Reliability Engineering and System Safety 131 (2014) 118–125

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Adversarial life testing: A Bayesian negotiation model

M.J. Rufo *, J. Martín, C.J. Pérez

Departamento de Matemáticas, Escuela Politécnica, Universidad de Extremadura, Avda. de la Universidad s/n, 10003 Cáceres, Spain

ARTICLE INFO

Article history:
Received 20 January 2014
Received in revised form
3 June 2014
Accepted 19 June 2014
Available online 8 July 2014

Keywords:
Bayesian analysis
Exponential families
Life testing
Reliability
Negotiation model

ABSTRACT

Life testing is a procedure intended for facilitating the process of making decisions in the context of industrial reliability. On the other hand, negotiation is a process of making joint decisions that has one of its main foundations in decision theory. A Bayesian sequential model of negotiation in the context of adversarial life testing is proposed. This model considers a general setting for which a manufacturer offers a product batch to a consumer. It is assumed that the reliability of the product is measured in terms of its lifetime. Furthermore, both the manufacturer and the consumer have to use their own information with respect to the quality of the product. Under these assumptions, two situations can be analyzed. For both of them, the main aim is to accept or reject the product batch based on the product reliability. This topic is related to a reliability demonstration problem. The procedure is applied to a class of distributions that belong to the exponential family. Thus, a unified framework addressing the main topics in the considered Bayesian model is presented. An illustrative example shows that the proposed technique can be easily applied in practice.

1. Introduction

Reliability demonstration of a system can be viewed as a Bayesian decision problem in which the acceptance of the system with respect to its reliability is analyzed (see [28]). Then, the principles of Bayesian decision theory can be applied (see, for instance, [3]). Papazoglou [28] points out that the decision about accepting or rejecting the system can be made on the basis of the existing prior assessment. However, the decision maker can be not sure about making an appropriate decision (decision under uncertainty). Therefore, additional information can be obtained through testing of the components or the system at a cost. The results can be combined with the existing information to provide updated information for the system reliability. In addition, as each test has a cost, then an added decision is whether to obtain further information. Some approaches related to decision analysis in engineering contexts are provided by Aven and Korte [2], Brito and de Almeida [6] and Pasanisi et al. [29].

Raiffa [32] introduced the term negotiation analysis, or joint decision making, which has its main roots in game theory and decision analysis (see also [35,36]). Negotiation can be defined as a process by which two or more parties try to reach compromises and to come to an agreement. Murtoaro and Kujala [24] pointed out in their work that negotiations occur in all domains of life. However, the basic structure of negotiations applied to different contexts is basically the same. Thus, any negotiation process has the following common characteristics (see [33]): first, there are two or more parties involved in the entire procedure; second, the payoffs for each party depend either on the consequences of the joint decisions or on alternatives external to the negotiations; third, the parties can reciprocally and directly exchange information (this communication can be honest or not); fourth, the parties can be creative in the decisions they make in order to arrive at a joint decision. See Tsay and Bazerman [38] and Agarwal [1] for more current reviews on negotiation analysis. Nowadays, there exists a wide literature that applies negotiation to different engineering fields. Examples of this statement are the papers by Yu et al. [39], Zhang et al. [40], Ethamo et al. [9] and McCalley et al. [23], among others.

Much of the literature on reliability and survival analysis deals with the topics of life testing and the analysis of failure data. Furthermore, in most cases involving industrial settings, life-testing is performed with the aim of making decisions (see, e.g., [37]). Lindley and Singpurwalla [19,20] proposed a Bayesian framework which involved two adversarial decision makers: the manufacturer and the consumer. The scenario in which the consumer demands a batch of items from the manufacturer is considered. The consumer could either accept or reject the batch provided by the manufacturer. In the last case, the two main issues brought up are: (i) whether the manufacturer should offer a sample to the consumer for testing, and (ii) how large it should be. Lindley and Singpurwalla [19] discuss the problem above in the context of acceptance sampling for quality control by using

* Corresponding author. Tel.: +34 927257220; fax: +34 927257203.
E-mail addresses: mrufo@unex.es (M.J. Rufo), jrmartin@unex.es (J. Martín), carper@unex.es (C.J. Pérez).

http://dx.doi.org/10.1016/j.ress.2014.06.007
0951-8320/

Bernoulli, Poisson and Normal distributions, whereas Lindley and Singpurwalla [20] consider exponential lifetimes in a reliability context.

In this paper, the procedure presented by Lindley and Singpurwalla [20] is extended and modified by considering a Bayesian sequential negotiation model in the context of life testing. Through the paper it is assumed that the manufacturer offers a product batch to the consumer. The decision of accepting or rejecting the product batch is based on the reliability of the product, which is represented by its lifetime. It is considered that the parametric lifetime distributions belong to a subclass of the exponential family. This family has been considered because it includes distributions widely used in practice and, especially in the reliability engineering field. Thus, a unified framework is presented. Different prior distributions and utilities for the manufacturer and the consumer are considered. The utilities are initially based on an observable lifetime. A Monte Carlo simulation-based approach is implemented to calculate the optimal sample size n that the manufacturer will offer the consumer. In most cases it is possible to find analytical expressions for the expectations involved in the process. Alternative approaches are proposed when these expressions can not be obtained analytically. This leads to an easy proposal to apply.

The outline of the paper is as follows. In Section 2, the model is described. Section 3 applies the developed methodology by considering a set of distributions coming from the exponential family. Hence, a unified framework is provided. Moreover, a straightforward approach is proposed in order to obtain an appropriate sample size. In Section 4, an application by considering the Weibull distribution is shown. The conclusion is presented in Section 5. Finally, an Appendix contains the theoretical developments.

2. Model description

Suppose that two parties (manufacturer and consumer) are negotiating the sale of a product based on its lifetime. Both the manufacturer and the consumer are involved in a negotiation process in which they are expected to come to an agreement. Initially, the manufacturer offers a product to the consumer. It is assumed that a batch composed by K units of this product is characterized by the unknown lifetimes (denoted by $X_i, i = 1, 2, \ldots, K$). In addition, given a parameter θ, it is considered that the random variables, X_i, are independent and identically distributed with distribution depending on the parameter θ. Note that this parameter is directly related to product characteristics like the product quality. This remark will be addressed further in the next section.

From the consumer viewpoint, she/he might accept or reject K units of the product based on her/his prior knowledge about θ, $\pi_C(\theta)$, and on her/his preferences represented by a utility function $\mathcal{U}_C(\cdot, \theta)$. Thus, the consumer will accept the product batch when it is satisfied that the prior expected utility of accepting is greater than or equal to the prior expected utility of rejecting, that is

$$\begin{aligned} E_{\pi_C(\theta)}[\mathcal{U}_C(A,\theta)] &= \int_\Theta \mathcal{U}_C(A,\theta)\pi_C(\theta)\, d\theta \\ &\geq \int_\Theta \mathcal{U}_C(R,\theta)\pi_C(\theta)\, d\theta \\ &= E_{\pi_C(\theta)}[\mathcal{U}_C(R,\theta)], \end{aligned} \tag{1}$$

where A and R denote the decisions for accepting and rejecting, respectively. If the inequality (1) is satisfied, then the process is finished.

If the consumer had rejected the product batch, then the manufacturer would have, at least, two options. The first one consists of offering her/him a lower price. The second option is to provide the consumer a sample of size n in order to carry out life testing. Therefore, she/he can update her/his beliefs about the product reliability. In any case, the manufacturer hopes to convince her/him about the product quality. Observe that this situation is connected with the considered one by Papazoglou [28] for reliability demonstration. This work is only focused on the second choice since both the consumer and the manufacturer are interested in the product reliability (lifetime). In this case, the consumer can modify the information about the parameter θ by considering the data sample, $\boldsymbol{x}$, combined with the prior distribution to obtain the posterior distribution $\pi_C(\theta|\boldsymbol{x}, n)$. Thus, the consumer will accept the batch when the following inequality is satisfied:

$$\begin{aligned} \mathcal{U}_C(A,\boldsymbol{x},n) &= \int_\Theta \mathcal{U}_C(A,\theta)\pi_C(\theta|\boldsymbol{x},n)\, d\theta \\ &\geq \int_\Theta \mathcal{U}_C(R,\theta)\pi_C(\theta|\boldsymbol{x},n)\, d\theta = \mathcal{U}_C(R,\boldsymbol{x},n). \end{aligned}$$

Note that $\mathcal{U}_C(A,\boldsymbol{x},n)$ and $\mathcal{U}_C(R,\boldsymbol{x},n)$ denote the expected utilities $E_{\pi_C(\theta|\boldsymbol{x},n)}[\mathcal{U}_C(A,\theta)]$ and $E_{\pi_C(\theta|\boldsymbol{x},n)}[\mathcal{U}_C(R,\theta)]$, respectively. In addition, the previous notation is used for the expected utilities since they depend on the parameter through the posterior distribution.

The previous problem is directly associated with the one from the manufacturer viewpoint, who has to decide an optimal sample size n. Specifically, the manufacturer decides n by calculating

$$\max_n \left\{ \int_{\boldsymbol{x}} (\mathcal{U}_M(A,\boldsymbol{x},n) I_{\mathcal{A}(n)}(\boldsymbol{x}) + \mathcal{U}_M(R,\boldsymbol{x},n) I_{\mathcal{R}(n)}(\boldsymbol{x}))\pi_M(\boldsymbol{x}|n)\, d\boldsymbol{x} \right\}, \tag{2}$$

where $\pi_M(\boldsymbol{x}|n)$ denotes the prior predictive distribution for the manufacturer, the set $\mathcal{A}(n)$ is given by

$$\mathcal{A}(n) = \{\boldsymbol{x} : \mathcal{U}_C(A,\boldsymbol{x},n) \geq \mathcal{U}_C(R,\boldsymbol{x},n)\}, \tag{3}$$

and $I_{\mathcal{A}(n)}(\boldsymbol{x})$ is the indicator function. Analogously, the expression for $\mathcal{R}(n)$ is

$$\mathcal{R}(n) = \{\boldsymbol{x} : \mathcal{U}_C(A,\boldsymbol{x},n) < \mathcal{U}_C(R,\boldsymbol{x},n)\}. \tag{4}$$

Finally, $\mathcal{U}_M(A,\boldsymbol{x},n)$ and $\mathcal{U}_M(R,\boldsymbol{x},n)$ denote the expected utilities for the manufacturer. They are obtained as the previous expected utilities for the consumer $\mathcal{U}_C(A,\boldsymbol{x},n)$ and $\mathcal{U}_C(R,\boldsymbol{x},n)$.

Note that, from the manufacturer viewpoint, the maximization of the expected utility with respect to her/his predictive distribution is involved.

Furthermore, observe that the optimal value n can be achieved at 0. This means that the consumer's preferences and the prior beliefs (prior distribution) are such that, under the manufacturer's judgements, no sample size can convince the consumer of the product reliability (see [37]).

The second part corresponding to the sequential negotiation process can be adequately represented by using an influence diagram, such as it is shown in Fig. 1. Influence diagrams are powerful graphical tools for dealing with Bayesian decision problems under uncertainty (see, e.g., [4,25]). They are acyclic directed graphs with three types of nodes and two types of arcs. Decision nodes, which are represented by a square or rectangle, chance or uncertainty nodes that are symbolized by a circle or an ellipse and value nodes that are represented as diamonds or hexagons. In relation to the arcs, they can be conditional arcs which are directed towards a chance or a value node and informational arcs that are directed toward a decision node. Arcs into a value or uncertainty node indicate functional and probabilistic dependence. Finally, arcs into a decision node indicate that when the decision is made, the values of the preceding nodes are known. A deeper description on influence diagrams can be found in Howard and Matheson [15] and Fikret and Senay [10].

For the influence diagram in Fig. 1, the manufacturer makes the decision to offer a random sample to the consumer which will