



The Institute for Integrating Statistics in Decision Sciences

Technical Report TR-2017-4

Advances in Bayesian Games and Decisions in Reliability

David Ríos Insua
ICMAT, Spain

Fabrizio Ruggeri
CNR-IMATI, Italy

Refik Soyer
*Department of Decision Sciences
The George Washington University, USA*

Simon P. Wilson
*School of Computer Science and Statistics
Trinity College, Ireland*

Advances in Bayesian Games and Decisions in Reliability

David-Rios Insua

ICMAT, Madrid, Spain

Fabrizio Ruggeri

CNR-IMATI, Milano, Italy

Refik Soyer

George Washington University, Washington, DC, USA

Simon Wilson

Trinity College, Dublin, Ireland

December 22, 2017

Abstract

Starting in the late 80's Bayesian methods have gained increasing attention in the reliability literature. The focus of most of the earlier Bayesian work in reliability involved statistical inference and thus the main focus was on modeling and analysis. Advances in Bayesian computing after the 90's have significantly contributed not only to the use of Bayesian inference but also to the implementation of Bayesian decision-theoretic approaches in reliability problems. In this review paper we present an overview of Bayesian methods to solve decision problems in reliability some of which involve two or more decision makers with conflicting objectives. We consider problems in areas such as design, life testing, preventive maintenance, reliability certification, warranty policies and others. In doing so, we present key aspects of the decision problems, give a brief review of earlier methods and finally discuss recent advances in Bayesian approaches to solve them.

Keywords: Optimal replacement, design of life tests, stopping rules, warranties, adversarial risk analysis, Markov chain Monte Carlo.

1 Introduction and Overview

In his review paper Barlow (1984) noted that "The mathematical theory of reliability has grown out of the demands of modern technology and particularly out of the experiences in World War II with complex military systems". Some of this earlier work in reliability was in the area of machine maintenance which involved renewal theory applications to replacement problems as in Lotka (1939). At about the same time, in 1944, von Neumann and Morgenstern published their celebrated *Theory of Games and Economic Behavior* which contained the axiomatization of utilities and laid down the foundations of decision theory and analysis. Abraham Wald (1950) was the first statistician to recognize the connection between game theory and the statistical theory of hypothesis testing. In 1954, the publication of Leonard Savage's *The Foundations of Statistics* completed the foundations of modern decision theory and Bayesian statistics developing an axiomatic basis of probability based on behavioral considerations in combination with utility theory.

As pointed out by Singpurwalla (2009), "Reliability is a key ingredient for making decisions that mitigate the risk of failure. The other key ingredient is utility". The Bayesian decision theoretic approach integrates both by quantifying uncertainty via probability and preferences via utility, providing a coherent framework for making decisions by subscribing to the principle of expected utility maximisation, see Lindley (1985) and French and Rios Insua (2000). This is well recognized by many authors in the Operations Research and reliability communities like Percy (2002), who points out that the limited availability of data in many reliability related decision problems makes the Bayesian approach inevitable. He mentions preventive maintenance and repair/replacement strategies, condition monitoring of systems and specification of warranty policies, etc. as example areas.

Reliability problems involving a single decision maker can be formulated and solved with the Bayesian decision-theoretic framework. However, as noted by Rios Insua et al. (2017), there are reliability problems that involve two or more decision makers, possibly with opposed interests. One of the earliest Bayesian approaches in adversarial reliability settings is due to Lindley and Singpurwalla (1991, 1993) who considered problems in acceptance sampling and life testing. Such problems involving adversarial components can be framed as games and are typically solved using (non cooperative) game-theoretic methods. Hausken (2002) provides examples of using the game-theoretic approach in system reliability analysis especially in relation with infrastructure reliability of public systems. Other adversarial application areas include warranty analysis, software testing, optimal maintenance and reliability demonstration.

Barlow (1984) noted that "Among statisticians working in reliability theory, perhaps the most significant trend is the growing recognition of the usefulness of the Bayesian approach to inductive inference." Since then, there has been a considerable increase in the use of Bayesian methods in reliability problems involving inference but also decision making. The objective of this paper is to focus on the latter and provide a review of recent advances in Bayesian methods for decision making in reliability, including Bayesian decision analysis methods to cope with adversarial issues in reliability.

The structure of our paper is as follows. The material in Section 2 focuses on Bayesian design for life tests including censored life testing experiments. Extensions such as accelerated life testing are considered. The main aspects of decision problems and the associated computational issues are discussed. Related problems of optimal stopping and planning of reliability demonstration tests are also presented. Section 3 refers to Bayesian maintenance policies for both repairable and non repairable systems. Different replacement protocols are presented and computation of optimal replacement intervals are discussed. Recent work on nonparametric Bayesian policies that are based on advances in Bayesian computing are also highlighted. The final part of the section considers sequential maintenance problems and outlines implementation of semi-Markov decision processes. The focus of Section 4 is on adversarial problems in reliability. Bayesian solutions are discussed for game theory and adversarial risk analysis setups for acceptance sampling plans and Bayesian methods for specifying

warranty policies are presented. Concluding remarks follow in Section 5.

2 Design of Life Testing Experiments

In reliability analysis experiments are conducted to obtain information about failure characteristics of systems or components of interest, to assess reliability and, if necessary, improve system performance based on information from the experiment. Such experiments are generally referred to as *life tests*. For example, we may be interested in learning about the life length of a component under a certain environment, its reliability at a given mission time, or the failure rate, and make a decision on whether to change the design of the component based on such information. Design problems in life testing involve the determination of one or more elements of the life test such as the testing environment, the number of items to be tested or the stopping rule. Optimal selection of the design variables require consideration of benefits from the life test as well as the associated costs with performing the test.

As noted by Polson and Soyer (2017), the Bayesian decision-theoretic approach to the optimal design problem requires the specification of three components:

1. A utility (loss) function: reflecting the consequences of selecting a specific design
2. A probability model: life distribution of the items in question
3. A prior distribution: reflecting a priori beliefs about all unknown quantities.

Let $a \in A$ denote the design variables, which may represent the number of items to be tested, the duration of testing, or binary actions such as stop/continue testing. Let the observed outcome of a life test be x and the probability model for x given parameter(s) θ be $p(x|\theta, a)$. We assume that the uncertainty about θ prior to the life test is described by a probability distribution $p(\theta)$. Finally, we denote the utility function associated with the consequences of selecting a specific design a as $u(x, \theta, a)$ which will, therefore, depend on both x and θ in general. Then, the Bayesian solution to the design problem is obtained by maximizing the expected utility, given by

$$E[u(x, \theta, a)] = \bar{u}(a) = \int \int u(x, \theta, a) p(x, \theta|a) d\theta dx, \quad (1)$$

with respect to the design variable a . Since (1) can be written as

$$\bar{u}(a) = \int \int u(x, \theta, a) p(\theta|x, a) p(x|a) d\theta dx, \quad (2)$$

where $p(\theta|x, a)$ is the posterior distribution of θ , $\bar{u}(a)$ is usually referred to as the *pre-posterior expected utility*. The optimal design is then obtained by solving

$$a^* = \arg \max_a \bar{u}(a). \quad (3)$$

Life test data are typically censored and provide partial failure/survival information. The most commonly used strategies are *failure truncated (Type II)* and *time truncated (Type I)* censoring. Under the Type II scenario, n items are tested until $k \leq n$ of them fail. The testing duration, which is the time of k -th failure, is an unknown quantity whereas n and k are pre-specified design variables. In the time truncated censoring, n items are put on test for a pre-specified τ units of time, that is, testing stops after τ . In this case test duration is known, but the number of items that will fail during τ units of time is an unknown quantity. The design variables are $a = (k, n)$ and $a = (\tau, n)$ in the Type II and Type I censoring, respectively. It is also common to accelerate life-tests by changing the

testing environment to induce early failures and, in this case, the testing environment is also part of the design variables a . Sometimes life tests consist of several stages and are conducted sequentially until a criterion known as *stopping rule* is satisfied; see, for example, Deely and Keats (1994).

In what follows, we first present the Bayesian optimal design of censored life tests and discuss some of its extensions and recent Bayesian work. This is followed by a discussion of optimal accelerated life testing design. Finally, we present optimal stopping problems in life testing and sketch optimal release problems in reliability analysis.

2.1 Optimal Bayesian life test design

As mentioned, by providing failure/survival data, life tests allow us to infer the failure characteristics of the components in question. Earlier Bayesian work on design of life tests include Thyregod (1975) who considered Type II censoring where the test results would be used to accept/reject a production lot. Using a cost-based utility function, an optimal sampling plan was determined specifying the failure truncation and the *acceptance rule* for the lot. In related work Barnett (1972) developed a Bayesian sequential life test procedure using a posterior probability based criterion. An alternative Bayesian sequential procedure was considered by Bancroft and Dunsmore (1978) using predictive distribution of observed lifetimes.

The concept of *gain in information* about the failure characteristic from a life test, discussed by Brooks (1982) and Barlow and Hsiung (1983), plays an important role in optimal design of Bayesian life tests. A general utility function for the design problem is given by

$$u(x, \theta, a) = g(x, \theta, a) - c(x, a), \quad (4)$$

where functions g and c are similar to the *gain* and *cost* functions of Bernardo (1997), respectively. It is important to note that g and c have the same units. For example, if $g(x, \theta, a) = -V(\theta|x, a)$, that is, g is set to be the negative of the posterior variance of θ , then the evaluation of $\bar{u}(a)$ in (2) requires assessing of the pre-posterior variance of θ . If we choose the gain function as $g(x, \theta, a) = \log p(\theta|x, a)$, where $p(\theta|x, a)$ is the posterior distribution of θ under design a , then the evaluation of $\bar{u}(a)$ requires assessing *Lindley's measure* (1956). The utility function (4) provides an additive structure to reflect the benefits and costs associated with the life test design a . It is possible to reflect such trade-offs by using alternative forms as well, as in Section 2.2. Here, we consider additive utility functions.

2.1.1 Design of failure truncated life tests

As previously mentioned in the failure truncated censoring, the design variable to be determined is $a = (k, n)$. We assume a specific form for (4) given by

$$u(T(x, a)) = -V(\theta|T(x, a)) - c_t T(x, a), \quad (5)$$

where $T(x, a)$ is the *total time on test* (TTOT) under the failure truncated scenario and c_t represents the cost of testing. The TTOT is given by

$$T(x, a) = \sum_{i=1}^{k-1} X_{(i)} + (n - k + 1)X_{(k)}, \quad (6)$$

where $X_{(1)} < \dots < X_{(k)}$ are the first k order-statistics from the distribution $p(x|\theta, a)$.

Assume that the life-length of the items follow an exponential distribution with failure rate θ with

gamma prior distribution

$$p(\theta) = \frac{b^d}{\Gamma(d)} \theta^{d-1} e^{-b\theta}$$

denoted as $\theta \sim Gam(d, b)$. The standard Bayesian updating implies that $\theta|T(x, a), a \sim Gam(d + k, b + T(x, a))$ and $V(\theta|T(x, a), a) = (d + k)/(b + T(x, a))^2$. The evaluation of the preposterior utility requires the predictive distribution of $T(x, a)$ given by

$$p(T(x, a)) = \int_{\theta} p(T(x, a)|\theta)p(\theta)d\theta,$$

where $(T(x, a)|\theta) \sim Gam(k, \theta)$. It can be shown that

$$p(T(x, a)) = \frac{\Gamma(d + k)}{\Gamma(d)\Gamma(k)} \frac{(1/b)(T(x, a)/b)^{k-1}}{[1 + (T(x, a)/b)]^{d+k}} \quad (7)$$

corresponding to a scaled *inverted beta* or *beta prime* density; see Dunsmore (1974). To obtain the optimal design, we need to maximize the pre-posterior utility $\bar{u}(a)$. This is equivalent to minimizing

$$E_{T(x,a)}[V(\theta|T(x, a)) + c_t T(x, a)], \quad (8)$$

where the expectation is taken with respect to the predictive distribution (7). Note that in this setup, the only design variable is $a = k$. Using properties of the scaled inverted beta distribution we can obtain that the preposterior loss (8) is

$$\frac{d(d + 1)}{b^2(d + k)(d + k + 1)} + c_t \frac{bk}{(d - 1)}, \quad (9)$$

for $d > 1$. The expected loss (9) can be easily minimized with respect to k to find the optimal design.

2.1.2 Design of time truncated life tests

We will assume the same utility function (5) for the time truncated censoring case with design variable $a = (\tau, n)$. The TTOT is now given by

$$T(x, a) = \sum_{i=1}^K X_{(i)} + (n - K)\tau,$$

where $X_{(i)}$ is the i -th order statistic of truncated exponential random variables with density function

$$p(x|\theta, \tau) = \frac{\theta e^{-\theta x}}{1 - e^{-\theta\tau}}, \quad (10)$$

for $x < \tau$. In this case both the $X_{(i)}$'s and the number K of failures are random quantities. Given K , it can be shown that $\sum_{i=1}^K X_{(i)}$ is a sum of independent truncated exponential densities as in (10) and the distribution of K is binomial with parameter $[1 - e^{-\theta\tau}]$, that is,

$$p(k|\theta) = \binom{n}{k} [1 - e^{-\theta\tau}]^k e^{-\theta\tau(n-k)}. \quad (11)$$

The sampling distribution of $T(x, a)$ and K was obtained by Bartholomew (1963) as

$$p(t, k|\theta) = \binom{n}{k} \frac{\theta^k}{(k-1)!} e^{-\theta t \sum_{i=0}^k \binom{k}{i} (-1)^i [\max\{0, t - \tau(n-k+i)\}]^{k-1}} \quad (12)$$

for $k > 0$ and $p(t, 0|\theta) = e^{-\theta n \tau}$, for $k = 0$. Assuming a gamma prior for $\theta \sim \text{Gam}(d, b)$ we can obtain the predictive distribution of $T(x, a)$ and K as

$$p(t, k) = \binom{n}{k} \frac{\Gamma(d+k)}{\Gamma(d)\Gamma(k)} \frac{b^d}{[b + h(t, k)]^{d+k}}, \quad (13)$$

where $h(t, k) = t \sum_{i=0}^k \binom{k}{i} (-1)^i [\max\{0, t - \tau(n-k+i)\}]^{k-1}$.

As with Type II censoring, the utility function is given by

$$u(T(x, a)) = -\frac{d+K}{b+T(x, a)} - c_t T(x, a). \quad (14)$$

Note that evaluating the preposterior expected utility $\bar{u}(a)$ involves obtaining the expectation of (14) with respect to the predictive distribution (13). This cannot be obtained analytically, but $\bar{u}(a)$ can be approximated as a Monte Carlo (MC) average by drawing samples of (K, θ) from $p(k, \theta) = p(k|\theta)p(\theta)$, and using these to draw samples from K independent truncated exponential random variables from (10) to evaluate $T(x, a) = \sum_{i=1}^K X_i + (n-K)\tau$. Note that $T(x, a)$ is a function of (τ, n) . Once MC draws $(K^{(s)}, T(x, a)^{(s)})$ are available, we can find the optimal design $a^* = (\tau^*, n^*)$ by minimizing the MC average

$$\frac{1}{S} \sum_s \frac{d+K^{(s)}}{b+T^{(s)}(x, a)} + c_t T^{(s)}(x, a)$$

with respect to (τ, n) .

The Bayesian optimal design setup presented for the exponential model under both censoring scenarios can be extended to other failure models. For example, Zhang and Meeker (2005) considered life test plans for the Weibull model with known shape parameter under Type II censoring and presented optimal plans using different criteria. More recently, Kundu (2008) considered life test designs under progressive censoring, whereas Hong et al. (2015) developed Bayesian designs for the *log-location-scale family* of distributions using Markov Chain Monte Carlo (MCMC) methods to evaluate expectations. Polson and Soyer (2017) proposed the use of augmented probability simulation (APS) methods for design of Bayesian life tests and implemented their approach to accelerated life testing problems.

2.2 Design of accelerated life tests

Accelerated life tests (ALTs) involve testing systems in an environment that is more severe than the use environment and employing the data collected in the accelerated environment to infer failure behavior in the use environment. The design problem in accelerated life testing is concerned with specification of the number and magnitude of the accelerated stress levels, and the number of items to be tested at such stress levels; see for example, Soyer (2007).

Original work in ALT design is due to Chernoff (1962). There is a considerable literature on Bayesian ALT designs dating back to Martz and Waterman (1978) and DeGroot and Goel (1979). Later work include Chaloner and Larntz (1990), who developed Bayesian designs for Type I censored tests using optimality criterion proportional to the expected asymptotic variance of the failure characteristics of interest; Menzefricke (1992), who consider design of Type II censored ALTs for log-

normally distributed lifetimes; Verdinelli et al. (1993,) who present optimal designs maximizing Shannon information; and Soyer and Vopatek (1995), who introduce *linear Bayesian* designs for ALTs. As pointed out by Soyer (2007), most of these Bayesian approaches are based on the theory of optimal Bayesian designs for linear models as in Chaloner (1984) and therefore, they either consider normal or lognormal failure models or use asymptotics. In the case of non Gaussian models such as exponential or Weibull, nonlinearities arise in the analysis. As a result, simulation methods are required to compute optimal designs. For example, Erkanli and Soyer (2000) use MCMC methods with nonparametric surface estimation to find optimal ALT designs and Zhang and Meeker (2006) consider large sample results for Bayesian ALT designs, as well as simulation-based methods. Nasir and Pan (2015) present simulation-based ALT designs for model discrimination.

A more recent treatment of the optimal design in ALTs can be found in Polson and Soyer (2017) who propose to use the APS method in Bielza et al. (1999) to compute optimal design a maximizing the preposterior expected utility (2). In the sequel, we discuss the APS approach of Polson and Soyer (2017). In doing so, we assume that the environment is characterized by a single stress and let the design a represent the stress level variable characterizing the accelerated test environment. It is possible to consider extensions to multiple stresses, as in Zhang and Meeker (2006), or testing at K accelerated levels of the stress variable.

2.2.1 APS model

Polson and Soyer (2017) point out that the evaluation of the preposterior expected utility $\bar{u}(a)$ via traditional MC techniques can become inefficient and hinder computing the optimal design $a^* = \arg \max_a \bar{u}(a)$, especially for high dimensional cases. To alleviate such inefficiencies, following Bielza et al. (1999), they propose an alternative approach by treating the design variable a as random and recasting the problem as one of drawing samples from the augmented probability model, defined through

$$\pi(x, \theta, a) \propto u(x, \theta, a)p(x, \theta|a)p(a), \quad (15)$$

where the distribution $p(a)$ is generally specified as a uniform distribution over the decision space. The "tilted" marginal distribution of a , that is, $\pi(a)$ is proportional to $\bar{u}(a)$ and therefore, the optimal design can be obtained by simulating samples from the marginal distribution of a and finding its mode. This can be done by using a MCMC scheme to draw from the augmented distribution (15). Polson and Soyer (2017) use a Gibbs sampler in their development by simulating from $\pi(x, \theta|a)$ and $\pi(a|x, \theta)$.

With higher dimensional a 's and flat expected utility surfaces, one can use $\pi(a)$ with a power type transformation as suggested by Mueller (1999). By drawing J samples $(x_j, \theta_j)_{j=1}^J$ for each design a , we can obtain

$$\pi_J(x^J, \theta^J, a) \propto \prod_{j=1}^J u(x_j, \theta_j, a)p(\theta_j, x_j|a), \quad (16)$$

where $\theta^J = (\theta_1, \dots, \theta_J)$ and $x^J = (x_1, \dots, x_J)$. It follows from (16) that $\pi_J(a) \propto u^J(a)$. As noted by Polson and Soyer (2017) the APS approach performs evaluation and optimization of $\bar{u}(a)$ simultaneously by treating the design variable a as a random quantity and simulating a , together with (x, θ) from the augmented probability model (15). In doing so, unlike the standard MC approach, the APS "tilts" MC draws to regions of high utility values and provides computational efficiency.

2.2.2 APS for ALT designs

Assume that failure times X_i under the stress environment a_i are exponentially distributed with failure rate θa_i . Assuming that n items are tested under environment a_i , the distribution of TTOT

$T(x, a_i) = \sum_{j=1}^n X_{ij}$ is gamma $Gam(n, \theta a_i)$. Using a gamma prior for $\theta \sim Gam(d, b)$, we can show that the predictive distribution is an inverted beta density

$$p(T(x, a_i)) = \frac{\Gamma(d+n)a_i/b}{\Gamma(n)\Gamma(d)} \frac{(T(x, a_i)a_i/b)^{n-1}}{(1+T(x, a_i)a_i/b)^{d+n}}.$$

The objective now is to select the accelerated environment $a_i > a_u$ to learn about the failure rate θa_u at the use stress. Polson and Soyer (2017) consider a single point design and assume a utility function

$$u(x, a) = \frac{1}{(a/a_u)^\alpha} e^{-kT(x,a)}, \quad (17)$$

where α and k are positive constants with $\alpha > 1$. Note that (17) is a *conjugate utility* function in the sense of Lindley (1976) and reflects the consequences of choosing a design a and observing a total time on test $T(x, a)$. It is desirable to test at a closer to a_u . However, smaller values of a will imply a lower failure rate and, thus, larger TTOT. In (17), α reflects the cost of selecting a design a away from a_u whereas k reflects the cost of testing. Without loss of generality we let $a_u = 1 < a$.

Using (17) and the gamma prior for θ , the augmented model (16) can be written as

$$\pi_J(T^J(x, a), \theta^J, a) \propto \prod_{j=1}^J \frac{1}{a^\alpha} e^{-kT_j(x,a)} (\theta_j a)^n (T_j(x, a))^{n-1} e^{-\theta_j a T_j(x,a)} \theta_j^{d-1} e^{-\theta_j b}, \quad (18)$$

where a is uniform over $(1, a_{max})$. The conjugate utility function allows us to design a Gibbs sampler as all full conditional distributions are available. Polson and Soyer (2017) show that the full conditional of a is $Gam(J(n - \alpha) + 1, s_J)$ where $s_J = \sum_{j=1}^J \theta_j T_j(x, a)$. Similarly, the full conditionals of the θ_j 's are $Gam(n + d, b + aT_j(x, a))$, for $j = 1, \dots, J$ and the full conditionals of the $T_j(x, a)$'s are $Gam(n, k + a\theta_j)$ for $j = 1, \dots, J$. By drawing iteratively from the full conditionals, the mode of the draws $a^{(g)}$'s histogram collapses on the optimal design. Polson and Soyer (2017) discuss implementation issues associated with APS as well as extension to multiple point designs.

2.3 Optimal stopping and optimal release

An important question in the testing of any product is when to finish testing and release it to the customer or place it on the market. This decision depends on how reliable the product is believed to be, as well as the costs and benefits of release: a decision to delay release will incur in costs for further testing and a potential loss of market advantage, but benefits of increased reliability, or reduced uncertainty about the reliability that will lead to lower expected repair costs or penalties for not meeting warranty or service agreement targets. This can be formulated as a decision problem in which:

- The actions are to release the product now, or test. If the decision is to test, this may be broken down further to consider how much further testing to do. Hence the set A of alternatives may simply take the form {test, release} or be the set of testing times $A = [0, \infty)$, with $a = 0$ corresponding to immediate release;
- A probability model $p(x | \theta)$ is assumed for the time to failure of the product with parameter θ representing the failure characteristics of the product with an elicited prior distribution $p(\theta)$;
- The consequences are usually the financial implications of the action, or its utility. They are usually based around the costs of testing, the costs of product failure after release and the potential costs of the loss of market advantage due to delaying release. The utility can be specified



Figure 1: Decision tree for single stage testing.

as $u(\theta, a)$ reflecting the consequences of releasing the product with failure characteristic θ following the testing decision a .

The simplest testing strategy in this context is that of single-stage testing, where testing is to be done for a period of time and then the product is released, regardless of the results of the test. The decision is to choose the optimal test time. This strategy can be represented by a decision tree (Figure 1). In this case, the decision is purely based on the utility and prior on θ , with the optimal decision being

$$\arg \max_a \bar{u}(a) = \int_{\theta} u(\theta, a) p(\theta) d\theta.$$

If the utility depends explicitly on the failure time x of the released product, $u(x, \theta, a)$, then as in (3), the optimal decision is:

$$\arg \max_a \bar{u}(a) = \int \int_{(x, \theta)} u(x, \theta, a) p(x | \theta) p(\theta) dx d\theta.$$

The 'full' solution would provide for multiple testing stages, with the option at the end of each stage to release or continue for another stage. This gives the decision tree in Figure 2. The optimal strategy now consists of a sequence of testing times a_1, a_2, a_3, \dots , with a_i depending on the prior and the data up to the last testing period x_1, \dots, x_{i-1} . This would include conditions under which release would occur after i testing stages. Unfortunately this optimal strategy is generally infeasible to compute, being a dynamic programming type of problem with nested maximizations and expectations, so that the utility at the n -th stage will depend in some way on the optimal action at the $(n + 1)$ th stage:

$$a_n^* = \arg \max_{a_n} E(u(X_n, \theta, a_n; a_{n+1}^*)),$$

where $X_n = (x_n, x_{n+1}, \dots)$ are the test data from the n th and subsequent stages. Only in special circumstances will such computations scale to even moderate n (Dunsmore and Wright, 1985).

The usual solution is to assume a Markov property, and solve the sequence of single stage problems, with the posterior distribution at stage n given x_1, \dots, x_{n-1} being used:

$$a_n^* = \arg \max_{a_n} E(u(x_n, \theta, a_n) | x_1, \dots, x_n) = \int_{(x_n, \theta)} u(x_n, \theta, a) p(x_n | \theta) p(\theta | x_1, \dots, x_{n-1}) dx_n d\theta.$$

The work of Barnett (1972) was one of the first to consider this idea.

An illustration of this approach is in McDaid and Wilson (2001), applied to software testing. The goal is to determine the time a to test the software. Uncertainty arises from $N(a)$, the number of bugs discovered in the software by time a , and then $\bar{N}(a) = N(\infty) - N(a)$, the number that are discovered after time a . Many probability models have been proposed for this (Singpurwalla and Wilson, 1999). A popular one is that of Goel and Okumoto (1979), in which $N(a)$ is a Poisson process with mean function $M(a) = \theta_1(1 - e^{-\theta_2 a})$, for parameters $\theta = (\theta_1, \theta_2)$ that represent the expected total number of bugs in the software and the discovery rate, respectively. Thus, $N(a)$ is Poisson distributed with

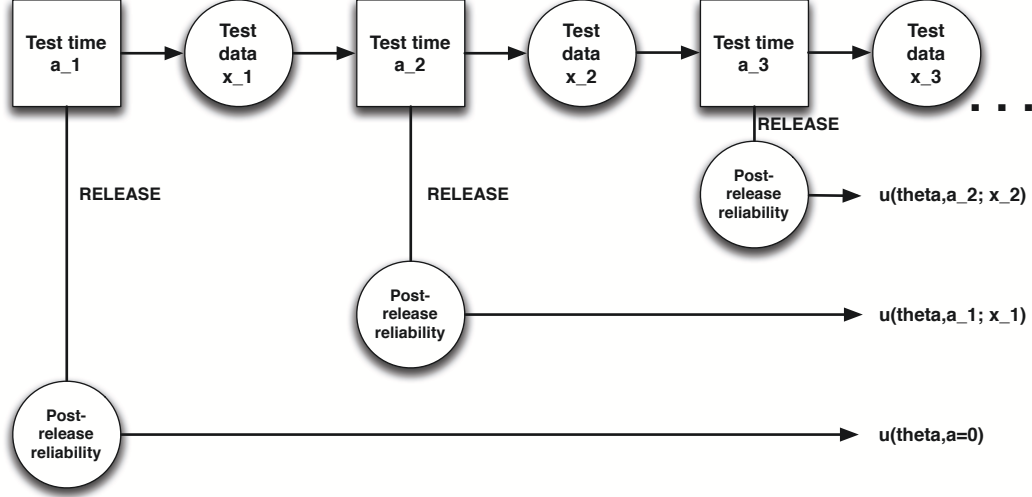


Figure 2: The decision tree for multiple-stage testing.

expected value $M(a)$. It is possible to consider extensions of this model by considering covariate information as in Ray et al. (2006).

A simple form for the utility of testing to time a , with $N(a)$ bugs discovered in testing and $\bar{N}(a)$ discovered after release, is:

$$u(x = (N(a), \bar{N}(a)), a, \theta) = B - CN(a) - D\bar{N}(a) - Fa, \quad (19)$$

where B is the profit from releasing the software without any testing, C is the cost of fixing a bug discovered in testing, D is the cost of fixing a bug post-release and F is the cost per unit time of testing, that includes both the testing costs as well as lost sales and market opportunity. In practice, D should be considerably larger than C . Gamma distributions were used as priors with parametric form $p(\theta) = \alpha^\beta \theta^{\beta-1} e^{-\alpha\theta} / \Gamma(\beta)$, with mean β/α and standard deviation $\sqrt{\beta}/\alpha$. McDaid and Wilson (2001) describe an elicitation process for these parameters based on these relationships. Here, we assume that such elicitation process has led to specifying a gamma prior with parameters (α_1, β_1) for θ_1 , and (α_2, β_2) for θ_2 . Given θ , $N(a)$ and $\bar{N}(a)$ are Poisson distributed. Averaging out over the prior on θ gives the expected values of $N(a)$ and $\bar{N}(a)$

$$E(N(a)) = \frac{\beta_1}{\alpha_1} \left[1 - \left(\frac{\alpha_2}{\alpha_2 + a} \right)^{\beta_2} \right] \quad (20)$$

$$E(\bar{N}(a)) = \frac{\beta_1}{\alpha_1} \left(\frac{\alpha_2}{\alpha_2 + a} \right)^{\beta_2}. \quad (21)$$

Hence the expected utility of testing to time a is:

$$\bar{u}(a) = B - C \frac{\beta_1}{\alpha_1} \left[1 - \left(\frac{\alpha_2}{\alpha_2 + a} \right)^{\beta_2} \right] - D \frac{\beta_1}{\alpha_1} \left(\frac{\alpha_2}{\alpha_2 + a} \right)^{\beta_2} - Fa,$$

and the optimal time a^* to test, which maximizes this function is:

$$a^* = \alpha_2 \left[\left(\frac{\beta_1 \beta_2 (D - C)}{\alpha_1 \alpha_2 F} \right)^{1/(\beta_2 + 1)} - 1 \right]; \quad (22)$$




Figure 3: Expected utility as a function of time (left) and optimal testing time as a function of post-release bug cost D (right).

assuming that $D > C$. When $D \leq C$, the optimal strategy is not to test and just repair all bugs post-release.

Figure 3 illustrates the case where the prior mean on a is 100 (we expect about 100 bugs in the code) and the prior mean on b is 0.01 (based on $\alpha_1 = 0.01, \beta_1 = 1, \alpha_2 = 100$ and $\beta_2 = 1$) and utility parameters are $B = 2000, C = 1, D = 20$ and $F = 0.5$. The left plot shows the expected utility as a function of a , and identifies $a^* = 516.4$ for an expected utility of 1333.6. The right plot shows how a^* changes as a function of D , the cost of fixing a bug post-release; this shows how the testing time should increase as the relative cost of fixing bugs after testing rises.

In this solution, release occurs regardless of the results of the testing and so there is no opportunity to learn about the software’s reliability from the testing results. McDaid and Wilson (2001) discuss other solutions that involve more than one stage of testing allowing for learning to take place, including the Markov sequence of single stage tests. Singpurwalla (1991) consider the Bayesian solution for a two-stage testing problem. Morali and Soyer (2003) discuss optimal stopping problem in software testing and investigate the possibility of developing one stage ahead optimal stopping rules using results from van Dorp et al. (1997).

2.4 Reliability demonstration test plans

The goal of reliability demonstration testing (RDT) is to accumulate enough evidence so that the reliability of the product under consideration has achieved or not a given level. This is in contrast to optimal testing in which the goal is to reach the optimal trade off between further testing and release. Demonstration testing will typically form part of product development, although it may also arise when a vendor wishes to convince a buyer or a regulator that its product meets reliability requirements, in which case it has links to adversarial life testing as discussed in Section 4.2.

The formulation of a solution to RDT requires a reliability model, a definition of a reliability metric and a pass/reject criterion. A pioneer Bayesian example was in Schafer and Singpurwalla (1970), who assumed exponential failure times with mean θ , a conjugate inverse gamma prior on θ and meeting a required mean time to failure θ_1 as a reliability criterion. A sequence of items was tested and failure times t_1, t_2, \dots were observed. After the failure of each item under test, the test stopped if the posterior probability $P(\theta > \theta_1 | t_1, \dots, t_n)$ either exceeded $1 - \alpha_2$, in which case the product

was passed, or fell below $1 - \alpha_1$, in which case the product was rejected, for specified threshold probabilities $0 < \alpha_2 < \alpha_1 < 1$.

Barnett (1972) advanced another important aspect of RDT by using a continuous pass/reject criterion, rather than one that could only be evaluated after each failure of the items on test. A large amount of subsequent work in the seventies extended this work in many ways, nicely reviewed in Martz and Waller (1982). Higgins and Tsoskos (1976) showed that RDTs could be very sensitive to the choice of prior on the reliability metric. Martz and Waller (1979) considered the case in which tests show no failure, extending work to the case of highly reliable systems.

Bayesian reliability demonstration tests have largely followed this format of metric, reliability model, prior specification and accept/reject criterion, either defined explicitly or based on a utility or loss function. After a hiatus in work in the 1980s and the 1990s, there has been an increase of activity in the past 15 years, and even Bayesian RDT has now been incorporated into official reliability testing standards (Yates, 2008).

More recent work has had more of a focus on using decision theory in the determination of when to stop testing and accept or reject, based on a loss function for a product with a given level of reliability, as well as other decisions such as optimal testing strategies. For example, Rahrourh et al. (2006) models a process of tests and decisions for a system with redundancy, taking into account the costs of testing as well as adding extra redundancy versus the benefits of the increased reliability. Jin and Matthews (2014) develop an approach to planning optimally the test, taking into account the costs of testing and measuring reliability.

Sun and Berger (1994) is a good example of the Bayesian approach to RDT. They considered a generalization of Schafer and Singpurwalla (1970) introducing an additional mature product goal $\theta_2 > \theta_1$ such that the product is rejected when the posterior probability that $\theta < \theta_2$ exceeds a threshold. They also considered a reliability model more general than the exponential, and introduced a loss function of the form:

$$l(\theta) = \begin{cases} 0, & \text{if } \theta_1 < \theta < \theta_2, \\ l, & \text{otherwise,} \end{cases}$$

for some positive loss $l > 0$. In this case, a closed-form solution for the accept/reject rule that minimizes the expected loss can be derived. For example, in the exponential reliability model case, a conjugate inverse-gamma prior on the mean time to failure θ with shape parameter a and scale parameter b , $p(\theta) = \frac{b^a}{\Gamma(a)} \theta^{-(1+a)} e^{-b/\theta}$, can be defined. Then, the posterior distribution of θ after testing N items to time t , of which n failed at times t_1, \dots, t_n , is also inverse gamma with shape parameter $b + n$ and scale parameter $a + V(t)$, where $V(t)$ is the total time on test to time t :

$$V(t) = (N - n)t + \sum_{i=1}^n t_i.$$

Hence the accept/reject rule becomes:

$$\begin{aligned} \text{If } q^*(\alpha_1) > \theta_1 & \text{ then stop testing and accept product;} \\ \text{If } q^*(1 - \alpha_2) \leq \theta_2 & \text{ then stop testing and reject product;} \\ \text{Otherwise continue testing,} \end{aligned}$$

where $q^*(\alpha)$ is the $100\alpha\%$ quantile of the posterior distribution. As the inverse gamma is related with

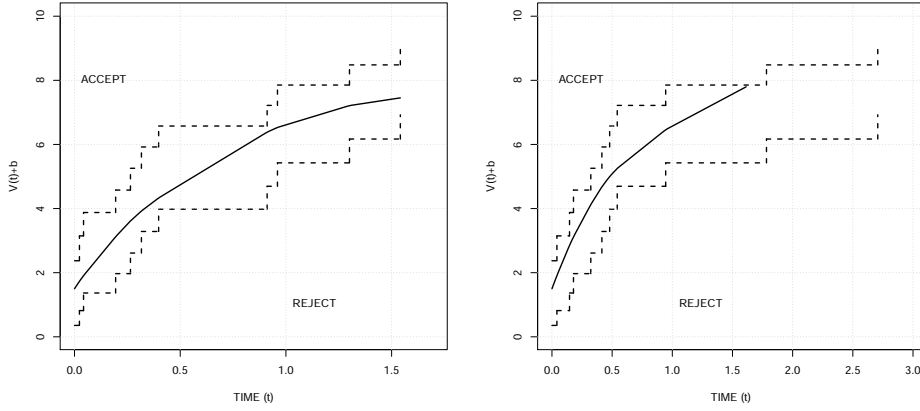


Figure 4: Two examples of the reliability demonstration test of Sun and Berger (1994). Upper dashed line is the accept threshold; lower dashed line, the reject threshold; solid line is $V(t) + b$.

the χ^2 distribution, it can be shown that the rule can be written in terms of χ^2 values:

$$\begin{aligned} \text{If } V(t) + b > \frac{\theta_1}{2} \chi^2(2(n+a), 1 - \alpha_1) \quad \text{then stop testing and accept product;} \\ \text{If } V(t) + b > \frac{\theta_2}{2} \chi^2(2(n+a), \alpha_2) \quad \text{then stop testing and reject product,} \end{aligned}$$

where $\chi^2(m, p)$ is the $100p\%$ quantile of the χ^2 distribution with m degrees of freedom.

Figure 4 illustrates how this rule works in practice, with $V(t) + b$ and the accept and reject thresholds, plotted as a function of time. It shows two examples with parameters $a = 2$, $b = 1.5$, $\theta_1 = 0.5$, $\theta_2 = 1.0$, $\alpha_1 = \alpha_2 = 0.05$. Ten units are placed on test. On the left, all ten units fail without the test stopping. On the right, the test reaches the accept boundary. Observe that the thresholds are step functions, changing value when a failure is observed, while $V(t) + b$ is strictly increasing, and that testing continues until $V(t) + b$ first hits one of the thresholds. Also note that $V(0) = 0$, hence $V(0) + b = b$ and so it is possible that the test is immediately passed or rejected at time 0 if b lies above the initial accept or below the initial reject boundary, further backing up the assertion of Higgins and Tsokos (1976) that the prior can play an important role in the outcome of an RDT.

3 Preventive Maintenance Policies

Since the seminal work of Barlow and Proschan (1965), preventive maintenance has become common practice for systems which are subject to deterioration as a result of usage and aging. The main objective of preventive maintenance is to prevent system failures to avoid costly service disruptions. Maintenance activities may take different forms including repairs, replacements as well as other practices that could prevent or delay system failure. The development of maintenance strategies have attracted considerable attention in the OR/MS literature; see for example, the review papers by Cho and Parlar (1991), Wang (2002) and Shafiee and Chukova (2013).

In this section we present the treatment of maintenance strategies from a Bayesian perspective. We first discuss maintenance concepts and policies for repairable systems. This is followed by a discussion of replacement policies, including replacement with minimal repair, where the Bayesian framework for optimal replacement is introduced. Recent work on parametric and nonparametric Bayesian replacement strategies is presented, including computational issues. Finally, sequential

maintenance problems are considered and semi Markov decision processes are introduced to develop Bayesian policies.

3.1 Policies for repairable systems

In reliability analysis, there is a major difference between repairable and non-repairable systems. The latter are to be replaced upon failure; light bulbs, covers of smart phones and window glasses are examples of non-repairable systems which lead to a sequence of lifetime distributions which, if i.i.d., give rise to a renewal process. Here we focus mostly on repairable systems which, in the event of a failure, can be repaired, for example, by replacing a component, and returned to regular operation. In some cases, the reliability of a system, after a “minimal” repair, returns to the same state as before the failure. On the other hand, “perfect” repairs bring the system reliability back to the state at the start of the operation. “Imperfect” repair is referred to the case in between these two; see Doyen and Gaudoin (2004). Failures of repairable systems are often described by means of non-homogeneous Poisson processes (NHPP). Whereas in Section 3.2 the focus will be on traditional decision-theoretic approaches, here we consider Bayesian policies based on posterior predictive distributions and other posterior estimates (which can also be justified from a decision-theoretic point of view).

Consider a NHPP $N(t)$ with intensity function $\lambda(t, \theta)$ and mean value (or cumulative intensity) function $\Lambda(t, \theta)$. Suppose we observe the system up to time τ and let n denote the observed number of failures at times $t_1 < t_2 \dots < t_n < \tau$. Then, the likelihood function of θ is given by

$$L(\theta; D) = \prod_{i=1}^n \lambda(t_i, \theta) \exp\{-\Lambda(\tau, \theta)\}, \quad (23)$$

where $D = (t_1, \dots, t_n)$. For example for the intensity function $\lambda(t, \theta) = \theta t$, with cumulative intensity $\Lambda(t, \theta) = \theta t^2/2$, the likelihood function (23) becomes

$$L(\theta; D) = \theta^n \prod_{i=1}^n t_i \exp\{-\theta \tau^2/2\}.$$

With a conjugate gamma prior $Gam(\alpha, \beta)$ on θ , the posterior distribution $p(\theta|D)$ is given by $Gam(\alpha + n, \beta + \tau^2/2)$.

Consider now a new copy of the system for which the question is “how long will it operate with rare chances of failing?” In terms of the reliability function $R(t)$, it means finding the largest T such

$$R(T|D) = Pr(N(T) = 0|D) = \int Pr(N(T) = 0|\theta)p(\theta|D)d\theta = \left\{ \frac{\beta + \tau^2/2}{\beta + \tau^2/2 + T^2/2} \right\}^{\alpha+n} > \epsilon$$

where ϵ is the risk threshold. The optimal T is given by

$$T^* = \sqrt{2(\beta + \tau^2/2)(\epsilon^{-1/(\alpha+n)} - 1)}.$$

In this case, $R(T|D)$ is seen as a posterior predictive probability. Similarly, we may consider the reliability of the system after time τ and the expected number of failures in future intervals, either for the current or a new copy of the system.

Pievatolo and Ruggeri (2004) considered forecasting gas escapes in the steel pipelines of a city network. Historical data of 33 failures during the period 1978-1997 for a network of 275 km was used to infer the behavior of the network for the following 5 years. Decisions included the possible replacement of pipelines with others less at risk and the size of the emergency and repair squads.

The authors provided not only point estimates of the system reliability and the expected number of gas escapes, but also their 95% credible intervals. The study predicted a limited number of gas escapes which made replacement financially non viable.

Posterior predictive probabilities were used also by Pievatolo et al. (2003) and Pievatolo and Ruggeri (2010) to assess if the reliability of underground train doors was compliant with the contract signed by a manufacturer and a transportation company. The decision involved whether an action was needed by the manufacturer to improve reliability to the contract level before warranty had expired. The authors used different NHPPs to describe failure data which was available both in the form of calendar time and kilometers run up to failure. Pievatolo et al. (2003) considered a univariate intensity function (in time) where kilometers and time were related either by a Gaussian model or a Gamma process, whereas Pievatolo and Ruggeri (2010) proposed a bivariate intensity function in time and kilometers.

Another example of the use of predictive probabilities in determining the replacement of an item is by Hermann and Ruggeri (2017) who consider data on the wear of cylinder liners in ships. Liners are approximately 10 meters high and have a thickness of 100 millimeters (mm); they are supposed to protect cylinders from breakage. If a cylinder breaks when the liner's wear is less than 4 mm, then the cost for the replacement of the cylinder is on the manufacturer of the liner; if it is more, the cost is on the ship's owner. Maintenance is very expensive because of the size of cylinders and the loss of operational days by the ship. By considering the problem from the ship-owner's point of view, a stochastic differential equation model was proposed for the thickness of the liner (with a threshold at 96 mm, then) and forecasts were developed for its evolution over time. In particular, it was possible to build 95% credible intervals for the thickness at future times which could be used for replacement decisions.

Cagno et al. (2000) considered the problem of which pipelines of a city gas network were to be replaced first to reduce the number of gas escapes. Eight subnetworks were determined by three covariates (diameter, depth and location of the pipes) assuming two levels each. Gas escapes were modeled by homogeneous Poisson processes and the decision was based on the largest parameter among the eight subnetworks, as estimated by respective posterior means. More recently, Arias et al (2015) considered the same problem by taking into account the uncertainty in modeling the prior distributions, which were originally obtained from qualitative assessment by company's experts using the Analytic Hierarchy Process and merging of opinions. The authors considered classes of priors based on empirical quantiles from the distribution of expert opinions, and concentrated on non dominated actions showing that a particular subnetwork was a posteriori the worst for all possible priors in the class.

3.2 Optimal replacement

For most preventive maintenance strategies, a major issue is the determination of a planned replacement (or maintenance) interval. This is performed in an optimal manner by considering the trade off between in-service failure and planned replacements costs. Typically, the planned replacement is less expensive than the in-service failure and subsequent service replacement. Two basic replacement protocols are the block and the age replacement policies.

Under a block replacement policy, the system in question is replaced at times $t_B, 2t_B, \dots$, irrespective of the age of the system; it is also replaced at the time of a failure. Planned replacements are specified in advance whereas the time of in-service replacement is unknown. If the cost of planned replacement is given by c_P , the cost of an in-service failure is c_F , with typically $c_F > c_P$, and $N(t_B)$ denotes the number of system failures for a time interval of length t_B , the cost of replacement per

unit time is given by

$$C(t_B) = \frac{c_P + c_F N(t_B)}{t_B}. \quad (24)$$

In (24), $N(t_B)$ is a renewal process. The optimal replacement interval t_B^* is obtained by minimizing the expected cost

$$E[C(t_B)] = \frac{c_P + c_F H(t_B)}{t_B}, \quad (25)$$

where $H(t_B)$ is the renewal function; see Cox (1962). Here it is assumed that after each failure the system is replaced by a new one or the repair applied to the system is perfect. Thus, this is referred to as the "good as new" scenario.

An alternative block replacement strategy is considered in Barlow and Hunter (1960) where the system is assumed to be minimally repaired upon failure, but replaced at times $t_B, 2t_B, \dots$ irrespective of its age. In this case after each failure it is assumed that the system can be restored to the state just prior to the failure. Thus, it is referred to as the "bad as old" scenario. Block replacement with minimal repair is applicable in those cases where replacement of the whole system is more costly, due to disruption of service, than its minimal repair. One such example is the replacement of railroad tracks where replacement of a track which is of miles in length is very costly compared to the repair; see Merrick and Soyer (2017). If we denote the minimal repair (MR) cost by c_R , then the cost per time $C(t_B)$ is given by (24) where c_F is replaced by $c_R < c_P$. In the MR case $N(t_B)$ is modeled by a NHPP. Similar to the "good as new" scenario, the optimal replacement interval is obtained by minimizing the expected cost

$$E[C(t_B)] = \frac{c_P + c_R \Lambda(t_B)}{t_B}, \quad (26)$$

where $\Lambda(t_B)$ is the cumulative intensity (or mean value) function of the NHPP. Other extensions of block replacement policies can be found in Sheu and Griffith (2002).

Under an age replacement policy, it is assumed that a planned replacement is performed when the age of the system reaches a specified time t_A or an in-service replacement is made when the system fails. As in the block replacement policy under the "good as new" scenario, the cost of in-service failure is assumed to be larger than the planned replacement cost, that is, $c_F > c_P$. However, in this case, unlike the block replacement policy, the length of the replacement cycle is random and is given by $\min(t_A, T)$ where T is the life time of the system. For an age replacement interval t_A , the cost per unit time is given by

$$C(T, t_A) = \begin{cases} c_P/t_A, & \text{if } T \geq t_A \\ c_F/T, & \text{if } T < t_A. \end{cases} \quad (27)$$

If $F(t)$ denotes the distribution function for T with failure density $f(t)$ and $\bar{F}(t) = 1 - F(t)$ is the reliability (or the survival) function, then the expected cost per unit time can be obtained as

$$E[C(T, t_A)] = \int_0^{t_A} \frac{c_F}{t} f(t) dt + \frac{c_P}{t_A} \bar{F}(t_A), \quad (28)$$

and the optimal interval is obtained by minimizing (28) with respect to t_A . Alternatively, the optimal age replacement interval can be obtained by minimizing the long-run average cost given by

$$\bar{C}(t_A) = \frac{c_P \bar{F}(t_A) + c_F F(t_A)}{t_A \bar{F}(t_A) + \int_0^{t_A} t f(t) dt}. \quad (29)$$

Note that in (29), the numerator represents the expected cost for the replacement cycle and the de-

nominator is the expected cycle length. Extensions of the basic age replacement policy can be found in Chien (2008).

3.2.1 Bayesian replacement strategies

As noted in Mazzuchi and Soyer (1996a), most implementations of replacement policies in the literature are based on the assumption that the failure characteristics of the system are specified. For example, in the block replacement with MR scenario, for the cumulative intensity function $\Lambda(t_B)$ in (26) a parametric form $\Lambda(t_B, \theta)$ is specified and the parameter θ is assumed to be known. A commonly used model is the *power law* with cumulative intensity function

$$\Lambda(t, \theta) = \alpha t^\beta, \quad (30)$$

where $\theta = (\alpha, \beta)$ and both parameters are positive. The intensity function is given by $\lambda(t, \theta) = d\Lambda(t, \theta)/dt = \alpha\beta t^{\beta-1}$, where $\beta > 1$ implies deterioration over time. Assuming that α and β are given, the optimal replacement interval t_B^* is

$$t_B^* = \left(\frac{c_P}{c_R \alpha (\beta - 1)} \right)^{1/\beta}.$$

The Bayesian decision theoretic approach requires that the uncertainty about θ to be specified probabilistically via a prior distribution $p(\theta)$ and the associated expected cost function to be minimized, where the expectation is taken with respect to all unknown quantities including the unknown parameters θ . This involves the minimization of

$$E[C(t_B)] = \frac{c_P + c_R E_\theta[\Lambda(t_B, \theta)]}{t_B}, \quad (31)$$

with respect to t_B in the block replacement with MR case and the minimization of

$$E[C(T, t_A)] = \int_0^{t_A} \frac{c_F}{t} f(t|\theta) p(\theta) d\theta dt + \int_{t_A}^{t_A} \frac{c_P}{t_A} \bar{F}(t|\theta) p(\theta) d\theta. \quad (32)$$

with respect to t_A , in the age replacement case where a parametric form $F(t|\theta)$ is specified for the system life distribution. When data D is available from previous replacement cycles, the uncertainty about θ is revised to obtain the posterior distribution $p(\theta|D)$ which replaces $p(\theta)$ in (31) and (32) to obtain the associated optimal intervals. Similarly, one can develop Bayesian block replacement strategies under the "good as new" case by making inference on the renewal function $H(t, \theta)$.

One of the earliest Bayesian approaches to replacement is due to Fox (1967), who considered age replacement policies when the failure model is Weibull. Its shape parameter was assumed to be known and adaptive optimal policies were developed by updating its scale parameter and minimizing the expected discounted cost over time. The author obtained asymptotic results for an infinite planning horizon. Sathe and Hancock (1973) developed Bayesian policies using a Weibull model where both the shape and scale parameters were treated as unknown quantities. The authors minimized the expected long-run average cost (29). An earlier attempt for developing Bayesian block replacement policies is by Bassin (1973) who used Bayesian point estimates in the MR scenario by using a power law model.

Mazzuchi and Soyer (1996a) developed Bayesian block and age replacement policies minimizing (31) and (32). A power law model (30) was used for the intensity function and an adaptive policy was developed by revising uncertainty about the parameters $\theta = (\alpha, \beta)$ as well as the optimal replacement interval after each cycle. In doing so, for a block replacement cycle of length t_B , where n minimal

repairs are performed at times $t_1 < t_2 < \dots < t_n < t_B$ the distribution of (α, β) was revised based on the likelihood function

$$L(\alpha, \beta; D) = \prod_{i=1}^n \lambda(t_i, \theta) \Lambda(t_B, \theta), \quad (33)$$

where $\lambda(t_i, \theta)$ is the intensity function of the power law model evaluated at t_i and $D = (t_1, \dots, t_n)$, see Pievatolo and Ruggeri (2004).

For the age replacement protocol, Mazzuchi and Soyer (1996a) assumed a Weibull failure model

$$\bar{F}(t|\alpha, \beta) = e^{-\alpha t^\beta} \quad (34)$$

and developed adaptive age replacement policies by revising uncertainty about the shape and scale parameters, α and β , respectively, after each cycle. During each age replacement cycle of t_A , uncertainty about α and β was revised based on the likelihood function

$$L(\alpha, \beta; t) = f(t|\alpha, \beta)I\{t < t_A\} + \bar{F}(t_A|\alpha, \beta)I\{t > t_A\}, \quad (35)$$

where $\bar{F}(t_A|\alpha, \beta)$ is given by the Weibull reliability function (34), $f(t|\alpha, \beta)$ is the corresponding density, and $I\{\cdot\}$ is the indicator function. Numerical integration methods were used to evaluate the expected cost (32).

Block replacement under the "good as new" scenario was considered by Mazzuchi and Soyer (1996b). The renewal function used was

$$H(t; \theta) = E[N(t|\theta)] = \sum_{n=1}^{\infty} F^{(n)}(t|\theta), \quad (36)$$

where $F^{(n)}(t|\theta)$ is the n -fold convolution of the failure model $F(t|\theta)$. A Weibull model was used as failure model and the renewal function (36) was approximated using the numerical approach proposed in Smeitink and Dekker (1990). A Monte Carlo (MC) approach was used to evaluate the expected cost (25) and determine the optimal replacement interval. An adaptive strategy was applied by updating uncertainty about the parameters after each cycle and obtaining the optimal interval accordingly. Revision of the parameters of the Weibull failure model (34), after a given replacement cycle of t_B , was based on the likelihood function

$$L(\alpha, \beta; D) = \left(\prod_{i=1}^n f(t_i - t_{i-1}|\alpha, \beta) \right) \left(\prod_{i=1}^n \bar{F}(t_B - t_n|\alpha, \beta) \right) \quad (37)$$

where $t_1 < \dots < t_n$ are the failure times (collectively designated as D).

The above adaptive replacement strategies of Mazzuchi and Soyer (1996a, 1996b) have been generalized to other replacement scenarios. Sheu et al. (1999) considered age replacement strategies with MR. Dayanik and Gurler (2002) developed strategies for more general MR protocols considered by Bechtel (1993) and a related adaptive preventive maintenance approach was discussed in Juang and Anderson (2004). Bayesian group replacement policies were studied by Wilson and Benmerzouga (1995) and Popova (2004).

Advances in computational Bayesian methods and particularly in MCMC approaches after 1990 have greatly enhanced the Bayesian analysis of repairable systems as well as the development of maintenance strategies. As an example, the effect of grinding on reliability of rail tracks was studied in Merrick et al. (2005) through modulated Poisson process models and block replacement strategies. In rail tracks, usage is measured in millions of gross tons (MGT) traversing on the rail. In order to prevent derailments caused by rail fractures that develop with heavy usage, "the cracked rail can

be either ground down, removing the metal surrounding the crack and leaving only solid metal, or welded, fusing the crack". As long as the initial crack does not lead to a complete fracture of the rail, the resulting repair is considered as minimal. A modulated Poisson process model (MPPM) is considered in Merrick et al. (2015) to describe the repair process. This is achieved by modulating the cumulative intensity function of a NHPP with a vector of covariates as suggested in Cox (1972a). More specifically, for a rail track i , the authors considered a cumulative intensity function

$$\Lambda_i(t, \theta, \beta, Z_i) = \Lambda_0(t, \theta) e^{-Z_i \beta}, \quad (38)$$

where Z_i is a vector of covariates and β is a regression parameter associated with Z_i . For the MPPM, the baseline cumulative intensity function $\Lambda_0(t, \theta)$ is modulated by the covariate vector Z_i . In Merrick et al. (2005), a power law model as in (30) was considered for the baseline cumulative intensity function and the covariate vector was assumed to be independent of usage. In their set up, the authors used grinding level as one of the rail specific covariates. Bayesian inference for the model was developed using a Gibbs sampler with adaptive rejection sampling steps, Dellaportas and Smith (1993).

More recent work in Bayesian replacement policies is in Belyi et al. (2017) who considered *bathhtub* failure rates.

3.2.2 Nonparametric replacement policies

Relaxations of the parametric assumptions in replacement models have been considered in the optimal maintenance literature. One of the earliest works is by Arunkumar (1972) who considered nonparametric age replacement strategies. Adaptive versions of nonparametric age replacement policies were developed in Frees and Ruppert (1985).

The nonparametric Bayesian framework has been proposed in the reliability literature to provide more flexibility in modeling uncertainty about the failure model $F(t|\theta)$ or the failure rate (or intensity rate) $\lambda(t, \theta)$. This is achieved by introducing a prior distribution over the class of failure models or the class of failure rate functions. Although the seminal work by Ferguson (1973), on Dirichlet process priors, and Antoniak (1974), on mixtures of Dirichlet processes, have contributed significantly to the development of Bayesian nonparametrics, as pointed out by Singpurwalla (2006, p. 244), the origins of Bayesian nonparametrics can be traced back to the works of Ramsey (1972) and Kraft and van Eeden (1964) on potency curve estimation in bioassay.

Earlier use of nonparametric methods in reliability modeling and survival analysis is presented in Kalbfleisch (1978) who assumed a gamma process prior to describe the cumulative failure rate; Dykstra and Laud (1981), who considered extended gamma processes for modeling nondecreasing failure rates; and Mazzuchi and Singpurwalla (1985), who proposed using an ordered Dirichlet prior for monotone failure rates. Although Bayesian nonparametric approaches have been used in statistical decision problems, as reviewed in Gutierrez-Pena and Walker (2005), their implementation in reliability decision problems have started only during the last fifteen years. This is mostly due to the advances in Bayesian computing and the associated MCMC methods which allow for simulation from posterior processes such as mixtures of Dirichlet processes, for example Escobar and West (1995), and failure rate processes, for example Laud et al. (1996).

One of the earliest development of nonparametric Bayesian replacement policies is by Merrick et al. (2003) who consider age and block replacement of machine tools consisting of multiple non-repairable components. The authors extend the proportional hazards model (PHM) for machine tool reliability assessment in Mazzuchi and Soyer (1989) by relaxing the parametric base line failure rate assumption. More specifically, under the parametric PHM of Cox (1972b) for T_i , the lifelength of

machine tool i , the failure rate function of the distribution of T_i is given by

$$\lambda_i(t, \theta, \beta, Z_i) = \lambda_0(t, \theta) e^{-Z_i' \beta} \quad (39)$$

where $\lambda_0(t, \theta)$ is the baseline failure rate and Z_i is a vector of covariates as in (38). This vector describes the operational environment of the machine tools including components such as cutting speed, feed rate or depth of cut. A strategy to relax the parametric assumptions in (39) is to assume a prior for the baseline failure rate function $\lambda_0(t, \theta)$ while treating β with a parametric prior. This yields a *semiparametric Bayesian* model for the failure rate. As noted by Merrick et al. (2003), the covariate vector Z_i does not capture potential variation in the individual characteristics of machine tools. To account for such heterogeneity, the authors consider a machine tool specific baseline failure rate $\lambda_0(t, \theta_i)$ where θ_i is the vector of unknown parameters associated with machine tool i and describe uncertainty about the θ_i 's through a prior distribution G whose form is treated as unknown. Indeed, a Dirichlet process prior is assumed as

$$G \sim DP(G_0, M), \quad (40)$$

where G_0 is a *best guess* baseline prior for G and M is the *strength of belief*, or precision parameter as in Ferguson (1973). Specifying $\lambda_0(t, \theta_i)$ conditional on θ_i gives a conditional parametric model for T_i with density $f(t_i | \theta_i, \beta, Z_i)$. The semiparametric Bayes model specification is completed by a parametric prior $p(\beta)$ for β . Merrick et al. (2003) point out that the distribution of T_i can be represented as an unknown mixture given by

$$f(t_i | G, \beta, Z_i) = \int f(t_i | \theta_i, \beta, Z_i) dG(\theta_i) \quad (41)$$

which provides the nonparametric structure of the model. Since G follows a Dirichlet process, this is referred to as a *Dirichlet process mixed model*. Given failure time and covariate data $D = (t_1, \dots, t_n, Z_1, \dots, Z_n)$ from n machine tools, the likelihood function of G and β is obtained as a product of the density functions in (41), that is, $L(G, \beta; D) = \prod_{i=1}^n f(t_i | G, \beta, Z_i)$. The posterior analysis of the model requires a Gibbs sampler with a Metropolis step to draw posterior samples from β . Since it is difficult to draw from the posterior full conditional distribution $p(G | \beta, D)$ directly, the Gibbs sampler proposed by Escobar and West (1995) is used to sample from the full conditionals of $\theta^{(n)} = (\theta_1, \dots, \theta_n)$ and β . Once posterior samples are drawn from $p(\theta^{(n)}, \beta | D)$, the expected posterior cost for the age replacement interval t_A can be approximated by the Monte Carlo average

$$E[C(T, t_A) | D, Z_i] = \frac{1}{S} \sum_{s=1}^S \int_0^{t_A} \frac{c_F}{t} f(t | \theta_i^s, \beta^s, Z_i) dt + \frac{c_P}{t_A} \bar{F}(t_A | \theta_i^s, \beta^s, Z_i) \quad (42)$$

for machine tool i , where (θ_i^s, β^s) are the posterior samples. The optimal age replacement interval t_A^* is then obtained by minimizing (42) with respect to t_A .

Following Mazzuchi and Soyer (1989), a Weibull baseline failure rate $\lambda_0(t, \theta_i) = \alpha_i \gamma t^{\gamma-1}$, where $\theta_i = (\alpha_i, \gamma)$, was used by Merrick et al. (2003) in the semiparametric PHM. The authors show that the semiparametric PHM fit the failure data better than the parametric model based on different model comparison criteria. Optimal age replacement intervals were obtained for different machine tools under the parametric and semiparametric models and significant differences were found between both approaches, with semiparametric policies providing more conservative results in most cases.

Nonparametric block replacement policies were also considered in Merrick et al. (2003) for a group of m machine tools. The authors used a Weibull failure model for machine tools as in the age replacement protocol. In this case, the cost under the common block replacement interval t_B for m

tools is defined through

$$C(t_B) = \sum_{j=1}^m \frac{c_P + c_F N_j(t_B | \theta_j, \beta, Z_j)}{t_B}, \quad (43)$$

where the $N_j(t_B | \theta_j, \beta, Z_j)$'s are (conditionally) independent renewal processes with respective renewal function $H_j(t_B | \theta_j, \beta, Z_j)$, $j = 1, \dots, m$. The determination of the optimal block replacement interval t_B^* requires evaluating the expected cost

$$E[C(t_B) | D, Z^{(m)}] = \int \sum_{j=1}^m \frac{c_P + c_F H_j(t_B | \theta_j, \beta, Z_j)}{t_B} p(\theta^{(m)}, \beta | D) d\theta^{(m)} d\beta. \quad (44)$$

The multi-dimensional integral in (44) can be approximated using an MC average as in (42) using draws from the posterior distribution $p(\theta^{(m)}, \beta | D)$. Note that the MC average involves evaluation of the renewal functions $H_j(t_B | \theta_j, \beta, Z_j)$, $j = 1, \dots, m$, for each posterior draw $(\theta_1^s, \dots, \theta_m^s, \beta^s)$. In Merrick et al. (2003), the renewal functions were estimated by simulating from the (conditionally) independent renewal processes and the optimal nonparametric Bayesian block replacement interval was obtained by minimizing the MC approximation to (44).

A similar nonparametric approach was considered in Merrick et al. (2005) to relax parametric assumptions in the MPPM (38) and model heterogeneity in rail tracks for developing replacement policies. More specifically, the authors specified a baseline cumulative intensity $\Lambda_0(t, \theta_i)$ in (38). The unknown distribution G of the θ_i 's is assumed to follow a Dirichlet process as in (40) while β is treated parametrically. The resulting semiparametric MPPM was used to develop nonparametric Bayesian block replacement policies with MR. In their development, Merrick et al. (2005) assumed a power law for $\Lambda_0(t, \theta_i)$, but the proposed approach can be implemented for other baseline cumulative intensity functions.

More recently, Merrick and Soyer (2017) considered an alternative semiparametric MPPM where the cumulative baseline intensity function $\Lambda_0(t)$ in (38) follows a gamma process prior

$$\Lambda_0(t) \sim G(c\Lambda_0^*(t), c), \quad (45)$$

where $\Lambda_0^*(t)$ is the mean function of the process and c is the precision parameter; see Kuo and Ghosh (2001). The nonparametric Bayesian analysis in their application was complicated by the fact that the available rail track failure data were interval censored. This implies that the posterior for $\Lambda_0(t)$ is not a gamma process and updating is not straightforward. Therefore, the authors developed a Gibbs sampler with a data augmentation step to be able to draw samples from the posterior distribution of $\Lambda_0(t)$ and obtained optimal block replacement intervals under the assumption of MR.

Damien et al. (2007) proposed a semiparametric Bayesian approach for developing optimal maintenance strategies for nuclear power plants. An important aspect of their development is the consideration of both preventive (replacement) and corrective (repair) maintenance. They also take into account down time costs. The maintenance policy involves replacing the item every T units of time and minimally repairing it upon failure at a cost c_m or c_d ($c_d > c_m$) depending on whether the failure causes disruption of power generation. The authors use a Bayesian semiparametric *accelerated failure time* (AFT) model assuming a *mixture of Polya tree* prior for the baseline reliability function $\bar{F}(t)$. The posterior Bayesian analysis is developed using MCMC methods and optimal replacement intervals are obtained. An alternative approach to nonparametric Bayesian failure rate modeling for nuclear power plants is introduced in Belyi et al. (2017) who considered the use of extended gamma processes for the failure rate functions.

3.3 Preventive maintenance via semi-Markov decision processes

Continuous Time Homogeneous semi-Markov Processes (CTHSMP or, simply, SMP) have been extensively used to model many real and theoretical cases, as outlined in Cinlar (1975) or Howard (2007). Their flexibility and reasonably tractable mathematical properties allow for their use in a wide range of applications as a powerful modeling tool. A particularly important applied area refers to maintenance decisions for HW systems, which has been pervaded by Continuous Time Markov Chain (CTMC) models, see Cano et al. (2010). SMPs provide additional flexibility due to more realistic modeling assumptions yet requiring a reasonable computational effort. However, contributions to SMPs tend to focus on probabilistic aspects, with relatively few references devoted to inferential issues, see Bhat and Miller (2002) for a description. Moreover, such inferential work tends to focus on classical approaches, with comparatively little attention paid to Bayesian methods, which, incidentally, have addressed mainly long-term equilibrium properties, see e.g. Marin et al. (2005). Flowgraph models have proved to be an important tool to deal with SMPs, see Huzurbazar (2005).

We focus on the use of Markov and, specially, semi-Markov decision processes (SMDP) in reliability and maintenance. White (1993) provides a review of MDP applications mentioning numerous ones in maintenance. More recent examples are in Chen and Trivedi (2005) and Huang and Guo (2011) referring to SMDPs. We place special emphasis on the propagation of uncertainty of the involved parameters, and how it affects the computation of various quantities relevant for decision making purposes.

We assume that the system evolves according to a SMP $\{X_t\}_{t \in T}$ with discrete state space $\{1, \dots, m\}$, with $\{1, \dots, l\}$ corresponding to ON states and $\{l + 1, \dots, m\}$ to OFF states. The parameters of the SMP are $(\boldsymbol{\nu}, P)$ where $\boldsymbol{\nu} = (\nu_1, \dots, \nu_m)^T$ are the parameters of the sojourn times $T_i, i = 1, \dots, m$ and $P = (p_{ij}), i, j = 1, \dots, m$, where p_{ij} is the transition probability from state i to state j with $\sum_{j=1}^m p_{ij} = 1$, and $p_{ii} = 0, \forall i$. Inference in SMPs has frequently been dealt from a classical perspective, neglecting the relevance of prior knowledge, and/or the uncertainty in parameters, and consequently, in predictions. Besides, SMPs have been often restricted to provide long-term results, as the transient period of the process typically entails additional computational complexity depending on the precision of the posterior distributions of the system parameters. Rios Insua et al. (2012) provide various alternative computational strategies.

We outline here decision making with SMPs under a Bayesian perspective. We focus the discussion on complex maintenance cases where we need to go beyond comparing several alternatives through discrete event simulation and pairwise comparison or ranking and selection methods, see Henderson et al. (2006) or Rios Insua et al. (2012). Assume that when entering state i , a decision maker (DM) will choose a maintenance action a from a (finite) space \mathcal{A}_i of alternatives, which may depend on state i . The system remains there for a sojourn time T_{ia} , with parameter ν_{ia} , which depends on the state and the decision a made. Upon leaving the state, the system will move to state j with probability $p_{ija} \geq 0$, with $\sum_{j=1}^m p_{ija} = 1$ and $p_{iia} = 0$. We assume that we have a prior distribution on the parameters $\boldsymbol{\nu}_a = (\nu_{ia})$ and $P_a = (p_{ija})$, for each action a , leading to a posterior $p(\boldsymbol{\nu}_a, P_a | D)$ over the parameters given the observed data D . For each maintenance decision made, we get a (possibly multi-objective) consequence $c(i, a, t_{ia})$ which depends on the time t_{ia} spent at state i and the action a . We evaluate the consequences with a utility function $u(c(i, a, t_{ia}))$, which might account for time effects such as discounting.

Consider the problem in which we manage the system until a time T has elapsed. Let $\mathbf{a} = (a_1, \dots, a_M)$ be the policy or sequence of maintenance actions that the DM adopts until time T ; $\boldsymbol{\tau} = (t_1, \dots, t_M)$, the sequence of times spent at various states visited; $\mathbf{x} = (x_1, \dots, x_M)$, the sequence of states visited, and $T_{M-1} = \sum_{i=1}^{M-1} t_i$. The utility globally obtained will be $u(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau})$. The

evolution of the system is described through

$$p_{\mathbf{a}}(\boldsymbol{\tau}, \mathbf{x} | \boldsymbol{\nu}_a, P_a) = \left[\prod_{i=1}^M p_{x_i x_{i+1} a_i} \right] \times \left[\prod_{i=1}^{M-1} f(t_{x_i} | \boldsymbol{\nu}_{x_i a_i}) \right] \times [1 - F(T - T_{M-1} | \boldsymbol{\nu}_{x_M a_M})].$$

The standard SMDP formulation would fix the $(\boldsymbol{\nu}, P)$ parameters at certain estimates $(\hat{\boldsymbol{\nu}}, \hat{P})$ and find the decisions \mathbf{a} providing maximum expected utility through

$$\max_{\mathbf{a}} \iint_{\sum t_i = T} u(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau}) p_{\mathbf{a}}(\boldsymbol{\tau}, \mathbf{x} | \hat{\boldsymbol{\nu}}_a, \hat{P}_a) d\mathbf{x} d\boldsymbol{\tau}.$$

Usually, it is assumed that the utility function is separable, that is, $u(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau}) = \sum_{i=1}^M u(x_i, a_i, t_i)$. Under such conditions, the solution to SMDPs with fixed parameters typically proceeds through some variant of dynamic programming, see Howard (2007), French and Rios Insua ((2000) or Ross (1992). For different approaches based on Q -learning, see Polson and Sorensen (2011) and references therein.

We actually do not require separability. For technical reasons, without loss of generality, assume that the utility function is positive. Then we should aim at solving

$$\max_{\mathbf{a}} \iiint \int_{\sum t_i = T} u(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau}) p_{\mathbf{a}}(\boldsymbol{\tau}, \mathbf{x} | \boldsymbol{\nu}_a, P_a) p(\boldsymbol{\nu}_a | data) p(P_a | data) d\mathbf{x} d\boldsymbol{\tau} d\boldsymbol{\nu}_a dP_a. \quad (46)$$

To evaluate the integral in (46) we may use APS as in the optimal design problem of Section 2.2.1 and define an auxiliary distribution with density

$$g(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau}, \boldsymbol{\nu}_a, P_a) \propto u(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau}) p_{\mathbf{a}}(\boldsymbol{\tau}, \mathbf{x} | \boldsymbol{\nu}_a, P_a) p(\boldsymbol{\nu}_a | D) p(P_a | D)$$

by treating the decision variable \mathbf{a} (the maintenance plan) as a random quantity. As previously discussed in Section 2.2.1, the optimal solution \mathbf{a}^* of (46) can be described as the mode of the marginal distribution of \mathbf{a} implied by the auxiliary distribution $g(\mathbf{x}, \mathbf{a}, \boldsymbol{\tau}, \boldsymbol{\nu}_a, P_a)$. Therefore, we could implement the approach described in Algorithm 1, with $\mathbf{h} = (\mathbf{x}, \mathbf{a}, \boldsymbol{\tau})$, to approximate the mode of the marginal in \mathbf{a} and, consequently, the optimal policy.

When posteriors for $(\boldsymbol{\nu}, P)$ are precise, as in the standard SMDP approach, we could suppress steps 2.2 and 2.3 and substitute them by $\hat{\boldsymbol{\nu}}, \hat{P}$ throughout. Step 3 collects the sampled decision after convergence is detected, possibly with procedures such as thinning to mitigate serial correlation. Step 4 finds the modes on the marginal distribution on decisions. Its structure will depend on problem specificities, typically requiring a discretization of policies over time.

Moreno et al. (2003) provide a related approach when $(\boldsymbol{\nu}, P)$ are fixed. Hoffman et al. (2009) provide an APS approach for Markov decision processes. The convergence of Algorithm 1 follows arguments similar to those in Bielza et al. (1999). An alternative approach may be based on using a discrete event simulation, combined with an optimization algorithm.

4 Adversarial Settings

As discussed in previous sections, reliability analysis deals with predicting how long a system will be functioning under given operational conditions and involves making decisions in relation with the maintenance, replacement, performance, design or redesign of such system. Many reliability problems, as those presented earlier, involve a single decision maker and may be appropriately dealt with at large through decision theoretic methods. However, there are reliability issues that may

Algorithm 1 Decision making with SMDP

1. Start from arbitrary values $(\mathbf{x}^0, \mathbf{a}^0, \boldsymbol{\tau}^0, \boldsymbol{\nu}_a^0, P_a^0) = (\mathbf{h}^0, \boldsymbol{\nu}_a^0, P_a^0)$.

Set $i = 0$.

Compute $u(\mathbf{h}^0)$.

2. Until convergence, iterate through

- 2.1. Generate a history \mathbf{h}^c from distribution $q_1(\cdot|\mathbf{h}^i)$

Compute

$$\gamma_1 = \min \left\{ 1, \frac{u(\mathbf{x}^c, \mathbf{a}^c, \boldsymbol{\tau}^c) p_{\mathbf{a}^c}(\boldsymbol{\tau}^c, \mathbf{x}^c | \boldsymbol{\nu}_a^i, P_a^i) q_1(\mathbf{h}^c | \mathbf{h}^i)}{u(\mathbf{x}^i, \mathbf{a}^i, \boldsymbol{\tau}^i) p_{\mathbf{a}^i}(\boldsymbol{\tau}^i, \mathbf{x}^i | \boldsymbol{\nu}_a^i, P_a^i) q_1(\mathbf{h}^i | \mathbf{h}^c)} \right\}.$$

Do

$$\mathbf{h}^{i+1} = \begin{cases} \mathbf{h}^c, & \text{with probability } \gamma_1, \\ \mathbf{h}^i, & \text{with probability } (1 - \gamma_1). \end{cases}$$

- 2.2. Generate P_a^c from distribution $q_2(\cdot|P_a^i)$

Compute

$$\gamma_2 = \min \left\{ 1, \frac{p_{\mathbf{a}^{i+1}}(\boldsymbol{\tau}^{i+1}, \mathbf{x}^{i+1} | \boldsymbol{\nu}_a^i, P_a^i) p(P_a^c | \text{data}) q_2(P_a^c | P_a^i)}{p_{\mathbf{a}^{i+1}}(\boldsymbol{\tau}^{i+1}, \mathbf{x}^{i+1} | \boldsymbol{\nu}_a^i, P_a^i) p(P_a^i | \text{data}) q_2(P_a^i | P_a^c)} \right\}.$$

Do

$$P_a^{i+1} = \begin{cases} P_a^c, & \text{with probability } \gamma_2, \\ P_a^i, & \text{with probability } (1 - \gamma_2). \end{cases}$$

- 2.3. Generate $\boldsymbol{\nu}_a^c$ from distribution $q_3(\cdot|\boldsymbol{\nu}_a^i)$

Compute

$$\gamma_3 = \min \left\{ 1, \frac{p_{\mathbf{a}^{i+1}}(\boldsymbol{\tau}^{i+1}, \mathbf{x}^{i+1} | \boldsymbol{\nu}_a^c, P_a^{i+1}) p(\boldsymbol{\nu}_a^c | \text{data}) q_3(\boldsymbol{\nu}_a^c | \boldsymbol{\nu}_a^i)}{p_{\mathbf{a}^{i+1}}(\boldsymbol{\tau}^{i+1}, \mathbf{x}^{i+1} | \boldsymbol{\nu}_a^i, P_a^i) p(\boldsymbol{\nu}_a^i | \text{data}) q_3(\boldsymbol{\nu}_a^i | \boldsymbol{\nu}_a^c)} \right\}.$$

Do

$$\boldsymbol{\nu}_a^{i+1} = \begin{cases} \boldsymbol{\nu}_a^c, & \text{with probability } \gamma_3, \\ \boldsymbol{\nu}_a^i, & \text{with probability } (1 - \gamma_3). \end{cases}$$

3. Once convergence is detected (say after iteration k), collect the next n sampled decisions $(\mathbf{a}^{k+1}, \dots, \mathbf{a}^{k+n})$.
 4. Use the sample $\{\mathbf{a}^{k+1}, \dots, \mathbf{a}^{k+n}\}$ to approximate the mode of the marginal distribution on decisions.
-

involve two or more actors with competing interests. Examples of adversarial situations in this field can be found in areas such as acceptance sampling (Lindley and Singpurwalla, 1991), life testing (Lindley and Singpurwalla, 1993), reliability demonstration (Rufo et al., 2014) and warranty analysis (Singpurwalla and Wilson, 1993).

These problems with adversarial components can be set up as games and are typically solved using game theory methods; see for example Gibbons (1992). A main drawback of such methodology in this application area is its underlying common knowledge assumption, assessed in e.g. Raiffa et al.(2002) or Lippman and McCardle (2012). Therefore, we shall also illustrate alternative approaches based on the recent framework of adversarial risk analysis (ARA), see Banks et al.(2015). We first present adversarial issues around warranty policies and then analyse comparatively the game theoretic and ARA approaches in relation with acceptance sampling.

4.1 Warranty policies

A warranty is an agreement made between the buyer and the seller of a good or service. In the case of a product, the seller agrees to rectify any fault through repair or replacement. This is for a certain amount of time, or usage, or both time and usage, after purchase. In the case of a service, a service

level agreement may be signed and the warranty compensates the buyer if the agreed level has not been maintained; this is again typically for a defined amount of time or other measure of use from purchase.

Aside from the legal aspects (Priest, 1981), from the consumer's perspective, a warranty acts as a guarantee of quality. From the seller's perspective, it is an important tool to attract and retain customers. The decision and game theory aspects of this problem are usually explored from the point of view of the seller, who is faced with determining the best type of warranty for the product, given knowledge about its reliability, the costs of repair, the relationship between sales and warranty size and the actions of its competitors.

Early reviews on warranties from a mathematical perspective are in Blischke (1990), Murthy (1990) and Blischke and Murthy (1991). The latter also looked at variants of the standard warranty, such as limited warranties, where repair costs are shared between buyer and seller, and extended warranties, where the consumer can pay to prolong the warranty period. Singpurwalla and Wilson (1994) looked at several game and decision-theoretic aspects, including warranties based on both time and usage.

The time-limited warranty, where the warranty is valid for a time t after purchase with no limit on usage, has received most of the attention in the literature. A decision-theoretic approach to the seller's warranty problem attempts to find the optimal warranty time t^* . The seller has a model for the number $N(t)$ of claims that will be made against a product by time t , quantified by a probability model $p_t(n|\theta) = P(N(t) = n|\theta)$. Typical models are point processes, such as the Poisson process, or a reliability model for time to failure can be defined, and then the number of failures is the corresponding renewal process if independent and identically distributed failure times are assumed (Ross, 1995). A prior distribution for the seller's uncertainty about θ may be assessed, which then can be updated with warranty claim data. As regards the utility, Blischke (1990) suggested a cost model for the seller for a single unit of a product with a warranty to time t from purchase of the form $S(1 + N(t))$, if the item is non-repairable (and so must be replaced if it fails), or $S + CN(t)$, if the item is repairable, where S is the cost of making and delivering the product to the customer and C is the average repair cost. A more comprehensive cost model would take into account the attractiveness to the consumer of a product at a price p with warranty length t . Written as a utility, this could then be of the form

$$u(N(t), \theta, (p, t)) = \pi(p, t) \left(p - S - \sum_{i=1}^{N(t; \theta)} C_i \right), \quad (47)$$

where $\pi(p, t)$ measures the probability that a consumer buys the product at a price p and warranty length t (a non-decreasing function of t) and C_i is the cost of resolving the i th claim (either a repair or replacement cost). The expected utility, assuming uncertainty over θ as well as $N(t)$, is then:

$$\bar{u}(t) = \pi(p, t) \left(p - S - \sum_n \left(\sum_{i=1}^n C_i \right) \int_{\theta} p_t(n|\theta) p(\theta) d\theta \right),$$

and then we look for $t^* = \arg \max_t \bar{u}(t)$. Simultaneous maximization over p and t is also possible. If the claim costs are identically distributed, then one can take the expectation over the claim costs to arrive at a simplified expression

$$\bar{u}(t) = \pi(p, t) \left(p - S - \sum_n n E(C) \int_{\theta} p_t(n|\theta) p(\theta) d\theta \right) = \pi(p, t) (p - S - E(C) E(N(t))).$$

The analysis of warranty costs has been well documented in the literature, e.g. the recent work

by Liu et al. (2015) or the review by Wu (2012), and so typically the seller has good information on the distribution of the C_i 's. The consumer response to the warranty size, appearing as the function $f(t)$ in (47), is more difficult to quantify. Jindal (2015) looked at what drives the purchase of extended warranties and derives a utility from the consumer's perspective. Chu and Chintagunta (2009) is more useful as it looks at a demand model for warranties from the seller's point of view.

Figure 5 illustrates the approach with a simple example, using the utility function in (47) with $p = 2$, $S = 1$ and $E(C) = 0.1$. A logistic function is used for the probability of purchase, $\pi(p, t) = 0.5 \exp(-0.01p + (t - 2)) / (1 + \exp(-0.01p + (t - 2)))$, so that the maximum chance of purchase is 0.5, and $N(t)$ is modelled as a Poisson process with mean value function $E(N(t)) = (1 + t)^\theta - 1$. Note that $\theta = 1$ corresponds to the homogeneous Poisson process, and so the times between successive claims have the same distribution, as might be expected if the claim was resolved by a replacement product or a repair that returns the product to an as-new state. If $\theta > 1$, there is an increasing rate of claims as the product ages, corresponding to a situation where a claim is resolved by an imperfect repair and there is an increasing rate of claims with time. The figure shows the optimal warranty period for fixed θ , which shortens as θ increases, reflecting the fact that the seller wishes to avoid the increasing claim rate that is implied by $\theta > 1$.

Two-dimensional warranties, where the warranty period covers up to a time t and usage m are common in some markets e.g. automobiles. The optimal warranty specification now looks to define the optimal pair (t^*, m^*) . One approach is to map (t, m) to a single index and optimise with respect to it; Oakes (1995) or Gertsbakh and Kordonsky (1998) provide examples. Alternatively, a bivariate model for time and usage can be defined. Mercer (1961) was an early attempt to do this. Eliashberg et al. (1997) modelled usage as a logistic function of time, while Singpurwalla and Wilson (1998) defined a model where usage was a stochastic process indexed by time. In the latter, the model was applied to the optimal two-dimensional warranty problem. More recently, Su and Wang (2016) proposed a bivariate Weibull model for time and usage to failure with survival function

$$R(t, m) = \exp \left(- \left\{ \left(\frac{t}{\alpha_T} \right)^{\beta_T/\delta} + \left(\frac{m}{\alpha_M} \right)^{\beta_M/\delta} \right\}^\delta \right),$$

where δ models their dependence. Two-dimensional renewal theory can be used to derive a distribution for $N(t, m)$, the number of claims by time t and usage m (Hunter, 1974). Singpurwalla and Wilson (1994) suggested a seller's utility function for a two-dimensional warranty that is an extension of (47) but where now the probability of purchase and number of claims depend on both time and usage,

$$u(N(t, m), \theta, (t, m)) = \pi(p, t, m) \left(p - S - \sum_{i=1}^{N(t, m; \theta)} C_i \right).$$

Then (t^*, m^*) are found by taking the expected utility with respect to $N(t, m)$ and θ , then maximising with respect to t and m .

Game theoretic aspects of warranties have also been explored. Emons (1988) looked at the moral hazard problem, where buyers take less care of a product if it has a warranty, leading to an adversarial interaction between the buyer and seller. The seller must assess the reduction in reliability due to the buyer's actions. Singpurwalla and Wilson (1994) discuss the interaction between a consumer and a producer by making the probability $\pi(p, t, m)$ of (47) depend on the buyer's utility for purchasing the product. All these require common knowledge assumptions that are then used to compute Nash equilibria and related refinements. We explore this and alternative approaches in the related problem of adversarial acceptance sampling.




Figure 5: Optimizing warranty length using the utility function of (47). Clockwise from top left: expected number of claims as a function of θ , probability of purchase, expected utility as a function of time for 4 values of θ , the optimal warranty length as a function of θ .

4.2 Adversarial acceptance sampling problems

We first formulate the general adversarial problem for acceptance sampling as in Lindley and Singpurwalla (1991) who deal with it from a game-theoretic perspective. As a motivation, consider a case with a manufacturer M (she) and a consumer C (he). In order to convince the consumer, the manufacturer offers sample products to the consumer who, based on the perceived quality, places an order or not. Thus, M needs to decide about n , the sample size that will be offered to the consumer. The outcome of the inspection is described through the data D , which depends on a parameter θ characterising the product quality, say the failure rate in a life testing context. The decision of C refers to accept (A) or reject (R) the batch, which will be based on the observed sample data D , eventually used by C to revise his uncertainty about θ . The utility function of C depends on his decision and the quality parameter θ . The utility function of M depends on her decision, the consumer's decision and the quality parameter θ . The problem may be described through a bi-agent influence diagram (BAID), Koller and Milch (2003) or Banks et al. (2015), represented in Figure 6, where round nodes

refer to random events, square nodes refer to decisions and exagonal nodes refer to evaluations. White nodes belong to M , grey nodes to C and striped nodes are shared by both agents. Arrows are interpreted as in standard influence diagrams.




Figure 6: Acceptance sampling. A Bi-Agent Influence Diagram.

We assume we are supporting the manufacturer. We illustrate first the game theoretic approach and then the ARA perspective.

For the game theoretic solution, due to the sequentiality of the decisions, the consumer sees the manufacturer decision and she does not need his judgements. But the manufacturer needs to know the consumer beliefs and preferences, which is the common knowledge condition in this problem. Then, the consumer should have available (and these should be available to the manufacturer) $p_C(\theta)$, which describes his beliefs about the product quality θ ; $p_C(d|\theta, n)$, which describes his beliefs about the experiment results d given the quality θ and the decision n of M ; and his utility function, $u_C(c, \theta)$. If so, he proceeds, for each d and n , by:

1. Inverting the arc θ - D and computing, by Bayes' formula,

$$p_C(\theta|d, n) \propto p_C(\theta)p_C(d|\theta, n). \quad (48)$$

2. Computing the expected utilities,

$$\psi_C(n, d, c) = \int u_C(c, \theta)p_C(\theta|d, n)d\theta.$$

3. Computing the optimal decision c , given d and n ,

$$c^*(d, n) = \arg \max_{c \in \{\mathcal{A}, \mathcal{R}\}} \psi_C(d, n, c).$$

The manufacturer knows this and, consequently, solves her problem. For this, she should have available $p_M(\theta)$, describing her beliefs about the quality θ ; $p_M(d|\theta, n)$, reflecting her beliefs about the experiment results given the quality θ and his decision n ; and her utility function $u_M(c, n, \theta)$. In this case, the manufacturer proceeds by:

1. Assessing the utilities of the attained results (the consumer decision)

$$\psi_M(n, d, \theta) = u_M(c^*(d, n), n, \theta).$$

2. Computing the expected utilities,

$$\psi_M(n, \theta) = \int \psi_M(n, d, \theta)p_M(d|\theta, n) dd.$$

3. Computing the expected utilities,

$$\psi_M(n) = \int \psi_M(n, \theta) p_M(\theta) d\theta.$$

4. Finally, computing her optimal decision through

$$n^* = \arg \max \psi_M(n).$$

Then n^* and $\{c^*(d, n^*)\}_d$ is a subgame perfect equilibrium, Gibbons (1992).

We move now to the ARA approach, avoiding common knowledge assumptions. Basically, we run a decision analysis problem for the manufacturer, which requires input coming from a simulation over the consumer's problem. We thus consider first the manufacturer's decision problem. To solve the problem, M should have available, as before, $p_M(\theta)$, $p_M(d|\theta, n)$ and $u_M(c, n, \theta)$, but also $p_M(c|d, n)$, which describes her beliefs about the customer decision c (accept, reject) given the experiment results d and her decision n . Then, the manufacturer proceeds by:

1. Computing the expected utilities,

$$\psi_M(n, d, \theta) = \sum_{c \in \{\mathcal{A}, \mathcal{R}\}} u_M(c, n, \theta) p_M(c|d, n).$$

2. Computing the expected utilities,

$$\psi_M(n, \theta) = \int \psi_M(n, d, \theta) p_M(d|\theta, n) dd.$$

3. Computing the expected utilities, to reduce node Θ ,

$$\psi_M(n) = \int \psi_M(n, \theta) p_M(\theta) d\theta.$$

4. Finally, computing her optimal decision through

$$n^* = \arg \max \psi_M(n).$$

However, in the above, $p_M(c|d, n)$ is nonstandard, since it entails strategic elements about the behavior of the consumer. To facilitate its assessment, we may actually consider his problem and simulate from it.

The consumer problem is shown in Figure 7. To solve it, the consumer should have available, as




Figure 7: Acceptance sampling. Consumer vision.

stated above, $p_C(\theta)$, $p_C(d|\theta, n)$, and $u_C(c, \theta)$. He does not need $p_C(n)$, since his decision is contingent upon seeing the manufacturer decision n at M . In this case, the consumer proceeds, for each d and n , as above, not being necessary to remove nodes D and M from the consumer ID. However, since we do not assume the common knowledge condition, we do not have the required elements $u_C(c, \theta)$, $p_C(\theta)$, $p_C(d|\theta, n)$. We may model our uncertainty around them, through random utilities and probabilities, which we designate by $F = (U_C(c, \theta), P_C(\theta), P_C(d|\theta, n))$. Since we do not need the denominator in Bayes' formula (48) for optimisation purposes, we may actually proceed, for each d and n , by

2' Computing the random functional

$$\Psi_C^*(n, d, c) = \int U_C(c, \theta) P_C(\theta) P_C(d|\theta, n) d\theta.$$

3' Computing the random optimal alternative, given d and n ,

$$C^*(d, n) = \arg \max_{c \in \{A, R\}} \Psi_C^*(d, n, c).$$

We, then, set $p_M(c|d, n) = Pr(C^*(d, n) = c)$ which feed into the manufacturer's problem.

Estimation of $C^*(d, n)$ would typically proceed by MC simulation, by sampling from the random utilities and probabilities, computing the corresponding optimal decisions, and then estimating by the MC frequencies. In general, $P_C(\theta)$, $P_C(d|\theta, n)$ could be based on $p_M(\theta)$, $p_M(d|\theta, n)$ with some uncertainty around them. In discrete cases, these could be Dirichlet distributions, whereas, in continuous cases, these could be Dirichlet processes. With regards to $U_C(c, \theta)$, we may have information about the consumer interests and use a parametric form for the utility function. Finally, we derive a distribution over the parameters. Banks et al. (2015) provide details on such assessments. Rios Insua et al. (2018) provide examples in Bernoulli acceptance sampling and exponential life testing.

5 Concluding Remarks

This article provides a review of Bayesian methods for making decisions in reliability. Decision making is a major aspect of reliability analysis which is frequently neglected by statisticians and applied probabilists whose main focus is often on modeling and inference. It is desirable to integrate decision making and statistics in many reliability problems and the Bayesian paradigm provides a coherent framework to do this. The primary objective of this review was to illustrate how the Bayesian decision-theoretic approach is applied in a variety of problems such as life testing, design of experiments, reliability certification, preventive maintenance, warranties and acceptance sampling and discuss recent advances in these areas.

Section 2 has presented Bayesian design of life tests including accelerated tests, stopping rules and reliability demonstration. Due to space limitations, design of burn-in tests is excluded. In reliability, burn-in testing is used to distinguish between "weak" and "strong" items so that weak items are eliminated and not released to consumers. One of the earlier Bayesian papers in burn-in testing design is by Clarotti and Spizzichino (1990). A more recent work is by Perlstein et al. (2001). Reliability based optimal design is another area which is not discussed since most of this work appear in engineering journals and the Bayesian approach has been considered in comparatively few articles, such as Gunawan and Papalambros (2006). Our discussion of preventive maintenance in Section 3 has not included condition-based maintenance policies which is an area also with limited Bayesian work; see for example, the recent review paper by Olde-Keizer et al. (2017). Another area where the development of Bayesian maintenance policies has not been considered is the *virtual age* or imperfect repair models of Zhang et al. (2015).

Besides its focus on Bayesian decision making, this article differs from the other reviews in reliability due to its consideration of adversarial issues in reliability discussed in Section 4. As noted by Rios-Insua et al. (2018), the ARA approach discussed here can be extended to other areas such as warranty policies and software testing which include more than two adversaries. Another potential area for extension is adversarial life testing with multiple stages. Solution of such sequential testing is challenging and will require the development of new computational methods. Last, but not least, it would be relevant to consider security issues in relation with reliability when adversaries attempt to reduce the reliability of our managed systems.

Acknowledgements

The work of DRI is supported by the Spanish Ministry of Economy and Innovation program MTM2014-56949-C3-1-Rand the AXA-ICMAT Chair on Adversarial Risk Analysis. DRI and FR acknowledge the support from the ESF-COST Action IS1304 on Expert Judgement. This work has also been partially supported by the Spanish Ministry of Economy and Competitiveness, through the “Severo Ochoa” Programme for Centres of Excellence in R&D (SEV-2015-0554). FR was supported by a Catedra de Excelencia from the Government of Madrid.

References

- Antoniak, C. E. (1974). Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Annals of Statistics*, Vol. 2, pp. 1152–1174.
- Arias, P., Martin, J., Ruggeri, F. and Suarez-Llorens, A. (2015). A robust Bayesian approach to optimal replacement policy for gas pipelines. *Entropy*, Vol. 17, pp. 3656–3678.
- Arunkumar, S. (1972). Nonparametric age replacement policy. *Sankhya Series A*, Vol. 34, pp. 251–256.
- Bancroft, G. A. and Dunsmore, I. R. (1978). Predictive sequential life testing. *Biometrika* 65, 3: 609–613.
- Banks D., Rios J. and Ríos Insua D. (2015). *Adversarial Risk Analysis*, CRC Press.
- Barlow R. E. and Hunter L. C. (1960). Optimum preventive maintenance policies. *Operations Research*, Vol. 8, pp. 90–100.
- Barlow R. E., Proschan, F. (1965). *Mathematical Theory of Reliability*. New York: John Wiley & Sons.
- Barlow R. E. and Hsiung, J. H. (1983). Expected information from a life test experiment. *The Statistician*, Vol. 32, pp. 35–45.
- Barnett, V.D. (1972). Bayesian sequential life test. *Technometrics* 14, 2: 453–468.
- Bartholomew, D. J. (1963). The sampling distribution of an estimate arising in life testing. *Technometrics*, Vol. 5, pp. 361–374.
- Bassin, W. M. (1973). A Bayesian optimal overhaul interval model for the Weibull restoration process. *Journal of the American Statistical Association*, Vol. 68, pp. 575–578.
- Belyi, D., Popova, E., Morton, D. and Damien, P. (2017). Bayesian failure-rate modeling and preventive maintenance optimization. *European Journal of Operational Research*, Vol. 262, pp. 1085–1093.

- Bernardo, J. M. (1997). Statistical inference as a decision problem: the choice of sample size. *The Statistician*, Vol. 46, pp. 151-153.
- Bhat, U.N. and Miller, G.K. (2002) *Elements of Applied Stochastic Processes*. Wiley.
- Bielza, C., Muller, P. and Insua-Rios, D. (1999). Decision analysis by augmented probability simulation. *Management Science*, Vol. 45, pp. 995-1007.
- Blischke, W.R. (1990). Mathematical models for analysis of warranty policies. *Mathematical and Computer Modelling*, Vol. 13, pp. 1-16.
- Blischke, W.R. and Murthy, D.N.P. (1991). Product warranty management I: a taxonomy for warranty policies. Technical Report, University of Southern California, Department of Decision Systems, Los Angeles, California.
- Brooks, R. J. (1982). On the loss of information through censoring. *Biometrika*, Vol. 69, pp. 137-44.
- Cagno, E., Caron, F., Mancini, M. and Ruggeri, F. (2000). Using AHP in determining the prior distributions on gas pipeline failures in a robust Bayesian approach, *Reliability Engineering and System Safety*, Vol. 67, pp. 275-284.
- Cano, J., Moguerza, J., and Rios Insua, D. (2010). Bayesian reliability, availability, and maintainability analysis for hardware systems described through continuous time Markov chains. *Technometrics*, Vol. 52, pp. 324-334.
- Chaloner, K. (1984). Optimal Bayesian experimental designs for linear models. *Ann. of Statist.*, Vol. 12, pp. 283-300.
- Chaloner, K. and Larntz, K. (1992). Bayesian design for accelerated life testing. *Jrnl. of Statist. Plan. & Inf*, Vol. 33, pp. 245-259.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, Vol. 10, pp. 273-304.
- Chen, D., Trivedi, K. (2005) Optimization for condition-based maintenance with semi-Markov decision process. *Reliability Engineering and Systems Safety*, Vol. 90, pp 25-29.
- Chernoff, H. (1962). Accelerated life designs for estimation. *Technometrics*, Vol. 4, pp.381-408.
- Chien, Y. H. (2008). A general age-replacement model with minimal repair under renewing free-replacement warranty. *European Journal of Operational Research*, Vol. 186, pp. 1046-1058.
- Cho, D. I. and Parlar, M. (1991). A survey of maintenance models for multi-unit systems. *European Journal of Operational Research*, Vol. 51, pp. 1-23.
- Chu, J. and Chintagunta, P.K. (2009). Quantifying the economic value of warranties in the U.S. server market. *Journal of Marketing Science*, Vol. 28, pp. 99-121.
- Cinlar, E (1975) *Introduction to Stochastic Processes*. Prentice Hall.
- Clarotti, C. A. and Spizzichino, F. (1990). Bayes burn-in decision procedures. *Probability in the Engineering and Informational Sciences*, Vol. 4, pp. 437-445.

- Cox, D. R. (1962). *Renewal Theory*. London: Methuen.
- Cox, D. R. (1972a). The statistical analysis of dependencies in point processes. In *Stochastic Point Processes*. Ed. P. A. W. Lewis, pp. 55-66, New York Wiley.
- Cox, D. R. (1972b). Regression models and life tables. *J. R. Statist. Soc. B*, Vol. 34, pp. 187-220.
- Damien, P., Galenko, A., Popova, E., and Hanson, T. (2007). Bayesian semiparametric analysis for a single item maintenance optimization. *European Journal of Operational Research*, Vol. 182, pp. 794-805.
- Dayanik, S. and Gurler, U. (2002). An adaptive Bayesian replacement policy with minimal repair. *Operations Research*, Vol. 50, pp. 552-558.
- Deely, J. J. and Keats, J. B. (1994). Bayes stopping rules for reliability testing with the exponential distribution. *IEEE Transactions on Reliability*, Vol. 43, pp. 288-293.
- DeGroot, M. H. and Goel, P. (1979). Bayesian estimation and optimal designs in partially accelerated life testing. *Naval Res. Logistics Quarterly*, Vol. 26, pp. 223-235.
- Dellaportas, P. and Smith, A. F. M. (1993). Bayesian inference for generalized linear and proportional hazards models via Gibbs sampling. *Appl. Statist.*, 42, No. 3, pp. 443-459.
- van Dorp, J. R., Mazzuchi, T. A and Soyer, R. (1997) Sequential inference and decision making during product development. *Journal of Statistical Planning and Inference*, Vol. 62, pp. 207-218.
- Doyen, L. and Gaudoin, O. (2004). Classes of imperfect repair models based on reduction of failure intensity or virtual age. *Reliability Engineering and System Safety*, Vol. 84, pp. 45-56.
- Dunsmore, I. R. (1974). The Bayesian predictive distribution in life testing models. *Technometrics*, Vol. 16, pp. 455-460.
- Dunsmore, I.R. and Wright, D.E. (1985). A decisive predictive approach to the construction of sequential acceptance sampling plans for lifetimes. *Journal of the Royal Statistical Society, Series C*, Vol. 34, pp. 1-13.
- Dykstra, R. L. and Laud, P. (1981). A Bayesian approach to reliability. *Annals of Statistics*, Vol. 9, pp. 356-367.
- Eliashberg, J., Singpurwalla, N.D. and Wilson, S.P. (1997). Calculating the reserve for a time and usage indexed warranty. *Management Science*. Vol. 43, pp. 966-975.
- Emons, W. (1988). Warranties, moral hazard and the lemons problem. *Journal of Economic Theory*, Vol. 46, pp. 16-33.
- Erkanli, A. and Soyer, R. (2000). Simulation-based designs for accelerated life tests. *Jrnl. of Statist. Plan. & Inf.*, Vol. 90, pp. 335-348.
- Escobar, M. D. and West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, Vol. 90, pp. 577-588.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Annals of Statistics*, Vol. 1, pp. 209-230.

- Fox, B. L. (1967). Adaptive age replacement. *J. of Math. Anal. and Appl.*, Vol. 18, pp. 365-376.
- Frees, E.W. and Ruppert, D. (1985). Sequential nonparametric age replacement policies. *The Annals of Statistics*, Vol. 13, pp. 650-662.
- French, S., Rios Insua, D. (2000) *Statistical Decision Theory*, Wiley.
- Gertsbakh, I.B. and Kordonsky, K.N. (1998). Parallel time scales and two-dimensional manufacturer and individual customer warranties. *IIE Transactions*, Vol. 30, pp. 1181-1189.
- Gibbons R. A. (1992). *A Primer in Game Theory*. Pearson Education Ltd.
- Goel, A.L. and Okumoto, K. (1979). Time-dependent error detection rate model for software reliability and other performance measures. *IEEE Transactions on Reliability*, Vol. 28, pp. 206-211.
- Gunawan, S. and Papalambros, P. Y. (2006). A Bayesian approach to reliability-based optimization with incomplete information. *Journal of Mechanical Design*, Vol. 128, pp. 909-918.
- Gutierrez-Pena, E. and Walker, S. G. (2005). Statistical decision problems and Bayesian nonparametric methods. *International Statistical Review*, Vol. 73, pp. 309-330.
- Henderson, S. and Nelson, B. (2006). *Simulation*. Elsevier.
- Hermann, S. and Ruggeri, F. (2017). Modeling Wear Degradation in Cylinder Liners. *Quality and Reliability Engineering International*, Vol. 33, pp. 839-851.
- Higgins, J.J. and Tsokos, C.P. (1976). On the behavior of some quantities used in Bayesian reliability demonstration tests. *IEEE Transactions on Reliability*, Vol. 25, pp. 261-264.
- Hoffman, M., Kueck, H., de Freitas, N. and Doucet, A. (2009). New inference strategies for solving Markov decision processes using reversible jump MCMC. In *Proceedings of the 25th Annual Conference on Uncertainty in Artificial Intelligence (UAI-09)*, pp. 223-231. AUAI Press.
- Hong, Y., King, C., Zhang, Y. and Meeker, W. Q. (2015). Bayesian life test planning for log-location-scale family of distributions. *Journal of Quality Technology*, Vol. 47, pp. 336-350.
- Howard, R. (2007) *Dynamic Probabilistic Systems. Volume 2: Semi-Markov and Decision Processes*. Dover.
- Huang, Y., Guo, X. (2011) Finite horizon semi-Markov decision processes with applications to maintenance systems. *European Journal of Operational Research*, Vol. 212, pp 131-140.
- Hunter, J.J. (1974). Renewal theory in two dimensions: Basic results. *Advances in Applied Probability*, Vol. 6, pp. 376-391.
- Huzurbazar, A. (2005). *Flowgraph Models for Multistate Time-to-Event Data*. Wiley.
- Jin, G. and Matthews, D. (2014). Reliability demonstration for long-life products based on degradation testing and a Wiener process model. *IEEE Transactions on Reliability*, Vol. 63, pp. 781-797.
- Jindal, P. (2015). Risk preferences and demand drivers of extended warranties. *Marketing Science*, Vol. 34, pp. 39-58.

- Juang, M. G. and Anderson, G. (2004). A Bayesian method on adaptive preventive maintenance problem. *European Journal of Operational Research*, Vol. 155, pp. 455–473.
- Kalbfleisch, J. D. (1978). Non-parametric Bayes analysis of survival time data. *Journal of the Royal Statistical Society, Ser. B*, Vol. 40, pp. 214-221.
- Kraft, C. H. and van Eeden, C. (1964). Bayesian bioassay. *Annals of Mathematical Statistics*, Vol. 35, pp. 886-890.
- Koller D. and Milch B. (2003). Multi-agent influence diagrams for representing and solving games, *Games and Economic Behavior*, Vol. 45, pp. 181-221.
- Kundu, D. (2008). Bayesian inference and life testing plan for the Weibull distribution in presence of progressive censoring. *Technometrics*, Vol. 50, pp. 144-154.
- Kuo, L. and Ghosh, S. K. (2001). Bayesian nonparametric inference for nonhomogeneous Poisson processes. *Technical Report*, North Carolina State University. Institute of Statistics Mimeo Series No. 2530.
- Laud, P.W., Smith, A. F. M. and Damien, P. (1996). Monte Carlo methods for approximating a posterior hazard rate process. *Statistics and Computing*, Vol. 6, pp. 77–83.
- Lindley, D. V. (1956). On a measure of information provided by an experiment. *Annals of Statistics*, Vol. 27, pp. 986–1005.
- Lindley, D. V. (1976). A class of utility functions. *The Annals of Statistics*, Vol. 4, pp. 1-10
- Lindley, D.V. (1985). *Making Decisions* (2nd edition). London: Wiley.
- Lindley, D. V. and Singpurwalla, N. D. (1991). On the evidence needed to reach agreed action between adversaries, with application to acceptance sampling, *Journal of the American Statistical Association*, Vol. 86, pp. 933-937.
- Lindley, D. V. and Singpurwalla, N. D. (1993). Adversarial life testing, *Journal of the Royal Statistical Society, Ser. B*, Vol. 55, pp. 837-847.
- Lippman S. A. and McCardle K. F. (2012). Embedded Nash bargaining: Risk aversion and impatience, *Decision Analysis*, Vol. 9, pp. 31-40.
- Liu, B., Wu, J. and Xie, M. (2015). Cost analysis for multi-component system with failure interaction under renewing free-replacement warranty. *European Journal of Operational Research*, Vol. 243, pp. 874–882.
- Marin, J.M., Pla, L.M., and Rios Insua, D. (2005). Forecasting for some stochastic process models related to sow farm management. *Journal of Applied Statistics*, Vol. 32, pp 797 812.
- Martz, H. F. and Waterman, M. S. (1978). A Bayesian model for determining the optimal test stress for a single test unit. *Technometrics*, Vol. 20, pp. 179-185.
- Martz, H.F. and Waller, R.A. (1979). A Bayesian zero-failure (BAZE) reliability demonstration testing procedure. *Journal of Quality Technology*, Vol. 11, pp.128–138.

- Martz, H.F. and Waller, R.A. (1982). *Bayesian Reliability Analysis* New York: John Wiley & Sons.
- Mazzuchi, T. A. and Singpurwalla, N. D. (1985). A Bayesian approach to inference for monotone failure rates. *Statistics and Probability Letters*, Vol. 3, pp. 135-142.
- Mazzuchi, T. A. and Soyer, R. (1989). Assessment of machine tool reliability using a proportional hazards model. *Naval Research Logistics*, Vol. 36, pp. 765-777.
- Mazzuchi, T. A. and Soyer, R. (1996a). A Bayesian perspective on some replacement strategies. *Reliability Engineering and System Safety*, Vol. 51, pp. 295-303.
- Mazzuchi, T. A., Soyer, R., (1996b). Adaptive Bayesian replacement strategies. In *Bayesian Statistics 5*, J. O. Berger, J. M. Bernardo, A. P. Dawid, A. F. M. Smith (Eds.), pp. 667-674.
- McDaid, K. and Wilson, S.P. (2001). Deciding how long to test software. *The Statistician* 50, 2: 117-134.
- Menzefricke, U. (1991). Designing accelerated life tests when there is Type II censoring. *Commun Statist.-Theory Meth.*, Vol. 21, pp. 2569-2590.
- Merrick, J. R., Soyer, R., and Mazzuchi, T. A. (2003). A Bayesian semiparametric analysis of the reliability and maintenance of machine tools. *Technometrics*, Vol. 45, pp. 58-69.
- Merrick, J. R., Soyer, R., and Mazzuchi, T. A. (2005). Are maintenance practices for railroad tracks effective ? *Journal of the American Statistical Association*, Vol. 100, pp. 17-25.
- Merrick, J. R. and Soyer, R. (2017). Semiparametric Bayesian optimal replacement policies: application to railroad tracks. *Appl. Stochastic Models Bus. Ind.*, Vol. 33, pp. 445-460.
- Morali, N. and Soyer, R. (2003). Optimal stopping in software testing. *Naval Research Logistics*, Vol. 50, pp. 88-104.
- Moreno, A., Virto, M.A., Martin, J., and Rios Insua, D. (2003) Approximate solutions to semi Markov decision processes through Markov chain Monte Carlo methods, *EUROCAST 2003*, pp. 151-162. Springer-Verlag.
- Muller, P. (1999). Simulation-based optimal design. *Bayesian Statistics 6*, (J. O. Berger, J. M. Bernardo, A. P. Dawid and A. F. M. Smith, Eds.), pp. 459-474.
- Murthy, D.N.P. (1991). Product warranty: a review and technology management implications. *Proceedings of the Second International Conference on Management of Technology*, Miami, Florida, February 1990.
- Nasir, E. A. and Pan, R. (2015). Simulation-based Bayesian optimal ALT designs for model discrimination. *Reliability Engineering and System Safety*, Vol. 134, pp. 1-9.
- von Neumann, J. and Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Oakes, D. (1995). Multiple time scales in survival analysis. *Lifetime Data Analysis*, Vol. 1, pp. 7-18.

- Olde Keizer, M. C. A., Flapper, S. D. P. and Teunter, R. H. (2017). Condition-based maintenance policies for systems with multiple dependent components: A review. *European Journal of Operational Research*, Vol. 261, pp. 405-420.
- Perlstein, D., Jarvis, W. H. and Mazzuchi, T. A. (2001). Bayesian calculation of cost optimal burn-in test durations for mixed exponential populations. *Reliability Engineering and System Safety*, Vol. 72, pp. 265-273.
- Pievatolo, A., Ruggeri, F. and Argiento, R. (2003). Bayesian analysis and prediction of failures in underground trains, *Quality Reliability Engineering International*, Vol. 19, pp. 327-336.
- Pievatolo, A. and Ruggeri, F. (2004). Bayesian reliability analysis of complex repairable systems. *Appl. Stochastic Models Bus. Ind.*, Vol. 20, pp. 253-264.
- Pievatolo, A. and Ruggeri, F. (2010). Bayesian modeling of train doors' reliability. in *Handbook of Applied Bayesian Analysis*, A. O'Hagan and M. West Eds, Oxford Press.
- Polson, N. and Sorensen, M. (2011). A simulation based approach to stochastic dynamic programming. *Applied Stochastic Models in Business and Industry*, Vol. 27, pp. 151-163..
- Polson, N. G. and Soyer, R. (2017). Augmented probability simulation for accelerated life test design. *Appl. Stochastic Models Bus. Ind.*, Vol. 33, pp. 322-332.
- Popova, E. (2004). Basic optimality results for Bayesian group replacement policies. *Operations Research Letters*, Vol. 32, pp. 283-287.
- Pratt, J.W. (1964). Risk aversion in the small and in the large. *Econometrica*, Vol. 32, pp. 122-136.
- Priest, G. (1981). A theory of consumer product warranty. *Yale Law Journal*, Vol. 90, pp.1297-1352.
- Rahrourh, M., Coolen, F.P.A. and Collen-Schrijner, P. (2006). Bayesian reliability demonstration for systems with redundancy. *Proceedings of the Institute of Mechanical Engineers, Part O: Journal of Risk and Reliability*, Vol. 220, pp. 137-145.
- Raiffa H., Richardson J. and Metcalfe D. (2002). *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Harvard University Press.
- Ramsey, F. L. (1972). A Bayesian approach to bioassay. *Biometrics*, Vol. 28, pp. 841-858.
- Ray, B. K., Liu, Z. and Ravishanker, N. (2006). Dynamic reliability models for software using time-dependent covariates. *Technometrics*, Vol. 48, pp. 1-10.
- Rios Insua, D., Ruggeri, F., Wiper, M, (2012). *Bayesian Analysis of Stochastic Process Models*, Wiley.
- Rios Insua, D., Ruggeri, F., Soyer, R. and Rasines, D. G. (2018). Adversarial issues in reliability. *European Journal of Operational Research*, forthcoming.
- Ross, S. (1992). *Applied Probability Models with Optimization Applications*. Dover.
- Ross, S.M. (1996). *Stochastic Processes*, 2nd edition. New York: John Wiley & Sons.

- Rufo, M. J., Martin, J. and Perez, C. J. (2014). Adversarial life testing: A Bayesian negotiation model, *Reliability Engineering and System Safety*, Vol. 131, pp. 118-125.
- Sathe, P. T. and Hancock, W. M. (1973). A Bayesian approach to the scheduling of preventive maintenance. *AIIE Transactions*, Vol. 5, pp. 172-179.
- Schafer, R.E. and Singpurwalla, N.D. (1970). A sequential Bayes procedure for reliability demonstration. *Naval Research Logistics Quarterly*, Vol. 17, pp. 55-67.
- Shafiee, M. and Chukova, S. (2013). Maintenance models in warranty: A literature review. *European Journal of Operational Research*, Vol. 229, pp. 561-572.
- Sheu, S. H. and Griffith, W. S. (2002). Extended block replacement policy with shock models and used items. *European Journal of Operational Research*, Vol. 140, pp. 50-60.
- Sheu, S. H., Yeh, R. H., Lin, Y. B. and Juang, M. G. (1999). A Bayesian perspective on age replacement with minimal repair. *Reliability Engineering and System Safety*, Vol. 65, pp. 55-64.
- Shi, Y. and Meeker, W. Q. (2012). Bayesian methods for accelerated destructive degradation test planning. *IEEE Transactions on Reliability*, Vol. 61, pp. 245-253.
- Singpurwalla, N. D. (1991). Determining an optimal time interval for testing and debugging software. *IEEE Transactions on Software Engineering*, Vol. 17, pp. 313-319.
- Singpurwalla, N.D. and Wilson, S.P. (1994). The warranty problem: its statistical and game theoretic aspects. *SIAM Review*, Vol. 35, pp. 17-42.
- Singpurwalla, N.D. and Wilson, S.P. (1998). Failure models indexed by two scales. *Advances in Applied Probability*, Vol. 30, pp. 1058-1072.
- Singpurwalla, N.D. and Wilson, S.P. (1999). *Statistical Methods in Software Reliability: Reliability and Risk*. New York: Springer.
- Singpurwalla, N. D. (2006). *Reliability and Risk: A Bayesian Perspective*. Wiley, New York.
- Smeitink, E., and Dekker, R. (1990). A simple approximation to the renewal function. *IEEE Transactions on Reliability*, Vol. 39, pp. 71-75.
- Soyer, R. and Vopatek, A. L. (1995). Adaptive Bayesian designs for accelerated life testing, in *Adaptive Designs*, (N. Flournoy and W. F. Rosenberger, Eds.), Institute of Mathematical Statistics, Lecture Notes, Vol. 25, pp. 263-275.
- Soyer, R. (2007). Accelerated life tests: Bayesian design, in *The Encyclopedia of Statistics in Quality and Reliability* (F. Ruggeri, R. S. Kenett and F. W. Faltin, Eds.), Vol. 1, pp. 12-20, John Wiley and Sons, UK.
- Su, C. and Wang, X.L. (2016). Modeling flexible two-dimensional warranty contracts for used products considering reliability improvement actions. *Proceedings of the Institution of Mechanical Engineers Part O – Journal of Risk and Reliability*, Vol. 230, pp. 237-247.
- Sun, D. and Berger, J.O. (1994). Bayesian sequential reliability for Weibull and related distributions. *Annals of the Institute of Statistical Mathematics*, Vol. 46, pp. 221-249.

- Thyregod, P. (1975). Bayesian single sampling plans for life-testing with truncation of the number of failures. *Scandinavian Journal of Statistics*, Vol. 2, pp. 61-70.
- Verdinelli, I., Polson, N. G., and Singpurwalla, N. D. (1993). Shannon information and Bayesian design for prediction in accelerated life testing. in *Reliability and Decision Making*, pp. 247-256. (R. E. Barlow, C. A. Clariotti and F. Spizzichino, Eds.). Chapman & Hall.
- Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, Vol. 139, pp. 469-489.
- White, D.J. (1993) A survey of applications of Markov decision processes. *Journal Operational Research Society*, Vol. 44, pp 1073-1096.
- Wilson, J. G. and Benmerzouga, A. (1995). Bayesian group replacement policies. *Operations Research*, Vol. 43, pp. 471-476.
- Wu. S. (2012). Warranty data analysis: a review. *Quality and Reliability Engineering International*, Vol. 28, pp. 795–805.
- Yates, S.W. (2008). Australian defence standard for Bayesian reliability demonstration. *Proceedings of the 54th Annual Reliability and Maintainability Symposium*, pp. 105–109.
- Zhang, M., Gaudoin, O. and Xie, M. (2015). Degradation-based maintenance decision using stochastic filtering for systems under imperfect maintenance. *European Journal of Operational Research*, Vol. 245, pp. 531–541.
- Zhang, Y., and Meeker, W. Q. (2005). Bayesian life test planning for the Weibull distribution with given shape parameter. *Metrika*, Vol. 61, pp. 237–249.
- Zhang, Y. and Meeker, W. Q. (2006). Bayesian methods for planning accelerated life tests. *Technometrics*, Vol. 48, pp. 49–60.