  $\{y^i\}$ , are *i.i.d.* according to distribution  $f(\theta)$ , where  $f(\theta)$  is a *pdf* parameterized by  $\theta \in \Theta$ .

To properly account for the *i.i.d.* assumption, we notice that under it, reasonable decision rules are “permutation-invariant” [7], that is, their output should not vary if the entries of  $\mathbf{y}$  are permuted. Such decision rules can be shown [7] to be functions of the *order statistics* of  $y^1, \dots, y^N$ , which correspond to the values of  $y^1, \dots, y^N$  sorted in ascending order, and are denoted by  $y^{(1)}, \dots, y^{(N)}$ .

We suggest to compress the data starting from the order statistics. To this end, we will use some *quantiles* of the order statistics: If  $0 < p < 1$ , then the  $(100p)$ -th sample quantile is a linear interpolation between  $y^{(Np)}$  and  $y^{(N(1-p))}$ . The computation of a small and fixed number of quantiles as the compressed data constitutes the first stage of TS, and the description of the first stage is given below.

<sup>1</sup>This special case can be easily extended to more than one parameter, *i.e.*,  $d > 1$ , by applying TS to estimate each parameter separately.

## A. First Stage

For the first stage, we set  $p = k/n$ , where  $k = 1, \dots, n$ , and obtain  $n$  quantiles. We then collect those quantiles into the compressed  $n$ -dimensional real-valued column vector  $\alpha = (\alpha_1, \dots, \alpha_n)^T$ . The first stage for both Bayes TS and minimax TS is the same.

## B. Second stage

For both formulations, the second stage starts by constructing non-linear features  $\phi(\alpha)$  using the compressed data  $\alpha$  and a fixed non-linear function  $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^m$  with  $n < m < N$ . The function  $\phi$  must be chosen by the user, and it depends on the type of parameter(s) that we want to estimate. In both formulations, we then design  $g$  as a linear function in composition with  $\phi$ , *i.e.*, as  $g(\alpha) = \beta^T \phi(\alpha)$ , where  $\beta$  is determined as explained in Section IV. If  $\theta = (\theta_1, \dots, \theta_d)^T$ , where  $d > 1$ , is the parameter vector to be estimated, then we can design a separate estimator  $g_i(\alpha) = \beta_i^T \alpha$  for each  $i = 1, \dots, d$ . To avoid numerical issues, in both Bayes and minimax formulations we add a regularization term of the form  $\lambda \|\beta\|_2^2$  to the optimization problem to be solved, where  $\lambda > 0$  is a very small number (*e.g.*,  $10^{-8}$ ).

*Remark 5.1:* Both  $n$  and  $\phi(\cdot)$  are user choices, which need to be selected according to the specific estimation problem.

## VI. SIMULATIONS

In this section, we illustrate the performance of the Bayes and minimax estimators derived in Section IV on the estimation of the two parameters of a Weibull distribution. Weibull distributions are widely used in reliability engineering [18, 19] to model the probability of failure of an equipment at a particular age, and are parameterized by: (i) a *scale* parameter, denoted by  $\eta > 0$ , and (ii) a *shape* parameter, denoted by  $\gamma > 0$ . Here,  $\eta$  and  $\gamma$  are unknown. Hence, this is a two dimensional parameter estimation problem, with  $\theta = (\eta, \gamma)^T$  being the unknown parameter. As described in Section V-B, we design separate estimators for  $\eta$  and  $\gamma$ , which are denoted by  $\hat{\eta}$  and  $\hat{\gamma}$  respectively.

The Weibull probability density function is

$$f(A) = \frac{\gamma}{\eta} \left( \frac{A}{\eta} \right)^{\gamma-1} \exp \left[ - \left( \frac{A}{\eta} \right)^\gamma \right], \quad A \geq 0. \quad (16)$$

We follow the TS approach described in Section V, where for the first stage we consider  $n = 10$  quantiles, and for the second stage, we choose  $\phi(\alpha)$  according to the following:

- For scale parameter  $\eta$ ,

$$\phi_i(\alpha) = \begin{cases} \alpha_i, & \text{if } 1 \leq i \leq n \\ \frac{\alpha_{i-n+1}}{\alpha_1}, & \text{if } n+1 \leq i \leq 2n-1. \end{cases} \quad (17)$$

The reason for choosing these combinations of quantiles to estimate  $\eta$  is that quotients of the (theoretical) quantiles of a distribution are proportional to its scale parameter (when the other parameters are kept fixed), which suggests that the quantities  $\phi_i(\alpha)$  above can be used for estimating  $\eta$ .

- For shape parameter  $\gamma$ , let

$$\psi_j(\alpha) = \begin{cases} \alpha_j, & \text{if } 1 \leq j \leq n \\ \frac{\alpha_{j-n}}{\alpha_n}, & \text{if } n+1 \leq j \leq 2n-1. \end{cases} \quad (18)$$

We let  $\phi(\alpha)$  consist of monomials of the  $\psi_j(\alpha)$ 's up to order 2. This choice is suggested by the fact that a small increase in  $\gamma$  (for  $\gamma > 1$ ) leads to a similar change of the Weibull distribution as a small decrease in  $\eta$ .

We should remark that for the Weibull distribution there exist exact expressions for  $\eta$  and  $\gamma$  in terms of any two quantiles [20]. However, using such expressions together with sample quantiles may lead to sub-optimal estimates, since they ignore the variability of the sample quantiles.

We will now look at results corresponding to both Bayes and minimax frameworks that we discussed in Section IV.

#### A. Bayes estimator

To illustrate the Bayes TS approach of Section IV-A, we study the effect of choosing different priors for the parameter  $\theta$ . In particular, two choices of prior are considered:

- 1) Independent uniform priors  $\mathcal{U}[1, 20]$  for  $\eta$  and  $\gamma$ .
- 2) Independent *uninformative* priors for  $\eta$  and  $\gamma$  with density [21]

$$f(x) = \begin{cases} \frac{c}{x}, & \text{if } a \leq x \leq b \\ 0, & \text{otherwise,} \end{cases} \quad (19)$$

where  $c = 1/\ln(b/a)$ . Here, we take  $a = 1$  and  $b = 20$  for both  $\eta$  and  $\gamma$ .

For each of the above priors, we find the Bayes TS estimator by drawing  $M_\theta = 1000$  values of  $\theta = (\eta, \gamma)^T$  (i.e., 1000 values of  $\eta$  and  $\gamma$ ) independently. For each sample of  $\theta$ ,  $N = 10000$  observations  $y^1, \dots, y^N$  are drawn according to the Weibull distribution parameterized by  $\theta$ . Once an estimator is found for each of the above priors, we evaluate its performance by drawing  $M_{\text{test}} = 1000$  values of  $\theta = (\eta, \gamma)^T$  (different from those used to find the estimator), and for each sample of  $\theta$ ,  $N = 10000$  observations are drawn according to the Weibull distribution parameterized by  $\theta$ .

We show a scatter plot of estimated values of  $\eta$  and  $\gamma$  vs. their true values in Figures 1 - 4. Figures 1 and 2 correspond to uniform prior, whereas Figures 3 and 4 correspond to the uninformative prior given by (19).

#### B. Minimax estimator

To illustrate the minimax TS approach of Section IV-B, we find it by drawing  $M_\theta = 1000$  values of  $\theta = (\eta, \gamma)^T$  independently from a proposal uniform  $\mathcal{U}[1, 20]$  distribution. For each sample of  $\theta$ ,  $N = 10000$  observations  $y^1, \dots, y^N$  are drawn according to the Weibull distribution parameterized by  $\theta$ . Once the estimator is found, we evaluate its performance by drawing  $M_{\text{test}} = 1000$  new values of  $\theta = (\eta, \gamma)^T$ , and for each sample of  $\theta$ ,  $N = 10000$  observations are drawn according to the Weibull distribution parameterized by  $\theta$ . We show the scatter plot of estimated values of  $\eta$  and  $\gamma$  vs. their true values in Figures 5 and 6.




Fig. 1: Scatter plot (blue dots) of Bayes TS estimates of  $\eta$  vs. its true value for a uniform prior. The red dashed line corresponds to an oracle estimate, which knows the true value of the parameter.




Fig. 2: Scatter plot (blue dots) of Bayes TS estimates of  $\gamma$  vs. its true value for a uniform prior. The red dashed line corresponds to an oracle estimate, which knows the true value of the parameter.

To assess the quality of the Bayes and minimax TS estimators of  $\eta$  and  $\gamma$ , we have numerically evaluated the *mean square error* (MSE) of the TS estimators, based on 1000 Monte Carlo simulations, and compared them with a numerical evaluation of their corresponding Cramér-Rao lower bounds (CRLB) [22]. The results are shown in Table I.

It is evident from Figures 1 - 6 that the estimates given by both Bayes and minimax TS estimators of the scale ( $\eta$ ) and shape ( $\gamma$ ) parameters are close to the true values. However, according to Table I, the TS estimates of  $\eta$  are fairly efficient (in the sense that their MSE is close to CRLB), while those of  $\gamma$  are not as reliable. The reason for the latter lies in the choice of the function  $\phi$  for  $\gamma$ . Also, one can see from Table I the effect of changing priors in the Bayes estimator: For the *uninformative* priors, estimates of  $\eta$  and  $\gamma$  are relatively much closer to true values than for the uniform priors. It is also interesting to note that, while no prior is needed for the minimax estimator, one can still get reliable estimates, as demonstrated in Figures 5 and 6. By more careful “feature engineering” of  $\phi$  in the second stage for  $\hat{\gamma}$ , we could obtain more reliable estimates of the shape parameter without using more complex functional approximators such as Deep Neural Networks. How to perform such feature engineering, based on a theoretical analysis of our novel TS estimators, is left for future research.

## VII. CONCLUSIONS

In this paper, we provided a decision-theoretical framework for the Two-Stage paradigm to estimate the parameters of a data generating model. In particular, we have derived

True Values		CRLB		MSE, Bayesian (Uniform Prior)		MSE, Bayesian (Uninformative Prior)		MSE, Minimax	
$\eta$	$\gamma$	$\hat{\eta}$	$\hat{\gamma}$	$\hat{\eta}$	$\hat{\gamma}$	$\hat{\eta}$	$\hat{\gamma}$	$\hat{\eta}$	$\hat{\gamma}$
2	2	$1.11 \times 10^{-4}$	$2.43 \times 10^{-4}$	$2.58 \times 10^{-4}$	$5.77 \times 10^{-2}$	$1.42 \times 10^{-4}$	$2.06 \times 10^{-2}$	$2.17 \times 10^{-4}$	$16 \times 10^{-2}$
2	8	$6.93 \times 10^{-6}$	$3.89 \times 10^{-3}$	$1.11 \times 10^{-5}$	$5.61 \times 10^{-2}$	$1.27 \times 10^{-5}$	$4.44 \times 10^{-2}$	$4.28 \times 10^{-4}$	$13.29 \times 10^{-2}$
4	2	$4.43 \times 10^{-4}$	$2.43 \times 10^{-4}$	$6.74 \times 10^{-4}$	$1.05 \times 10^{-1}$	$6.07 \times 10^{-4}$	$2.43 \times 10^{-2}$	$8.38 \times 10^{-4}$	$14.67 \times 10^{-2}$
4	8	$2.77 \times 10^{-5}$	$3.89 \times 10^{-3}$	$3.84 \times 10^{-5}$	$6.40 \times 10^{-2}$	$4.33 \times 10^{-5}$	$3.96 \times 10^{-2}$	$1.72 \times 10^{-3}$	$8.59 \times 10^{-2}$
8	2	$1.77 \times 10^{-3}$	$2.43 \times 10^{-4}$	$2.26 \times 10^{-3}$	$1.89 \times 10^{-1}$	$2.27 \times 10^{-3}$	$2.35 \times 10^{-2}$	$3.307 \times 10^{-3}$	$18.605 \times 10^{-2}$
8	8	$1.11 \times 10^{-4}$	$3.89 \times 10^{-3}$	$1.58 \times 10^{-4}$	$7.901 \times 10^{-2}$	$1.76 \times 10^{-4}$	$4.51 \times 10^{-2}$	$6.77 \times 10^{-3}$	$8.39 \times 10^{-2}$

TABLE I: MSE of Bayes and minimax TS estimators of the scale and shape parameters, and their corresponding CRLBs.




Fig. 3: Scatter plot (blue dots) of Bayes TS estimates of  $\eta$  vs. its true value for an uninformative prior. The red dashed line corresponds to an oracle estimate, which knows the true value of the parameter.




Fig. 5: Scatter plot (blue dots) of minimax TS estimates of  $\eta$  vs. its true value. The red dashed line corresponds to an oracle estimate, which knows the true value of the parameter.




Fig. 4: Scatter plot (blue dots) of Bayes TS estimates of  $\gamma$  vs. its true value for an uninformative prior. The red dashed line corresponds to an oracle estimate, which knows the true value of the parameter.




Fig. 6: Scatter plot (blue dots) of minimax TS estimates of  $\gamma$  vs. its true value. The red dashed line corresponds to an oracle estimate, which knows the true value of the parameter.

Bayes and minimax formulations of TS, and showed how to implement them for *i.i.d.* data, by considering nonlinear functions of the sample quantiles of the data as a first stage, and then optimizing for a linear function as the second stage. We have also evaluated the performance of these novel estimators on simulated data from a Weibull distribution.

#### REFERENCES

- [1] G. Casella and R. Berger, *Statistical Inference, 2nd Ed.* Duxbury, 2002.
- [2] D. Massart, B. Vandeginste, S. Deming, Y. Michotte, and L. Kaufman, *Chemometrics: A textbook.* Elsevier, 1988.
- [3] C. Gouriéroux and A. Monfort, *Statistics and Econometric Models.* Cambridge University Press, 1996.
- [4] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms.* Cambridge University Press, 2012.
- [5] L. Ljung, *System Identification: Theory for the User, 2nd Edition.* Prentice Hall, 1999.
- [6] T. Söderström and P. Stoica, *System Identification.* Prentice Hall, 1989.
- [7] E. Lehmann and G. Casella, *Theory of Point Estimation, 2nd Ed.* Springer-Verlag, 1998.
- [8] T. Söderström and P. Stoica, *Instrumental Variable Methods for System Identification.* Springer-Verlag, 1983.
- [9] C. Gouriéroux and A. Monfort, *Simulation-Based Econometric Methods.* Oxford University Press, 1997.
- [10] S. Garatti and S. Bittanti, "A new paradigm for parameter estimation in system modeling," *International Journal of Adaptive Control and Signal Processing*, vol. 27, no. 8, pp. 667–687, 2013.
- [11] S. Garatti and S. Bittanti, "Estimation of white-box model parameters via artificial data generation: a two-stage approach," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 11409–11414, 2008.
- [12] M. Ledoux, *The Concentration of Measure Phenomenon.* AMS, 2001.
- [13] L. L. Cam, "Sufficiency and approximate sufficiency," *The Annals of Mathematical Statistics*, vol. 35, no. 4, pp. 1419–1455, 1964.
- [14] C. Gouriéroux, A. Monfort, and E. Renault, "Indirect inference," *Journal of Applied Econometrics*, vol. 8, no. S, pp. S85–118, 1993.
- [15] A. B. Owen, *Monte Carlo theory, methods and examples.* in press 2013. <https://artowen.su.domains/mc/>.
- [16] S. Boyd and L. Vandenberghe, *Convex Optimization.* Cambridge University Press, 2004.
- [17] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.
- [18] P. R. Tadikamalla, "Age replacement policies for Weibull failure times," *IEEE Transactions on Reliability*, vol. R-29, no. 1, pp. 88–90, 1980.
- [19] W. Li, "Incorporating aging failures in power system reliability evaluation," *IEEE Transactions on Power Systems*, vol. 17, no. 3, pp. 918–923, 2002.
- [20] J. D. Cook, "Determining distribution parameters from quantiles," [https://www.johndcook.com/quantiles\\_parameters.pdf](https://www.johndcook.com/quantiles_parameters.pdf), January 2010.
- [21] D. Sun, "A note on noninformative priors for Weibull distributions," *Journal of Statistical Planning and Inference*, vol. 61, no. 2, pp. 319–338, 1997.
- [22] V. K. Murthy and G. B. Swartz, "Estimation of Weibull parameters from two-order statistics," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 37, no. 1, pp. 96–102, 1975.