



Innovative Applications of O.R.

## Adversarial issues in reliability

David Ríos Insua<sup>a,\*</sup>, Fabrizio Ruggeri<sup>b</sup>, Refik Soyer<sup>c</sup>, Daniel G. Rasines<sup>d</sup><sup>a</sup>ICMAT, CSIC-UAM-UCM-UC3M, Madrid, Spain<sup>b</sup>IMATI, CNR, Milano, Italy<sup>c</sup>Department of Decision Sciences, George Washington University, Washington, DC, USA<sup>d</sup>Department of Mathematics, Imperial College, London, UK

## ARTICLE INFO

## Article history:

Received 14 April 2017

Accepted 16 October 2017

Available online 24 October 2017

## Keywords:

Reliability analysis

Acceptance sampling

Life testing

Decision analysis

Adversarial risk analysis

## ABSTRACT

Many reliability problems involve two or more agents with conflicting interests whose decisions affect the performance of the system at hand. Examples of such problems relevant in management practice abound and include acceptance sampling, life testing, software testing, optimal maintenance, reliability demonstration, warranties and insurance. Most earlier attempts in such problems have focused on game theoretic approaches based on Nash equilibria and related concepts. However, these require strong common knowledge assumptions which do not frequently hold in practice. We provide an alternative framework based on adversarial risk analysis to deal with such problems which avoids the strong common knowledge assumptions of game theory. We illustrate the framework through acceptance sampling and life testing problems.

© 2017 Elsevier B.V. All rights reserved.

## 1. Introduction

Reliability analysis (Singpurwalla, 2006) refers to the assessment of how long a system will be functioning under given operational conditions so as to make management decisions in relation with maintenance, replacement, performance and/or redesign. From a foundational perspective, many reliability problems involve a single decision maker and may be appropriately framed with decision analytic methods, see e.g. French and Ríos Insua (2000).

However, there are reliability cases that involve two or more actors with competing interests whose decisions affect the performance of the system. Examples of adversarial situations in reliability analysis can be found in areas such as acceptance sampling (Lindley & Singpurwalla, 1991); life testing (Lindley & Singpurwalla, 1993); warranty analysis (Singpurwalla & Wilson, 1993); and optimal release policies in software testing (Zeephongsekul & Chiera, 1995). More recent work in this area considers adversarial issues in optimal maintenance as in Jackson and Pascual (2008); service and warranty contracts (Esmailia, Gamchia, & Asgharizadeh, 2014); system reliability (Hausken, 2008); and reliability demonstration (Rufo, Martin, & Perez, 2014). These problems with adversarial elements can be set up as games and are typically solved using (non cooperative) game theory methods based on Nash equilibria and related concepts, see, for example, Gibbons (1992).

A drawback of the game theoretic methodology in this application area is its underlying common knowledge assumption, criticised in e.g. Raiffa, Richardson, and Metcalfe (2002) or Lippman and McCardle (2012) and analysed in detail in Hargreaves-Heap and Varoufakis (2004). Most versions of non-cooperative game theory assume that agents not only know their own payoffs, preferences, beliefs, and possible actions, but also those of their opponents. When there is uncertainty in the game, it is assumed that players have common probabilities over the so-called types, as in games of incomplete information (Harsanyi, 1967). These common knowledge assumptions allow for symmetric joint normative analysis in which players maximise their expected utilities, expecting other players to proceed similarly. Then, their decisions are anticipated and predated by Nash equilibria related concepts. However, in many of the adversarial reliability contexts outlined above, players will not typically have such knowledge of their opponent's problem elements, for security, safety and secrecy reasons, this being aggravated in highly competitive markets as participants conceal information.

Adversarial Risk Analysis (ARA) provides a way forward, as common knowledge is no longer required (Ríos Insua, Ríos, & Banks, 2009). We support one of the participants, viewing her problem as a decision analytic one, but employ the game-theoretical structure, and other information available, to estimate the probabilities of the opponent's actions. We present here how the ARA approach may be used to solve adversarial reliability issues, focusing first on an acceptance sampling problem with two actors presented by Lindley and Singpurwalla (1991): a manufac-

\* Corresponding author.

E-mail addresses: [david.rios@icmat.es](mailto:david.rios@icmat.es) (D. Ríos Insua), [fabrizio@mi.imati.cnr.it](mailto:fabrizio@mi.imati.cnr.it) (F. Ruggeri), [soyer@gwu.edu](mailto:soyer@gwu.edu) (R. Soyer), [dg416@ic.ac.uk](mailto:dg416@ic.ac.uk) (D.G. Rasines).

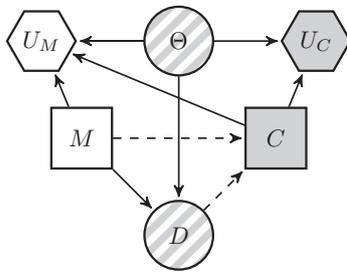


Fig. 1. Bi-agent influence diagram for the AS problem.

turer  $M$  (she) is trying to sell a batch of items to a consumer  $C$  (he) who may either accept or reject the batch provided by  $M$ .  $C$ 's decision depends on the evidence provided by  $M$  to  $C$ , based on a sample from an inspection that  $M$  may perform. The decision that  $M$  faces is whether to offer a sample to  $C$  and, if so, the size of such sample. Both  $M$  and  $C$  are assumed to be expected utility maximizers. In the original reference, the problem is represented as a game tree and the authors provide a solution assuming that  $M$ , who decides before  $C$ , knows  $C$ 's preferences and beliefs, and that they share other relevant distributions. This is a very strong common knowledge assumption and ARA will allow us to overcome such issue.

In Section 2, we provide a general ARA formulation and solution to this problem. We develop the ARA approach for a Bernoulli acceptance sampling problem in Section 3 and a life testing problem in Section 4. We finally discuss in Section 5 how other adversarial reliability applications may be dealt with through ARA.

## 2. An adversarial risk analysis approach to acceptance sampling

We start by formulating the acceptance sampling (AS) problem through a bi-agent influence diagram (BAID), as in Koller and Milch (2003) or Banks, Ríos, and Ríos Insua (2015). This is represented in Fig. 1. Square nodes represent decisions; circle nodes represent uncertainties; hexagonal nodes represent evaluations. White nodes refer to elements affecting only the manufacturer problem; grey nodes, affect only the consumer problem; striped nodes affect both agents. This BAID serves as a template for the model in Section 3. After minor modifications, it will also serve for that in Section 4.

The decision node  $M$  represents the manufacturer's decision about  $n$ , the sample size that will be offered to the consumer. The random node  $D$  denotes the data, that is, the outcome of the inspection, which depends on a parameter  $\theta$  describing the product quality, e.g. through the proportion of defective items in a batch in a quality control setting. The decision node  $C$  represents the consumer's decision to accept ( $\mathcal{A}$ ) or reject ( $\mathcal{R}$ ) the batch. Such decision will be based on the observed sample data  $D$ , which may be used by the consumer to revise his uncertainty about  $\theta$ . The utility  $u_C$  that he attains depends on his decision  $C$  and the quality parameter  $\theta$ . The manufacturer's utility  $u_M$  depends on her decision  $M$ , the consumer's decision  $C$  and the quality parameter  $\theta$ .

We assume we are supporting the manufacturer in making her decision. Note that in the ARA jargon (Banks et al., 2015) this would correspond to a sequential Defend–Attack problem, with  $M$  as the Defender and  $C$  as the Attacker.

### 2.1. The game theoretic solution

For comparison purposes, we first provide the game theoretic solution, although based on a BAID. In this case, due to the sequential nature of the problem, the consumer sees the manufacturer's decision, and thus he does not need her judgements. To move forward, the consumer should have available:

- The distribution  $p_C(\theta)$ , which represents his beliefs about the product quality  $\theta$ ;
- The distribution  $p_C(d|\theta, n)$ , which models his beliefs about the experiment result  $d$  given the quality  $\theta$  and the decision  $n$  of  $M$ ; and,
- Finally, his utility function  $u_C(c, \theta)$ .

Then, the consumer proceeds, for each  $d$  and  $n$ , by:

1. Inverting the arc  $\theta - D$  and computing, by Bayes' formula,
 
$$p_C(\theta|d, n) \propto p_C(\theta)p_C(d|\theta, n). \tag{1}$$
2. Computing the expected utilities

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

to reduce node  $\Theta$ .

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 game theoretic approach requires the manufacturer to know the consumer's beliefs  $p_C$  and preferences  $u_C$ , which is the common knowledge condition in this problem. Then, the manufacturer is capable of performing the above calculations and predict which decision will the consumer implement. She switches now to her problem. In order to solve it,  $M$  should have also available:

- The distribution  $p_M(\theta)$ , which describes her beliefs about the quality  $\theta$ ;
- The distribution  $p_M(d|\theta, n)$ , which models her beliefs about the experiment results, given the quality  $\theta$  and her decision  $n$ ;
- Her utility function  $u_M(c, n, \theta)$ .

In this case, the manufacturer proceeds by:

1. Assessing the utilities of the attained results (related with the consumer decision  $c^*(d, n)$ , which, recall, is known to her under the common knowledge assumption)

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

2. Computing the expected utilities

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

to reduce node  $D$ .

3. Computing the expected utilities

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

to reduce node  $\Theta$ .

4. Finally, computing her optimal decision through

$$n_{GT}^* = \arg \max \psi_M(n).$$

We may see that  $(n_{GT}^*, \{c^*(d, n_{GT}^*)\}_d)$  is a subgame perfect equilibrium Menache and Ozdaglar (2011).  $n_{GT}^*$  would be the sample size to be offered by the manufacturer to the consumer and  $c^*(d, n_{GT}^*)$  would be the corresponding consumer decision, if  $d$  is the observed data.

### 2.2. The ARA approach

We move now to the ARA approach which avoids the above common knowledge assumption: the manufacturer will not usually have access to the consumer's beliefs and preferences which allow her to find out his optimal decision. In turn, we shall model the manufacturer's uncertainty about the consumer's judgements and simulate from the consumer's decision making problem to form a well-founded forecast of his decision. We then use this probabilistic forecast as a key input to the decision analysis problem faced by the manufacturer.

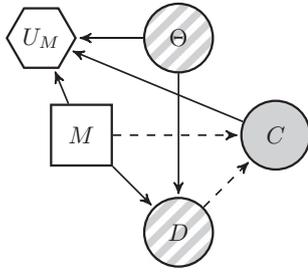


Fig. 2. Manufacturer vision of the AS problem.

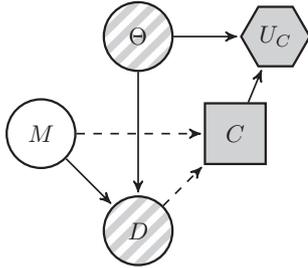


Fig. 3. Consumer vision of the AS problem.

2.2.1. Solving the manufacturer problem

In the ARA approach, we first consider the manufacturer's decision problem, described in Fig. 2. Observe that now C is uncertain to the the manufacturer.

In order to solve it, M should have available, as before,  $p_M(\theta)$ ,  $p_M(d|\theta, n)$  and  $u_M(c, n, \theta)$ , but also

- The distribution  $p_M(c|d, n)$ , which describes her beliefs about the customer decision c (accept, reject) given the experiment results d and her decision n.

The procedure to be adopted by the manufacturer essentially coincides with that in Section 2.1, except for the first step. There, under common knowledge, we knew the reaction  $c^*(d, n)$  of the consumer. Here, without common knowledge, we must forecast such reaction, through  $p_M(c|d, n)$ , and replace the first stage by:

1' Computing the expected utilities

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

to eliminate node C.

The other steps 2–4 would coincide leading to  $n_{ARA}^*$ , the size of the sample to be offered by the manufacturer to the consumer. It is important to stress that, in general,  $n_{CT}^* \neq n_{ARA}^*$ , as may be seen in examples, as in Section 3.

Note, however, that the forecast  $p_M(c|d, n)$  is nonstandard, as it requires strategic thinking about the behaviour of the consumer. To facilitate M's assessment of C's behaviour, we consider C's problem.

2.2.2. Simulating the consumer problem

The consumer's problem is reflected in Fig. 3, where now the decision M is uncertain to the consumer, although he observes it prior to making his own decision. In order to solve this ID, the consumer should have available, as stated above,  $p_C(\theta)$ ,  $p_C(d|\theta, n)$ , and  $u_C(c, \theta)$ . He does not actually 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 n and d, as above, given that he does not need to remove nodes D and M from his ID.

However, since we do not assume common knowledge, we lack the required ingredients  $u_C(c, \theta)$ ,  $p_C(\theta)$ ,  $p_C(d|\theta, n)$ . Suppose, we may model our uncertainty about them through

random utilities and probabilities, which we designate  $F = (U_C(c, \theta), P_C(\theta), P_C(d|\theta, n))$ . Now, observing that we do not need the denominator in Bayes' formula (1) for optimisation purposes, we may proceed, for each d and n, by

2' Computing the random functional

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

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

$$C^*(d, n) = \operatorname{argmax}_{c \in \{A, R\}} \Psi_C^*(d, n, c).$$

We, then, make  $p_M(c|d, n) = \Pr(C^*(d, n) = c)$ , which we feed into the manufacturer's problem in Section 2.2.1. Observe that the random functional  $\Psi_C^*(d, n, c)$  in Step 2' would be a scale transformation of  $\int U_C(c, \theta) P_C(\theta) P_C(d|\theta, n) d\theta$ , the random expected utility. The assessment of  $C^*(d, n)$  may proceed by Monte Carlo simulation, through the steps in Algorithm 1 in the Appendix.

In general, note that  $P_C(\theta)$  and  $P_C(d|\theta, n)$  could be, respectively, based on  $p_M(\theta)$  and  $p_M(d|\theta, n)$  with some uncertainty around them. In discrete cases, we could model them with Dirichlet distributions, whereas, in continuous cases, we could deal with them through Dirichlet processes. With respect to  $U_C(c, \theta)$ , we could have information about the consumer interests, model a parametric form for the utility function and, finally, derive a distribution over the corresponding parameters.

We illustrate the approach in an important case.

3. A Bernoulli model for acceptance sampling

We consider a Bernoulli acceptance sampling model based on Lindley and Singpurwalla (1991), where the items in the sample offered by the manufacturer may be defective with probability  $\theta$ . Based on the results of such sample, the consumer accepts or rejects the whole lot.

For the manufacturer, we assume that,

- The sampling model is binomial with  $p_M(d|\theta, n) \sim \text{Bin}(n, \theta)$ ;
- The quality parameter  $\theta$  follows a beta distribution,  $p_M(\theta) \sim \beta e(\beta_1, \beta_2)$ ;
- The utility function  $u_M(c, n, \theta)$  is that in Lindley and Singpurwalla (1991):
  - $u_M(A, n, \theta) = b_1 + b_2\theta + b_4n$ ,
  - $u_M(R, n, \theta) = b_3 + b_4n$ ,

where  $b_1 > b_3 > b_1 + b_2$  and  $b_2, b_4 < 0$ . In it,  $b_4$  represents the unit cost of providing each sample unit, therefore being negative. The parameter  $b_2$  represents a penalty for defectiveness; the higher  $\theta$ , the worse the corresponding cost is and, therefore,  $b_2$  will be negative. The utility also reflects that when there is rejection, the value of  $\theta$  is not relevant. The inequality  $b_1 > b_3$  portrays the preference for having the items accepted rather than rejected, and  $b_3 > b_1 + b_2$  reflects that it is worse to have a lot accepted with the worst quality, than have it rejected, for reputation reasons.

In order to facilitate the assessment of  $p_M(c|d, n)$ , which is also required:

- We do not assume uncertainty about  $p_C(d|\theta, n)$  and consider it to be the same binomial model  $\text{Bin}(n, \theta)$  as before.
- We have uncertainty about  $p_C(\theta)$ . The random distribution is designated  $P_C(\theta)$ . We model it as a beta distribution  $p_C(\theta) \sim \beta e(\alpha_1, \alpha_2)$ , with uncertainty over the parameters modelled through uniform distributions  $\alpha_1 \sim U \in [a_{11}, a_{12}]$ , and  $\alpha_2 \sim U \in [a_{21}, a_{22}]$ . This uncertainty induces the random distribution. Note that (Lindley and Singpurwalla (1991)) assumes a

beta distribution with parameters  $\alpha_1$  and  $\alpha_2$ , known by the manufacturer.

- As for  $U_C(c, \theta)$ , we adopt the same linear functional form in Lindley and Singpurwalla (1991), given by
  - $u_C(A, \theta) = a_1 + a_2\theta$ ,
  - $u_C(R, \theta) = a_3$ ,

where  $a_1 > a_3 > a_1 + a_2$  and  $a_2 < 0$ . The consumer does not carry out the costs of the  $n$  sample units, but the rest of the utility function formally coincides with that of the manufacturer. The constraints reflect that if  $\theta$  were known, acceptance would occur for  $\theta < \frac{a_3 - a_1}{a_2}$ , a break even value between 0 and 1, and rejection would hold, otherwise.

However, we do not know the values of the parameters and assume that  $a_1 \sim \mathcal{U}[0, A_1]$ ,  $a_2 \sim \mathcal{U}[A_2, 0]$  and  $a_3|a_1, a_2 \sim \mathcal{U}[a_1 + a_2, a_1]$ . This uncertainty leads to the random utility function. Note that, based on these, we may compute interesting quantities like  $Pr(\theta < \frac{a_3 - a_1}{a_2})$ , which reflects the probability that the consumer will accept, as perceived by the manufacturer.

Note that we use linear utility functions with the same functional form as in Lindley and Singpurwalla (1991) to facilitate comparisons. In many cases, it could be that utilities are non-linear, for example if we cater for risk aversion or proneness attitudes (French & Ríos Insua, 2000). This is easily acomodated by the proposed computational schemes below, albeit at a bigger computational effort.

We specify now the simulation-optimisation scheme in Section 2. First, we estimate the probabilities of various consumer decisions  $\hat{p}_M(A|d, n)$  (and  $\hat{p}_M(R|d, n) = 1 - \hat{p}_M(A|d, n)$ ) with Algorithm 2 in the Appendix. Then, we plug these into the manufacturer's problem

1. Compute the expected utilities

$$\psi_M(n, d, \theta) = \hat{p}_M(A|d, n)(b_1 + b_2\theta - b_3) + (b_3 + b_4n).$$

2. Compute the expected utilities

$$\psi_M(n, \theta) = \sum_{d=0}^n \left[ \hat{p}_M(A|d, n)(b_1 + b_2\theta - b_3) \binom{n}{d} \theta^d (1-\theta)^{n-d} \right] + (b_3 + b_4n) = g_1(n, \theta) + (b_3 + b_4n).$$

3. Compute the expected utilities

$$\psi_M(n) = \int g_1(n, \theta) \beta e(\beta_1, \beta_2) d\theta + (b_3 + b_4n).$$

4. Compute  $M$ 's optimal decision

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

Observe that in computing the expected utilities to solve  $M$ 's problem, the estimated probabilities  $\hat{p}_M(A|d, n)$  are not a function of  $\theta$ . Thus in evaluating  $\psi_M(n)$  we may change the order of the summation and integration and first evaluate the integral with respect to  $\theta$ . We can show that

$$\psi_M(n) = \sum_{d=0}^n \left[ \hat{p}_M(A|d, n) \binom{n}{d} \frac{B(\beta_1 + d, \beta_2 + n - d)}{B(\beta_1, \beta_2)} \times \left[ b_1 - b_3 + b_2 \frac{\beta_1 + d}{\beta_1 + \beta_2 + n} \right] \right] + (b_3 + b_4n),$$

which we may represent as  $\psi_M(n) = h_1(n) + (b_3 + b_4n)$ . This enables us to perform computations more efficiently.

**Example.** We consider a numerical illustration of the above approach, based on Lindley and Singpurwalla (1991).

- For the manufacturer, we use, on one hand,  $b_1 = 8$ ,  $b_2 = -14$ ,  $b_3 = 4$ ,  $b_4 = -.02$ . On the other,  $\beta_1 = 1$ ,  $\beta_2 = 3$ , reflecting that  $E(\theta) = 0.25$ .

- For the consumer, we use  $a_1 \sim \mathcal{U}[0, 12]$  and  $a_2 \sim \mathcal{U}[-15, 0]$ . We also assume that he believes that  $\theta \sim \beta e(\alpha_1, \alpha_2)$ , with  $\alpha_1 \sim \mathcal{U}[2, 5]$ ,  $\alpha_2 \sim \mathcal{U}[, 5, 2]$ . These uncertainties, respectively, induce the uncertainty in the utility function and the probability distribution. Observe that the values chosen by Lindley and Singpurwalla (1991),  $\alpha_1 = 3$ ,  $\alpha_2 = 1$ ,  $a_1 = 10$ ,  $a_2 = -13$  and  $a_3 = 6$ , are in the support of the above distributions.

With the parameters chosen in Lindley and Singpurwalla (1991), the optimal batch size is  $n_{CT}^* = 16$ . However, note that other parameters within such intervals lead to different solutions to theirs. For example, when  $\alpha_1 = 2$ ,  $\alpha_2 = 2$ ,  $a_1 = 10.5$ ,  $a_2 = -13.5$  and  $a_3 = 5.5$ , the optimal batch size is 8. Thus, by acknowledging our uncertainty about the consumer's beliefs and preferences, we avoid very strong common knowledge assumptions that may lead to erroneous decisions, with very different consequences. For a study on the robustness of the game theoretic and the ARA solutions in a security context see Ríos Insua, Ruggeri, Alfaro, and Gomez (2016).

In our case, given the above uncertainty, the forecasts we made for the consumer acceptance decision are reflected in Table 1, with Monte Carlo sample size  $K = 10^7$ . For example, when the offered sample size is  $n = 4$  and  $d = 3$  pieces are defective, the estimated acceptance probability is  $\hat{p}_M(A|d = 3, n = 4) = 0.26$  (and the rejection probability is estimated at  $\hat{p}_M(R|d = 3, n = 4) = 0.74$ ).

Table 2 displays the expected utilities of the manufacturer decisions.

Therefore, the optimal alternative would be  $n_{ARA}^* = 9$ , although we can see that the expected utility function  $\psi_M(n)$  is quite flat around such solution. Note that we reach a different decision than (Lindley and Singpurwalla, 1991).

Observe that the optimisation approach proposed effectively tries various solutions until an optimum is detected. We could perform more systematic approaches typically based on a simulation metamodel to be optimised once fitted possibly with algorithms requiring only objective function evaluations, as in the classical Nelder–Mead method. Both are described in detail in Ríos Insua, Ríos Insua, Martin, and Jimenez (2008).

#### 4. Adversarial exponential life testing

We now consider the adversarial life testing model in Lindley and Singpurwalla (1993). The items offered by the manufacturer to the consumer will have exponential life times with mean  $1/\lambda$ . This problem is called *reliability demonstration* and the decision is to accept or reject a batch of items based on the observed durations of a few offered items. It is also termed acceptance sampling for life lengths.

Standards like MIL STD 781C assume a common prior for the manufacturer and the consumer over  $\lambda$ . In the above mentioned paper, Lindley and Singpurwalla allow for different priors and utilities, although the manufacturer must essentially know the consumer utility and prior, a strong common knowledge assumption which we weaken here. We also provide a more natural formulation of the problem in which the manufacturer needs to decide both the size of the batch offered and the maximum test duration. Finally, we consider different, more realistic, utility functions.

We assume, therefore, that certain items are characterized by their unknown life lengths  $X_i$ , for  $i = 1, 2, \dots$ . The manufacturer may offer the consumer a sample of  $n$  items for testing over a period of duration  $T$ . Based on such observations the consumer decides whether to accept or not a batch of  $N$  items. The manufacturer needs to decide the values of  $T$  and  $n$ . Both the manufacturer and the consumer assume that lifetimes can be reasonably described as exponential, that is,  $p_M(X_i|\lambda)$ ,  $p_C(X_i|\lambda) \sim \text{Exp}(\lambda)$ ,  $i = 1, 2, \dots$  and independent.

**Table 1**  
Acceptance probabilities for various manufacturer decisions and experimental results.

	$n = 0$	1	2	3	4	5	6	7	–
$\hat{p}_M(\mathcal{A} d, n)$	$d = 0$	x	0.4	0.49	0.55	0.61	0.65	0.68	0.71
	$d = 1$	x	0.22	0.34	0.42	0.49	0.54	0.58	0.62
	$d = 2$	x	x	0.19	0.29	0.37	0.44	0.49	0.53
	$d = 3$	x	x	x	0.16	0.26	0.33	0.4	0.45
	–	x	x	x	x	0.14	0.23	0.3	0.36

**Table 2**  
Expected utilities of various manufacturer decisions.

	$n = 1$	2	3	4	5	6
$\psi_M(n)$	4.25	4.325	4.374	4.408	4.43	4.444
–	7	8	9	10	11	12
$\psi_M(n)$	4.453	4.456	4.457	4.456	4.451	4.444

We describe now the required ingredients for the manufacturer’s decision making problem, where  $(n, T)$  are the decisions to be made and  $x = (x_1, \dots, x_n)$  are the observed lifetimes of the sample units by time  $T$ , which serves as truncation time:

- The lifetime parameter  $\lambda$  follows an inverse gamma prior distribution,  $p_M(\lambda) \sim IG(\alpha_1, \beta_1)$ ;
- For the manufacturer utility, we adopt the functional form:
  - $u_M(\mathcal{A}, (n, T)) = -b_1T - b_2n + b_3N$ ,
  - $u_M(\mathcal{R}, (n, T)) = -b_1T - b_2n$ ,

where  $b_1$  represents the testing cost per unit time,  $b_2$  the cost of each sampling unit and  $b_3$  the unit selling price. When the customer accepts, there is an income of  $b_3N$  which compensates the expenditure of  $-b_1T - b_2n$ ; such income is not realised when there is rejection. We are assuming risk neutrality in income for the manufacturer, but the approach extends easily when we cater for risk aversion or proneness.

In order to facilitate the assessment of  $p_M(c|x, (n, T))$ , which the manufacturer requires, we also consider that:

- We have uncertainty about  $p_C(\lambda)$ . The random distribution is designated  $P_C(\lambda)$  and modelled as an inverse gamma  $IG(\alpha_2, \beta_2)$ , with  $\alpha_2 \sim \mathcal{U}[\alpha_2^1, \alpha_2^2]$  and  $\beta_2 \sim \mathcal{U}[\beta_2^1, \beta_2^2]$ . The uncertainty in  $\alpha_2$  and  $\beta_2$  induces the randomness of  $P_C(\lambda)$ . Note that [Lindley and Singpurwalla \(1993\)](#) assumes an inverse gamma model for  $p_C(\lambda)$  with parameters  $(\alpha_2, \beta_2)$  known by the manufacturer.
- For the consumer utility, we adopt the following functional form:
  - $u_C(\mathcal{A}, y) = a_1 \sum_{i=1}^N y_i^p - a_2$ ,
  - $u_C(\mathcal{R}, y) = a_3$ ,

where the vector  $y = (y_1, \dots, y_N)$  contains the life length of the items acquired by C. The parameter  $-a_2$  represents the disutility to the consumer of an item that does not function and would aggregate installation costs, among others. We have  $a_1 > 0$  and  $p \in (0, 1]$  to model the effect of time over the utility that the consumer receives from each item. The case  $p < 1$  suggests risk aversion, whereas  $p = 1$  corresponds to risk neutrality. The parameter  $a_3$  represents opportunity losses. We assume uncertainty about the utility function parameters with  $a_1 \sim \mathcal{U}[a_1^1, a_1^2]$  and  $a_1^1 > 0$ ,  $p \sim \beta e[\alpha_3, \beta_3]$ ,  $a_2 \sim \mathcal{U}[a_2^1, a_2^2]$  and  $a_3 \sim \mathcal{U}[a_3^1, a_3^2]$ . This uncertainty induces the random utility model.

Observe that the BAID describing this problem does not fully coincide with that shown in [Fig. 1](#), since the utility that C receives

depends now on the future lifetimes that he will observe. We consider this to be more realistic in this case.

Suppose that during a test of duration  $T$ , the first  $i$  offered items fail at times  $x_1, x_2, \dots, x_i$ , smaller than  $T$ , whereas the remaining  $n - i$  items last longer than  $T$ . Conditional on  $\alpha_2, \beta_2$ , the posterior  $p_C(\lambda|d, m)$  is inverse gamma with parameters  $(\alpha_2 + i, \beta_2 + \sum_{j=1}^i x_j + (n - i)T)$ . We now apply our simulation-optimisation scheme, where  $d = (x_1, \dots, x_i, i)$ , summarise the data observed in the experiment, and  $m = (n, T)$ . First, we assess the probabilities of the consumer accept–reject decisions with [Algorithm 3](#) in the Appendix. Once  $\hat{p}_M(\mathcal{A}|d, m)/K$  (and  $\hat{p}_M(\mathcal{R}|d, m) = 1 - \hat{p}_M(\mathcal{A}|d, m)$ ) is obtained, we proceed to the optimisation phase as follows:

1. Compute the expected utilities
 
$$\psi_M(m, d, \lambda) = \hat{p}_M(\mathcal{A}|d, m)b_3N - b_1T - b_2n.$$
2. Compute the expected utilities
 
$$\psi_M(m, \lambda) = \sum_{i=0}^n p(i|\lambda) \int_0^T \dots \int_0^T \psi_M(m, (x_1, \dots, x_i, i), \lambda) \times f(x_1, \dots, x_i|i, \lambda) dx_1 \dots dx_i.$$
3. Compute the expected utilities
 
$$\psi_M(m) = \int \psi_M(m, \lambda) p_M(\lambda|\alpha_1, \alpha_2) d\lambda.$$
4. Compute her optimal decision through
 
$$(n^*, T^*) = \arg \max \psi_M(n, T).$$

Then,  $n^*$  would be the optimal sample size to be offered to the customer and  $T^*$ , the optimal testing time for the sample. In this case, we would typically use a direct optimisation method like Nelder–Mead’s coupled with a regression metamodel to find the manufacturer’s decision in step 4.

**5. Discussion**

We have outlined an ARA framework to deal with adversarial issues in reliability based on BAIDs. We illustrated them with a binomial acceptance model and modified it to cope with an exponential life testing model. In the first case, the model was based on [Lindley and Singpurwalla \(1991\)](#) to allow for comparisons, showing that the game theoretic and the ARA approach actually lead to different solutions, with ARA demanding less stringent assumptions. In the second one, we departed from the formulation in [Lindley and Singpurwalla \(1993\)](#) to address more realistic cases. We have discussed and illustrated several issues concerning adversarial beliefs. For illustrations on utility modelling, including adversarial aspects, see [Wang and Bier \(2013\)](#) and [Ortega, Radovic, and Ríos Insua \(2017\)](#). In many application areas, utilities would be based not only on reliability related parameters, but other criteria like price, comfort or maintenance costs.

Our approach extends to other acceptance and life testing problems and, indeed, to other adversarial reliability areas. For example, in warranties ([Singpurwalla & Wilson, 1993](#)), we may implement the proposed procedures without the need to assume the

manufacturer's knowledge of the consumer's beliefs and preferences. Another important application area is software reliability testing, see e.g. Zeephongsekul and Chiera (1995) and Ozekici and Soyer (2003).

There are other issues that could be addressed. First, we have limited the discussion to two agents: the manufacturer and the consumer. Numerous adversarial reliability problems involve multiple agents. Relevant variants include: a manufacturer facing several consumers; several manufacturers and one consumer; several manufacturers and several consumers; and, several manufacturers, a regulator and several consumers. When more than two agents are considered, there could be various degrees of cooperation among the agents. Ideas from cooperative game theory could then be incorporated. Such extensions have applications in supply chain problems (e.g. Cachon and Zipkin (1999)).

Finally, it would be interesting to revise various quality control standards in the light of the type of models here described. The introduction of insurance concepts for the consumer, so as to protect from very low quality items is also relevant. It would also be of interest to apply this type of ideas to auctions in which several manufacturers offer a product to a buyer with reliability as one of the criteria in consumer utilities, when we support one of the auctioneers.

**Acknowledgements**

The work of DRI is supported by the Spanish Ministry of Economy and Innovation program MTM2014-56949-C3-1-R, and the AXA-ICMAT Chair on Adversarial Risk Analysis. DRI and FR are grateful to the support of the ESF-COST Action IS1304 on Expert Judgement. FR was supported by a Catedra de Excelencia program of the Government of Madrid.

**Appendix A. Algorithms**

We outline here the basic algorithms mentioned in the paper.

*Algorithm 1.* It approximates the probabilities of both consumer decisions (accept or reject) in the basic Acceptance Sampling problem, given the data  $(d, n)$  that he observes.

---

**Algorithm 1:** Acceptance sampling. Simulating the Attacker problem.

---

```

For Each  $(d, n)$ 
   $count(d, n) = 0$ 
  For  $k = 1, \dots, K$ 
    Generate  $(U_c^k(c, \theta), P_c^k(\theta), P_c^k(d|\theta, n)) \sim F$ 
    Compute  $\Psi_c^{*k}(n, d, c)$ 
    Compute  $C_k^*(d, n)$ 
    If  $C_k^*(d, n) = \mathcal{A}$ 
      |  $count(d, n) = count(d, n) + 1$ 
    end
  end
   $\hat{p}_M(\mathcal{A}|d, n) = count(d, n)/K ; \hat{p}_M(\mathcal{R}|d, n) = 1 - \hat{p}_M(\mathcal{A}|d, n)$ 
end

```

---

*Algorithm 2.* It specifies Algorithm 1 for the Bernoulli acceptance case.

---

**Algorithm 2:** Bernoulli acceptance sampling. Simulating the Attacker.

---

```

For Each  $(d, n)$ 
   $count(d, n) = 0$ 
  For  $k = 1, \dots, K$ 
    Generate  $a_1^k \sim \mathcal{U}[0, A_1], a_2^k \sim \mathcal{U}[A_2, 0], a_3^k \sim \mathcal{U}[a_1^k + a_2^k, a_1^k]$ 
    Generate  $\alpha_1^k \sim \mathcal{U}[a_{11}, a_{12}], \alpha_2^k \sim \mathcal{U}[a_{21}, a_{22}]$ 
    Compute  $\psi_c^{*k}(n, d, \mathcal{A}) = a_1^k + a_2^k \frac{\alpha_1^k + d}{\alpha_1^k + \alpha_2^k + n}$ 
    If  $\psi_c^{*k}(n, d, \mathcal{A}) > a_3^k$ 
      |  $count(d, n) = count(d, n) + 1$ 
    end
  end
   $\hat{p}_M(\mathcal{A}|d, n) = count(d, n)/K; \hat{p}_M(\mathcal{R}|d, n) = 1 - \hat{p}_M(\mathcal{A}|d, n)$ 
end

```

---

*Algorithm 3.* It estimates the consumer decision probabilities in the exponential life testing problem.

---

**Algorithm 3:** Exponential life testing. Simulating the Attacker.

---

```

For Each  $(d, m)$ 
   $count(d, m) = 0$ 
  For  $k = 1, \dots, K$ 
    Generate  $a_1^k \sim \mathcal{U}[a_1^1, a_2^1], p^k \sim \beta e[\alpha_3, \beta_3], a_2^k \sim \mathcal{U}[a_2^1, a_2^1],$ 
     $a_3^k \sim \mathcal{U}[a_3^1, a_3^1]$ 
    Generate  $\alpha_2^k \sim \mathcal{U}[\alpha_2^1, \alpha_2^1], \beta_2^k \sim \mathcal{U}[\beta_2^1, \beta_2^1]$ 
    Compute  $\psi_c^k(d, m, \mathcal{A}) =$ 
     $a_1^k N \int \int y^{p^k} p_C(y|\lambda) p_C(\lambda|d, m, \alpha_2^k, \beta_2^k) dy d\lambda - a_2^k$ 
    If  $\psi_c^k(d, m, \mathcal{A}) > a_3^k$ 
      |  $count(d, m) = count(d, m) + 1$ 
    end
  end
   $\hat{p}_M(\mathcal{A}|d, m) = count(d, m)/K; \hat{p}_M(\mathcal{R}|d, m) = 1 - \hat{p}_M(\mathcal{A}|d, m)$ 
end

```

---

Note that

$$\int_0^\infty \int_0^\infty y^{p^k} p_C(y|\lambda) p_C(\lambda|d, m, \alpha_2^k, \beta_2^k) dy d\lambda = \frac{\Gamma(p^k + 1) \Gamma(\alpha_2^k + p^k)}{\Gamma(\alpha_2^k) (\beta_2^k)^{p^k}},$$

which may be used to make the simulation scheme more efficient.

**References**

Banks, D., Ríos, J., & Ríos Insua, D. (2015). *Adversarial risk analysis*. CRC Press.  
 Cachon, G., & Zipkin, P. (1999). Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science*, 45, 936–953.  
 Esmaeilia, M., Gamchia, N. S., & Asgharizadeh, E. (2014). Three-level warranty service contract among manufacturer, agent and customer: A game-theoretical approach. *European Journal of Operational Research*, 239, 177–186.  
 French, S., & Ríos Insua, D. (2000). *Statistical decision theory*. Arnold.  
 Gibbons, R. A. (1992). *A primer in game theory*. Harlow: Pearson Education Ltd.  
 Hargreaves-Heap, S., & Varoufakis, Y. (2004). *Game theory: A critical introduction*. London: Routledge.  
 Harsanyi, J. C. (1967). Games with incomplete information played by Bayesian players, I-III. Part I. *Management Science*, 14, 159–182.  
 Hausken, K. (2008). Strategic defense and attack for series and parallel reliability systems. *European Journal of Operational Research*, 186, 856–881.  
 Jackson, C., & Pascual, R. (2008). Optimal maintenance service contract negotiation with aging equipment. *European Journal of Operational Research*, 189, 387–398.  
 Koller, D., & Milch, B. (2003). Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45, 181–221.

- Lindley, D. V., & 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*, 86, 933–937.
- Lindley, D. V., & Singpurwalla, N. D. (1993). Adversarial life testing. *Journal of the Royal Statistical Society, Series. B*, 55, 837–847.
- Lippman, S. A., & McCardle, K. F. (2012). Embedded Nash bargaining: Risk aversion and impatience. *Decision Analysis*, 9, 31–40.
- Menache, I., & Ozdaglar, A. (2011). *Network games: Theory, models and dynamics*. Morgan Claypool.
- Ortega, J., Radovic, V., & Ríos Insua, D. (2017). Utility elicitation. In L. C. Dias, A. Morton, & J. Quigley (Eds.), *Handbook of judgement elicitation*. Springer.
- Ozekici, S., & Soyer, R. (2003). Reliability of software with an operational profile. *European Journal Operational Research*, 149, 459–474.
- Raiffa, H., Richardson, J., & Metcalfe, D. (2002). *Negotiation analysis: The science and art of collaborative decision making*. Cambridge, MA: Harvard University Press.
- Ríos Insua, D., Ríos, J., & Banks, D. (2009). Adversarial risk analysis. *Journal of the American Statistical Association*, 104, 841–854.
- Ríos Insua, D., Ríos Insua, S., Martín, J., & Jimenez, A. (2008). *Simulación*. Madrid: RAMA.
- Ríos Insua, D., Ruggeri, F., Alfaro, C., & Gomez, J. (2016). Robustness for adversarial risk analysis. In M. Doumpos, C. Zopounidis, & E. Grigoroudis (Eds.), *Robustness analysis in decision aiding, optimization, and analytics*. In *International series in operations research and management science*: 241. Springer.
- Rufo, M. J., Martín, J., & Perez, C. J. (2014). Adversarial life testing: A Bayesian negotiation model. *Reliability Engineering and System Safety*, 131, 118–125.
- Singpurwalla, N. D. (2006). *Reliability and risk: A Bayesian perspective*. New York: Wiley.
- Singpurwalla, N. D., & Wilson, S. (1993). The warranty problem: Its statistical and game theoretic aspects. *SIAM Review*, 35, 17–42.
- Wang, C., & Bier, V. (2013). Expert elicitation of adversary preferences using ordinal judgments. *Operations Research*, 61(2), 372–385.
- Zeepongsekul, P., & Chiera, C. (1995). Optimal software release policy based on a two-person game of timing. *Journal of Applied Probability*, 32, 470–481.