

Equitable and adversarial pricing for items under warranty

NOZER D. SINGPURWALLA[†] AND BYEON JOON SEUNG

George Washington University, Washington, DC 20052, USA

[Received on 20 October 2004; revised on 3 March 2005; accepted on 9 March 2005]

Most manufactured goods entail the involvement of three parties, a manufacturer, an intermediary and a customer. Usually, the manufacturer warrants his/her product against defects to the intermediary, who in turn extends this warranty to the customer. Here, we consider the manufacturer's warranty to the intermediary providing the free replenishment of failed items within a specific period. Specifically, we focus on the matter of how manufacturer and intermediary arrive upon a unit price per item that the manufacturer can levy against the intermediary. First, we discuss two equitable scenarios; one in which both manufacturer and intermediary strive for fairness with a common probability distribution about the number of unit failures within a prescribed warranty period, and the other wherein manufacturer and intermediary have different distributions, but they arrive upon a consensus by a process of co-operation. Then, we discuss the adversarial scenario wherein each strives to maximize his/her expected utility without knowledge of the other's utilities and without using a common distribution. This latter scenario is perhaps the more realistic one, in the sense of what is done in practice. However, a discussion of the equitable scenario may be useful in matters of litigation pertaining to warranty claims and violations.

Keywords : utility; opinion pool; consensus; Bayesian methods; decision making.

1. Preamble

Our interest in the topic of warranties has been spawned by three scenarios wherein one of us (NDS) was invited to serve as a consultant or as an expert witness. As a professional discipline, *warrantology* covers a spectrum of issues in risk analyses, like economics, game theory, law, operations research, probability, reliability and statistics, to name a few; see Singpurwalla (2003). Thus, e.g. the first scenario we were involved with entailed the design of a 2D warranty, involving age and use, in order to minimize the risk of a financial loss due to an excessive number of in-warranty failures. This work, done on behalf of an automobile manufacturer, was motivated by engineering and economic considerations. It entailed aspects of probabilistic modelling, game theory and operations research. The technical details are given in Singpurwalla & Wilson (1992, 1993).

Work on the second scenario, also done on behalf of an automobile manufacturer, was mandated by US consumer protection laws. It entailed contingency financing to cover the cost of in-warranty failures that would be incurred in the future by the manufacturer. The law was designed to ensure that consumers do not get penalized due to a manufacturer's inability to honour warranty claim payments. The problem which boils down to how much money should be put aside now to meet future obligations entails the statistical analyses of some complex multivariate time series data on historical claims, and also an analyses of reliability data. This work has been reported in Chen *et al.* (1995, 1996).

[†] Email: nozer@gwu.edu

The third scenario was brought about by litigation against a computer chip manufacturer by a consumer. The consumer was an intermediary, in the sense that the consumer installed the chips in computers which were then sold to the US Government. The basis of the litigation was the consumer's loss of goodwill and reputation due to an excessive number of chip failures during the warranty period. The manufacturer, who was willing and prepared to honour all warranty claims, alleged consumer misuse and maintained that the actual failure rate was well within industry norms, as a line of a defence. This case, which was settled out of court, entailed aspects of reliability, utility, game theory and optimum pricing. The details are given in Singpurwalla (2000).

The adversarial litigation scenario mentioned above, brought to surface the matter of equity and fairness in the presence of uncertainty that needed to be addressed. This is the main purpose of this paper which has been motivated by the above need. Its aim is to describe a process by which a manufacturer and a consumer can agree on a price for items under a specified warranty which both see as being fair. We label such pricing as being 'equitable'. In actuality, of course, pricing is 'adversarial' in the sense that a manufacturer will choose to charge as much as is possible and a consumer will endeavour to pay as little as possible. But when the manufacturer or the consumer chooses to litigate, legal counsel may want to know what would constitute a fair price. Our hope is that the methods of this paper will help address this and related questions.

The content of this paper, though technical, is not difficult. It uses existing probabilistic, statistical and operations research tools and technologies. The statistical technologies have a Bayesian flavour. Legal scholars not comfortable with symbolic notation and a mathematical style of presentation may find the bulk of the material given here difficult to follow. For this, we apologize. However, our hope is that verbiage this preamble and the content of our introductory/overview (Section 2) will enable them to gain an appreciation about the essence of what we are endeavouring to put forth.

2. Introduction

Generally, most manufactured goods entail the involvement of three parties, a manufacturer, say M , an intermediary I who is generally an installer and a consumer, say C . As an example, consider the case of computer chips. They are manufactured by a chipmaker who sells them to a computer manufacturer who in turn installs the chips in a computer which is then sold to a consumer. The intermediary here is the computer manufacturer. A more familiar example is the case of automobiles which entails the carmaker as the manufacturer, the dealer as the intermediary and the public as the consumers.

Often, the manufacturing and installation activities are under the protection of a warranty. The manufacturer warranties his/her product against defects to the intermediary, who in turn extends this warranty to the consumer. The intermediary's warranty may or may not be supplemented to cover deficiencies of installation. The purpose of this paper is to articulate on aspects of the manufacturer's warranty. Specifically, we consider the matter of how M and I arrive upon the unit price per item, say x , that M can levy against I . We focus on price rather than the nature of the warranty (such as the duration of the warranty) because the latter is often decided on by market forces and other macro-economic and game theoretic considerations, e.g. most automobiles come with a 5-year, 50 000 miles warranty, so that the warranty per se is not specific to an M and an I . In other words, if M and I do not offer C a warranty that is prevalent in the market-place, then C will seek an alternate source for the product. Thus, all that remains to be negotiated between M and I , and in the sequel also C , are the prices at which the goods and goods plus services are exchanged.

Despite the fact that what we are considering here is only a small—albeit important—part of the overall problem of specifying optimal warranties, the stated task involves several aspects. These are questions about fairness, utilities, uncertainties and game theoretic considerations between M and I , and between M and I as a team (or partners) against C , the third party. Notwithstanding the above, determining the optimum x is, at a minimum, the first step in a large and complex process.

In Section 3, we introduce notation and develop a strategy for M and I to arrive at a fair price x , were both M and I to know k , the number of unit failures within a prescribed warranty period, say τ . Whereas an ability to know k precisely is an idealization, the material of Section 3 lays out a general framework for developing the material of subsequent sections.

Section 4 considers the case wherein M and I operate in the face of uncertainty so that k is unknown. However, both M and I strive for equity or fairness—i.e. M does not intend to overcharge I , and I does not intend to underpay M —and both use a common distribution for their uncertainty about k . We discuss here two scenarios, one in which both M and I agree at the outset on a common probability distribution for describing uncertainty about k , and the other wherein M and I start off with different distributions, but driven by the motive of co-operation and equity arrive at a consensus distribution by pooling their distributions.

Section 5 considers the more realistic adversarial scenario wherein neither do M and I share their utilities nor are they willing to use a common distribution for their uncertainty about k . Each strives to maximize his/her expected utility without a knowledge of the other's probabilities and utilities.

Section 6 concludes the paper by placing the current work in the context of the existing literature on warranty analysis, and by pointing out directions of future work.

3. Fair price under perfect information

We start by introducing some preliminaries and notation. Let us first focus attention on the actions by M . Suppose that it costs c monetary units for M to produce 1 unit of an item in question; c is assumed to be fixed and known. M sells this item to I at a unit price of x , where x , our decision variable, needs to be determined. The above transaction is to be undertaken under the protection of a warranty whose nature is as follows. If the item fails within τ time units since installation by I , then it is to be replaced free of charge by M . Similarly, if any re-installed item fails before τ , then M is to supply I with a new one, and so on, until time τ . A failed unit has no value to M whose inventory of items is large, say N . We emphasize that the essence of any warranty is the free replenishment of failed items by M during a specified period τ . Like the cost c , τ is also assumed to be known and fixed.

We now turn our attention to I . Suppose that it costs d monetary units for I to install the item in question, and suppose that I charges the customer C for every installed item w monetary units. One would expect that w equals, or is greater than, $x + d$. We assume that like c and τ , w and d are also known, so that the only unknown of the scenario considered here is x . We also suppose, as it is reasonable to do so, that C inherits M 's warranty to I , via I , i.e. it is the consumer C who is the real beneficiary of the warranty. The assumption that w and d are known is not unrealistic. Like τ , the former is determined by market forces and the latter by I , according to I 's modus operandus. Before proceeding further, we wish to stress the point that the cost of installation of all failed items is borne by I and I alone. The warranty contract obligates M only to provide free replacement of failed units made by M .

With c , τ , d and w fixed, M and I need only negotiate the price x that M wishes to charge I for, and which I is willing to accept. Besides utilities, the key factor which goes into determining x is K , the number of failures experienced by C in the time interval $[0, \tau]$. When M and I negotiate x , K is of course unknown, and it is the uncertainty about K and the manner in which this uncertainty is described that make the proposed problem of interest to us here. However, problems involving uncertainties are solved by conditioning on the unknown, and invoking the law of total probability; i.e. we first solve the problem assuming that K is known, say as k , and then averaging over all possible values of k . The aim of this section is to do the former; i.e. to address the problem of determining an optimum x , assuming that $K = k$, and k is known.

3.1 M 's and I 's profit and admissible prices

When $K = k$, and k is known, M 's profit will be $(x - c - ck) = x - c(k + 1)$. Similarly, I 's profit will be $(w - x - d - dk) = (w - x) - d(k + 1)$. Clearly, if x is such that any of the above become negative, then M and I will not enter into a contract involving an exchange of the items. When both profit models result in positive values for any particular value of x , then, in principle, this value of x should be acceptable to both M and I . However, since it is possible that a particular value of x may result in a higher profit for $M(I)$ than for $I(M)$, the party experiencing the lower profit will strive to negotiate the x . Specifically, M will want to increase x , whereas I will want to lower it. Thus, M and I will strive to find that value of x at which there is no room for further negotiation. We call this value of x an *equilibrium price*, and denote it by $x^*(k)$. Any attempt by M to increase $x^*(k)$ will be resisted by I and correspondingly, any attempt by I to lower $x^*(k)$ will be objected to by M . In determining $x^*(k)$, M and I will engage in a two-person zero sum game, and this boils down to first finding that value of x , say $M(k)$, at which M 's profit is zero, and $I(k)$, the value of x at which I 's profit is zero. Clearly,

$$M(k) = c(k + 1) \quad \text{and} \quad I(k) = w - d(k + 1),$$

where $M(k)$ is the lowest price that M can offer without suffering a loss and $I(k)$ the highest price that I can offer to pay without experiencing a loss. Both $M(k)$ and $I(k)$ are linear functions of k ; see Fig. 1 wherein it is assumed that $I(0) > M(0)$.

There are several noteworthy features of Fig. 1. First, $M(k)$ and $I(k)$ intersect at $k^* = (w/(c + d) - 1)$. Since $I(0) > M(0)$ —by assumption— $w > c + d$, and thus k^* is necessarily positive. However k^* should be an integer, and to make it so, we let k^* be the *largest* integer less than or equal to $(w/(c + d) - 1)$. The second point to note is that k^* is in some sense a critical value because when $k > k^*$, there is no value of x that is acceptable to both M and I ; i.e. for $k > k^*$, all values of x are *inadmissible*. When $k = k^*$, there exists (for all intents and purposes) a unique value of x that satisfies both M and I , namely $x^*(k^*) = M(k^*) = c(k^* + 1) = I(k^*)$. Values of x that are acceptable—though not desirable—to both M and I are deemed *admissible*. Thus, when $k = k^*$, there is a unique admissible value of x , namely $c(k^* + 1)$. For $k > k^*$, there does not exist an admissible x ; similarly, when $I(0) < M(0)$, irrespective of the value of k . Thus by far, the most interesting case is when $I(0) > M(0)$ and $k < k^*$. There are several admissible values of x , and these are illustrated by the bold vertical lines of Fig. 1. Our aim is to find, for each value of k , that x , $x \in [M(k), I(k)]$, yields an equilibrium price $x^*(k)$. But before we do so, it is useful, at least for the sake of completeness, to consider the case of $I(0) = M(0)$. When such is the case, the admissible region of Fig. 1 collapses to a point taking the value c or $(w - d)$ because $I(0) = M(0)$ implies that




FIG. 1. Admissible values of x for $I(0) > M(0)$ (bold vertical lines).




FIG. 2. Unique admissible value of x for all values of k when $M(0) = I(0)$.

$c = (w - d)$. We now have a unique equilibrium price of $x^*(k) = c$ only if $k = 0$; see Fig. 2. The implication of this figure is that when $w = c + d$, i.e. when I charges C the cost of the item plus the cost of the installation, I should pay M only what it costs M to produce the item, and the item should be guaranteed not to fail.

3.2 Equilibrium prices when $I(0) > M(0)$

For the purpose of discussion we suppose, as before, that $K = k$, and with reference to Fig. 1, note that for this value of k the admissible values of x are in the closed set $[M(k), I(k)]$. Clearly, any value of x which is greater than $M(k)$ is attractive to M , the larger the better. Similarly, any value of x that is smaller than $I(k)$ is attractive to I , the smaller the better. However, x cannot get smaller than $M(k)$, since values of $x < M(k)$ are inadmissible; similarly, for values of x greater than $I(k)$. Thus, the best (i.e. the largest) unit price that M can hope to receive from I is $I(k)$; similarly, the best (i.e. the smallest) unit price that I can hope to receive from M is $M(k)$, i.e. whereas $I(k)$ [$M(k)$] is most (least) attractive to M , it is least (most) attractive to I . This means that M and I have to arrive

at some consensus as to what price in the interval $[M(k), I(k)]$ would be a meaningful compromise. An obvious answer is the mid-point of the above interval; i.e. $x^*(k) = (I(k) + M(k))/2$. However, this choice for x needs to be defended, else it would be deemed ad hoc.

A defensible choice for x would entail a consideration of the utilities of M and I for money (i.e. for monetary units). For any price $x \geq M(k)$, M would experience a price advantage of $(x - M(k))$; let $U_M(x - M(k))$ denote the utility of this price advantage to M . Similarly, for $x \leq I(k)$, I would experience a price advantage of $(I(k) - x)$; let $U_I(I(k) - x)$ denote the utility of this price advantage to I . In Fig. 3, we illustrate several plausible functional forms of $U_M(\cdot)$ and $U_I(\cdot)$, linear, concave or convex.

As can be seen from the above figures, $U_M(x - M(k))$ increases in x , for $x \in [M(k), I(k)]$, starting from 0 and achieving a maximum of $U_M(I(k) - M(k))$. Similarly, $U_I(I(k) - x)$ decreases in x , starting from $U_I(I(k) - M(k))$ to 0. A price equilibrium will be achieved when these two utility functions cross. Figure 4 illustrates the crossing point (which gives us an equilibrium price) when the utility of money is linear. By a linear utility, we mean $U_M(x - M(k)) = (x - M(k))$ and $U_I(I(k) - x) = I(k) - x$.

The linear utilities shown in Fig. 4 are symmetric; i.e. $U_M(x - M(k)) = U_I(I(k) - x)$ for all $x \in [M(k), I(k)]$. Similarly, the convex utilities are shown by the dotted lines of Fig. 4. It is easy to




FIG. 3. Utilities of M and I for their respective price advantages.




FIG. 4. Symmetric utilities and equilibrium price.

verify that if the utilities of M and I are symmetric, then the equilibrium price is

$$x^*(k) = \frac{I(k) + M(k)}{2}.$$

If however, M has a concave utility, and I 's utility is convex, then the equilibrium price will be to the left of $(I(k) + M(k))/2$; see Fig. 5. A concave utility for M signals is the fact that M 's desirability for profits tends to level off once the profit surpasses a certain range of values. If M has a concave utility and I a convex one, then the equilibrium price will move to the left of $(I(k) + M(k))/2$. Verify that when $x^*(k) = (I(k) + M(k))/2$, and the utilities are linear, their common value is $(I(k) - M(k))/2$; it is less (more) when the utilities are convex (concave) and symmetric—see Fig. 4.

Since $M(k) = c(k + 1)$ and $I(k) = w - d(k + 1)$, we may verify that $x^*(k)$ is, for $c > d$, a linearly increasing function of k ; i.e.

$$x^*(k) = \frac{c + w - d}{2} + \frac{k(c - d)}{2}.$$




FIG. 5. Asymmetric utilities and equilibrium price.




FIG. 6. Equilibrium price as a function of k .

This means that in the case of symmetric utilities, the equilibrium price levied by M on I increases in k , the implication being that the brunt of failure has to be borne by both M and I ; Fig. 6 illustrates this behaviour of $x^*(k)$ versus k . Verify that when $k = k^*$, $x^*(k^*) = c(k^* + 1) = cw/(c + d)$, since $k^* = (w/(c + d) - 1)$.

This completes our discussion on fair pricing under perfect information.

4. Fair pricing in the presence of uncertainty

The assumption of Section 3, namely, that K is known to both M and I , is an idealization. However, this idealization is necessary to set the stage for conditionalization which is the essence of the material of this section. Here, we assume that K is unknown so that both M and I are uncertain about K . Suppose that $\pi_M(k)$, $k = 0, 1, \dots, N$, describes M 's uncertainty about K and $\pi_I(k)$, I 's uncertainty. It is often the case that $\pi_M(k)$ and $\pi_I(k)$ are different, though we can foresee situations wherein $\pi_M(k) = \pi_I(k) = \pi(k)$, for each value of k ; i.e. M and I share a common distribution for K . When such is the case how should M and I arrive at an equilibrium price? This case is considered below.

4.1 M and I share a common distribution for K

We have seen before that when K is assumed to be known, as k , the equilibrium price is $x^*(k)$. This equilibrium price yields a common utility to both M and I , which in the case of linear utilities yields the value $(I(k) - M(k))/2$; see Fig. 4. In what follows, we shall refer to this common utility as an *equilibrium utility*, and denote it by $U(x^*(k))$. What should the equilibrium price, say x^* , be when K is unknown and both M and I agree that their uncertainty about K is best encapsulated by $\pi(k)$?

The problem posed here is best addressed as one of decision making under uncertainty involving $(k^* + 1)$ choices, $x^*(0), x^*(1), \dots, x^*(k), \dots, x^*(k^*)$; see Fig. 6. The decision problem is portrayed in Fig. 7, wherein the decision node—the rectangular box—involves M and I as the joint decision makers and the random node R —the circle—pertains to the number of failures that will be experienced. The decision node is referred to as a joint decision node because both M and I are striving to take actions that are fair to each vis-a-vis their specified utilities (assumed declared to each other).

Suppose that M and I assume that the number of failures that will occur is i , and based on this supposition select $x^*(i)$ as the equilibrium price. Then, their equilibrium utility corresponding to this choice for price will be $U(x^*(i))$. Now suppose that the actual number of failures that occur upon the termination of the warranty is k , where $k \neq i$. Had they known in advance that k would occur, their equilibrium price would have been $x^*(k)$ and their equilibrium utility $U(x^*(k))$. Thus, in choosing $x^*(i)$ when in fact $x^*(k)$ should have been chosen there has been a compromise in fairness, and one way to capture this compromise is to consider the *squared error regret* $(U(x^*(i)) - U(x^*(k)))^2$. But




FIG. 7. Decision tree for equilibrium price when K is unknown.

$K = k$ can occur with probability $\pi(k)$, and thus the expected squared error regret when M and I choose $x^*(i)$ as

$$\sum_{k=0}^N [U(x^*(i)) - U(x^*(k))]^2 \pi(k).$$

However, there is an additional caveat to the above expression. Since M and I would not have chosen any $x^*(k)$ for $k > k^*$, the distribution of K , $\pi(k)$ will have to be truncated at k^* ; i.e. $\pi(k)$ for $k > k^*$ should be zero. Let us denote this truncated distribution by $\pi^*(k)$, $k = 0, 1, \dots, k^*$. With this modification in place, the *expected squared error regret* when M and I choose $x^*(i)$ is

$$\sum_{k=0}^{k^*} [U(x^*(i)) - U(x^*(k))]^2 \pi^*(k);$$

observe that the squared error regret is zero whenever $k = i$. Clearly, M and I will choose that $x^*(i)$, $i = 0, 1, \dots, k^*$, for which the above is a minimum, i.e. x^* is that $x^*(i)$ which yields

$$\min_{(x^*(i); i=0,1,\dots,k^*)} \sum_{k=0}^{k^*} [U(x^*(i)) - U(x^*(k))]^2 \pi^*(k).$$

This is our proposed solution for fair pricing when K is unknown and M and I share a common distribution for K , namely $\pi(k)$; $\pi^*(k)$ is $\pi(k)$ truncated at k^* . As an example, if $\pi(k)$ is assumed to be a Poisson distribution with mean μ , then $\pi^*(k)$ will be of the form

$$\pi^*(k) = \left(\sum_{j=0}^{k^*} \frac{e^{-\mu} \mu^j}{j!} \right)^{-1} e^{-\mu} \mu^k / k!, \quad k = 0, 1, \dots, k^*.$$

The Poisson distribution has infinite support and M 's inventory of items N is finite. Thus, unless N is very large the Poisson model for $\pi(k)$ may not be a suitable one. A more realistic model is the binomial with parameter p , $p \in [0, 1]$; i.e. K has mean (expectation) Np . With this choice

$$\pi(k) = \binom{N}{k} p^k (1-p)^{N-k}, \quad k = 0, 1, \dots, N$$

and

$$\pi^*(k) = \left(\sum_{j=0}^{k^*} \binom{N}{j} p^j (1-p)^{N-j} \right)^{-1} \binom{N}{k} p^k (1-p)^{N-k}, \quad k = 0, 1, \dots, k^*.$$

Other possible choices for $\pi(k)$ are the geometric and the *discrete Laplace* with support $0, 1, \dots, N$. The geometric has infinite support and suffers from the same criticism as the Poisson; furthermore, its probability mass at $k = 0$ is zero. This means that both M and I are certain that there will be at least one failure, and so $x^*(0)$ as a choice for the equilibrium price is eliminated. Here,

$$\pi(k) = (1-p)^{k-1} p, \quad k = 1, 2, \dots,$$

for some $p \in [0, 1]$ and

$$\pi^*(k) = \left(\sum_{j=1}^{k^*} (1-p)^{j-1} p \right)^{-1} (1-p)^{k-1} p, \quad k = 1, 2, \dots, k^*.$$

The discrete Laplace whose probability mass function is illustrated in Fig. 8 below, is the consequence of the sum of two independent discrete uniform distributions, each having support $0, 1, \dots, N/2$, where N is even.

It can be verified that in the case of the discrete Laplace

$$\begin{aligned} \pi(k) &= (k+1) \left(\frac{2}{N} \right)^2, \quad k = 0, 1, \dots, N/2, \\ &= (N+1-k) \left(\frac{2}{N} \right)^2, \quad k \left(\frac{N}{2} + 1 \right), \dots, N. \end{aligned}$$

Thus,

$$\pi^*(k) = \left(\sum_{j=0}^{k^*} \pi(j) \right)^{-1} \pi(k), \quad k = 0, 1, \dots, k^*.$$

Before closing this section, it may be helpful to note that in the case of the binomial model for k , it is $\pi(k|N, p)$ that is given as

$$\pi(k|N, p) = \binom{N}{k} p^k (1-p)^{N-k}, \quad k = 0, 1, \dots, N,$$

where p is the *propensity* of a unit's failure in $[0, \tau]$, so that if $\pi_M(p)$ denotes M 's prior distribution for p , then

$$\pi_M(k|N) \stackrel{\text{def}}{=} \pi_M(k) = \int_0^1 \binom{N}{k} p^k (1-p)^{N-k} \pi_M(p) dp.$$




FIG. 8. Probability mass function of the discrete Laplace distribution.

Similarly,

$$\pi_I(k|N) \stackrel{\text{def}}{=} \pi_I(k) = \int_0^1 \binom{N}{k} p^k (1-p)^{N-k} \pi_I(p) dp,$$

where $\pi_I(p)$ is I 's prior distribution for p .

4.2 The reconciliation (or fusing) of distributions

In this section, we consider the scenario wherein M and I do not have, at the outset, a common distribution for K so that $\pi_M(k) \neq \pi_I(k)$, $k = 0, 1, \dots$; the $\pi_M(k)$ and $\pi_I(k)$ have been defined at the end of Section 4.1. We then propose four strategies via which M and I can arrive at a common distribution $\pi(k)$. Once this is done, M and I can invoke the strategy of Section 4.1 to arrive at an equilibrium price with minimum squared error regret.

4.2.1 Pooling distributions by weighted linear combinations. A naive approach for M and I to arrive at a consensus distribution $\pi(k)$ is to take a weighted linear combination of $\pi_M(k)$ and $\pi_I(k)$ with weights w_1 and w_2 , where w_1 and w_2 are mutually agreed upon by both M and I , i.e.

$$\pi(k) = w_1 \pi_M(k) + w_2 \pi_I(k), \quad k = 0, 1, 2, \dots,$$

with $w_1 + w_2 = 1$ and $w_i \geq 0$, $i = 1, 2$. It is reasonable to suppose that $w_1 = w_2 = 1/2$ unless of course M (or I) feels that $\pi_M(k)$ ($\pi_I(k)$) is a more realistic description of the uncertainty about K in which case $w_1(w_2) > w_2(w_1)$. The formula for $\pi(k)$ given above is called the *linear opinion pool*. An advantage of the linear opinion pool is its simplicity; its disadvantage is that it lacks a theoretical foundation. Furthermore, a main difficulty in using an opinion pool is the choice of weights. The method given below circumvents this difficulty, though it too suffers from the disadvantage of lacking a normative foundation. The linear opinion pool dates back to Laplace (cf. Genest & Zidek, 1986).

4.2.2 Pooling distributions by reaching a consensus. The basic theme underlying the approach here is that once M is appraised of $\pi_I(k)$, M may be inclined to revise $\pi_M(k)$ and vice versa. This process of revision continues, on an iterative basis, for several cycles until the revision stops. When this happens a consensus is said to be achieved. The procedure is formally described below, and the conditions for convergence to a consensus are summarized below.

Let $F_M(k)$ denote the distribution function corresponding to $\pi_M(k)$; i.e. $F_M(k) = \sum_{j \leq k} \pi_M(j)$. Similarly, let $F_I(k)$ corresponding to $\pi_I(k)$. Let P_{MI} be a weight assigned by M to $F_I(k)$; similarly, P_{IM} is the weight that I assigned to $F_M(k)$. Then, $(1 - P_{MI}) \stackrel{\text{def}}{=} P_{MM}$ is the weight that M assigns to $F_M(k)$ and $(1 - P_{IM}) \stackrel{\text{def}}{=} P_{II}$ is the weight that I assigns to $F_I(k)$. These weights are chosen by M and I prior to being informed of $F_M(k)$ and $F_I(k)$ and are based on M 's (I 's) opinion of the expertise of I (M) in assessing K . Once the weights are specified, $F_M(k)$ changes to $F_{M1}(k)$ and $F_I(k)$ to $F_{I1}(k)$ in the following manner:

$$F_{M1}(k) = P_{MM} F_M(k) + P_{MI} F_I(k)$$

and

$$F_{I1}(k) = P_{IM} F_M(k) + P_{II} F_I(k).$$

The above equations can be written in matrix form as

$$\begin{bmatrix} F_{M1}(k) \\ F_{I1}(k) \end{bmatrix} = \begin{bmatrix} P_{MM} & P_{MI} \\ P_{IM} & P_{II} \end{bmatrix} \begin{bmatrix} F_M(k) \\ F_I(k) \end{bmatrix},$$

or as

$$F_1 = P F,$$

where P is the matrix of weights, $F^\top = (F_M(k), F_I(k))$ and $F_1^\top = (F_{M1}(k), F_{I1}(k))$; the superscript \top denotes a transpose.

Since $F_M(k)$ and $F_I(k)$ have been revised to $F_{M1}(k)$ and $F_{I1}(k)$, respectively, M and I will want to go through another iteration of revisions, and in the interest of self-consistency continue to use P as their matrix of weights. This in turn leads us to the relationship

$$F_2 = P F_1,$$

where $F_2^\top = (F_{M2}(k), F_{I2}(k))$. The iterative revision process continues in this manner so that, in general,

$$F_n = P F_{n-1} = P^{(n)} F, \quad \text{for } n = 2, 3, \dots$$

The process stop at that n for which $F_n = F_{n+1}$; i.e. $F_{Mn}(k) = F_{M(n+1)}(k)$ and $F_{In}(k) = F_{I(n+1)}(k)$, for $k = 0, 1, \dots$; i.e. further revisions will not change the distributions of M and I . However, the distributions of M and I will converge to each other if and only if there exists a distribution F^* such that

$$\lim_{n \rightarrow \infty} F_{Mn}(k) = F^*(k)$$

and

$$\lim_{n \rightarrow \infty} F_{In}(k) = F^*(k), \quad \text{for } k = 0, 1, 2, \dots;$$

i.e. both components of F_n converge to the same limit as $n \rightarrow \infty$. When such is the case *consensus is said to be reached*, and the consensus distribution is $F^*(k)$, where

$$F^*(k) = \psi_1 F_M(k) + \psi_2 F_I(k), \quad \text{with } \psi_1 + \psi_2 = 1.$$

The weights ψ_1 and ψ_2 are given as

$$\lim_{n \rightarrow \infty} P_{MM}^{(n)} = \lim_{n \rightarrow \infty} P_{IM}^{(n)} = \psi_1$$

and

$$\lim_{n \rightarrow \infty} P_{II}^{(n)} = \lim_{n \rightarrow \infty} P_{MI}^{(n)} = \psi_2,$$

where

$$P^{(n)} = \begin{bmatrix} P_{MM}^{(n)} & P_{MI}^{(n)} \\ P_{IM}^{(n)} & P_{II}^{(n)} \end{bmatrix}$$

is the matrix P multiplied by itself n times. A simple condition for achieving consensus is that there exists a positive integer n such that every element in at least one column of $P^{(n)}$ be positive (cf. DeGroot, 1974, Theorem 1). The rationale behind this theorem is a result in the theory of Markov chains about the existence of stationary distributions. The matrix P is a 2×2 stochastic matrix and thus it can be regarded as a one-step transition matrix of a Markov chain with two states and stationary transition probabilities.

If a consensus is indeed achieved, then a simple way to obtain ψ_1 and ψ_2 is by solving the system of linear equations given by

$$\begin{bmatrix} \psi_1 \\ \psi_2 \end{bmatrix} = \begin{bmatrix} P_{MM} & P_{MI} \\ P_{IM} & P_{II} \end{bmatrix} \begin{bmatrix} \psi_1 \\ \psi_2 \end{bmatrix},$$

under the added condition that $\psi_1 + \psi_2 = 1$.

An equivalent condition which ensures consensus is that the states of the said Markov chain communicate with each other and are aperiodic. For this to happen, P_{MM} and P_{II} should not be 1.

It is important to note that the technique of this section assumes that once M and I select their weights, they revise them only through the process of iteration involving a multiplication of the transition matrix P . There is no provision for either M or I to deviate from $P_{MM}^{(n)}$, $P_{MI}^{(n)}$, $P_{IM}^{(n)}$ and $P_{II}^{(n)}$. Finally, the method is ad hoc, in the sense that there is no law of probability that leads one to it.

4.2.3 Pooling under the spectre of a supra Bayesian. Besides lacking a normative foundation, the linear opinion pool as well as the method of consensus pooling suffer from a practical disadvantage, namely, a failure to incorporate opinion dependence between M and I . It is very likely that M and I possess common evidence, information or experience pertaining to K , and this should make their assessments $\pi_M(k)$ and $\pi_I(k)$ bear a relationship to each other. The absence of a group leader or an external decision maker is one practical reason for circumventing the incorporation of opinion dependence.

A way to rectify this deficiency is to invent a fictitious decision maker who represents the ‘synthetic personality’ of M and I . Keeney and Raiffa (1976) call this fictitious personality a *supra Bayesian*, and suppose that the supra Bayesian is endowed with altruistic qualities. Once a supra Bayesian is introduced into the picture, the pooling is achieved via a standard application of Bayes’ Law. Specifically,

$$\pi(k; \pi_M(k), \pi_I(k)) \propto \pi(k) \mathcal{L}(k; \pi_M(k), \pi_I(k)).$$

Here, $\pi(k)$ is the supra Bayesian prior distribution for K , and $\mathcal{L}(k; \pi_M(k), \pi_I(k))$ is the *likelihood* of k for the specified $\pi_M(k)$ and $\pi_I(k)$ treated as (fixed) data.

Since the supra Bayesian is assumed to be altruistic, $\pi(k)$ should be a ‘non-informative prior’. One possibility is to assume that $\pi(k) = 1/N$, where N is finite. The delicate choice of the likelihood would fall on M and I , who now instead of reaching a consensus on a distribution for K need to reach a consensus on the likelihood and be open minded about their choices. More about how to construct the likelihood is given later, but for now it suffices to say that the introduction of the supra Bayesian has facilitated a use of the normative method. The pooling operator is Bayes’ rule and the supra Bayesian posterior distribution $\pi(k; \pi_M(k), \pi_I(k))$ is the consensus distribution.

An approach for constructing the likelihood has been outlined by Lindley (1983). Specifically, if μ_M and μ_I denote the expected values yielded by $\pi_M(k)$ and $\pi_I(k)$, respectively, and σ_M^2 and

σ_I^2 the respective variances, then the likelihood is based on a multivariate normal with mean vector $(\mu_M, \mu_I)^\top$ and covariance matrix V , where

$$V = \begin{bmatrix} \sigma_M^2 & \sigma_{MI}^2 \\ \sigma_{IM}^2 & \sigma_I^2 \end{bmatrix};$$

$\sigma_{MI}^2 = \sigma_{IM}^2$ is a covariance term arrived upon by M and I via a process of consensus. If M and I agree to disagree then σ_{MI}^2 is negative; otherwise, it is positive. Values of μ_M and μ_I close to each other signal a positive covariance. It is easier for M and I to arrive upon a consensus as to whether they agree or disagree with each other than to arrive upon a consensus about their entire distributions, or the weights that should be assigned to their respective distributions in a linear opinion pool or the iterated method of pooling described in Section 4.2.2.

To pin down a specific value of σ_{MI}^2 ($=\sigma_{IM}^2$), M and I may proceed by focussing on the coefficient of correlation

$$\rho_{MI} \stackrel{\text{def}}{=} \frac{\sigma_{MI}^2}{\sigma_M \sigma_I}.$$

Since $-1 < \rho_{MI} < +1$, M and I may find it convenient to arrive upon a value of ρ_{MI} by considering the extent of their agreement or disagreement about the nature of K . If both M and I feel that their assessments about K do not bear a relationship to each other, then they will set $\rho_{MI} = 0$; they will set $\rho_{MI} = +1$ (-1) if they are in complete agreement (disagreement) with each other. The above three values of ρ_{MI} are unlikely to be used in practice. What appears to be more meaningful are values of ρ_{MI} between 0 and $+1$, and between -1 and 0, with the former (latter) reflecting the extent of agreement (disagreement) between M and I . Agreement between M and I can be attributed to the amount of information which is commonly known to both. Thus, a consensus between M and I for ρ_{MI} between 0 and $+1$ is easier to foresee than a consensus about values of ρ_{MI} between 0 and -1 . The latter encapsulates the degree of disagreement between M and I which can generally be attributed to the fact that either M has more information about K than I or vice versa. There could, of course, be other causes that are responsible for the disagreement between M and I , but these have to be eliminated because we are supposing here that M and I are striving for fairness. Thus, the testimonies about K provided by M and I are supposedly objective and free of biases. It is important to emphasize that the specification of ρ_{MI} is to be done *after* the μ_M , μ_I , σ_M^2 and σ_I^2 have been revealed by M and I to each other. It is only then that M and I can compare their sources of information and arrive at a consensus about ρ_{MI} .

Once ρ_{MI} is specified, we are able to obtain V^{-1} , the inverse of V , where

$$V^{-1} = \begin{bmatrix} (\sigma_{MM}^2)^{-1} & (\sigma_{MI}^2)^{-1} \\ (\sigma_{IM}^2)^{-1} & (\sigma_{II}^2)^{-1} \end{bmatrix}.$$

Following Lindley (1983), it can now be seen that the supra Bayesian posterior distribution for K is *approximately* a Gaussian (normal) with mean μ^* and variance σ^2 , where

$$\mu^* = w_1 \mu_M + w_2 \mu_I$$

and

$$\sigma^2 = ((\sigma_{MM}^2)^{-1} + (\sigma_{MI}^2)^{-1} + (\sigma_{IM}^2)^{-1} + (\sigma_{II}^2)^{-1})^{-\frac{1}{2}}.$$

The weights w_1 and w_2 are given as

$$w_1 = \frac{(\sigma_{MM}^2)^{-1} + (\sigma_{MI}^2)^{-1}}{(\sigma_{MM}^2)^{-1} + (\sigma_{MI}^2)^{-1} + (\sigma_{IM}^2)^{-1} + (\sigma_{II}^2)^{-1}}$$

and

$$w_2 = \frac{(\sigma_{IM}^2)^{-1} + (\sigma_{II}^2)^{-1}}{(\sigma_{MM}^2)^{-1} + (\sigma_{MI}^2)^{-1} + (\sigma_{IM}^2)^{-1} + (\sigma_{II}^2)^{-1}};$$

note that $w_1 + w_2 = 1$.

The supra Bayesian posterior distribution is the consensus distribution that we are seeking. The said distribution is approximately a normal because K is discrete. A discrete version of the result given above is possible; however, its derivation would be cumbersome and the final result less elegant than the one given above.

4.2.4 The logarithmic opinion pool; marginalization and external Bayesianity. Where the linear opinion pool satisfies the properties of *unanimity* and *marginality*, it fails to satisfy the property of *external Bayesianity*. Furthermore, the linear opinion pool is the only pooling method that satisfies the property of marginality. Thus, if external Bayesianity was important to M and I , then an alternative to the linear opinion pool is called for. One such alternative is the *logarithmic opinion pool* wherein

$$\pi(k) = A((\pi_M(k))^{w_1}(\pi_I(k))^{w_2});$$

A is a normalizing constant which ensures that $\sum_k \pi(k) = 1$, and the weights w_i are such at $w_1 + w_2 = 1$. It is common to take $w_i = 1/2, i = 1, 2$.

The property of unanimity states that if both M and I agree on the distribution of K , i.e. if $\pi_M(k) = \pi_I(k)$, $k = 0, 1, 2, \dots$, then $\pi(k) = \pi_M(k)$. The marginality property states that if θ is a vector of uncertain quantities, and if interest were to be focussed on θ_j , which is one element of θ , then this property ensures that irrespective of how θ_j is obtained, by combining all the experts' marginal distributions for θ_j or by combining all the experts' joint distributions for θ , and then obtaining the marginal of θ_j , the answer remains the same.

More germane to us here is the principle of external Bayesianity. Suppose that M and I have used any one of the proposed methods here to obtain $\pi(k)$. Suppose now that new information has been made available to both M and I that is relevant to K . Then, M and I have two choices for incorporating this added information:

- (i) Update $\pi_M(k)$ and $\pi_I(k)$ —presumably via Bayes' Law—and then pool these updated distributions to obtain $\pi^{(1)}(k)$ using any one of the proposed methods; i.e. update and pool.
- (ii) Update $\pi(k)$ in the light of the new information to obtain $\pi^{(2)}(k)$; i.e. pool and update.

The property of external Bayesianity requires that $\pi^{(1)}(k) = \pi^{(2)}(k)$, $k = 0, 1, \dots$

Cooke (1991) has unified the linear opinion pool and the logarithmic opinion pool by proposing the following generalization :

$$\pi(k) = A[(\pi_M(k))^r w_1 + (\pi_I(k))^r w_2]^{1/r},$$

where $r \in [0, 1]$; w_1 and w_2 are weights with $w_1 + w_2 = 1$ and \mathcal{A} is a normalizing constant which ensures that $\sum_k \pi(k) = 1$. When $r = 1$, the above results in the linear opinion pool. When $r \downarrow 0$, we get the logarithmic opinion pool. The values of r between 0 and 1 result in other rules of combination.

As a final comment, the linear and the logarithmic opinion pools point out the existence of an *impossibility theorem* along the lines of Arrow's (1951) work on choice theory. Specifically, a combining rule cannot simultaneously satisfy a number of seemingly compelling desiderata.

4.3 Expected and realized utilities under fair pricing

In Section 4.1, we have seen how M and I arrive upon an equilibrium price x^* when K is unknown and both M and I describe the uncertainty about K via $\pi(k)$, $k = 0, 1, \dots, N$. Recall that this x^* is arrived upon by minimizing the expected squared error regret. This is contrary to what M and I will do when they are adversaries—the scenario considered in Section 5. When M and I are adversarial each will arrive upon a price that will maximize his/her expected utility; i.e. M and I conduct themselves as *homo economicus*. Thus, the question now arises as to the effect of x^* on the actual and the expected utilities of M and I , when both are altruistic. In what follows, we assume linear utilities for both M and I .

Let $U_M(x^*|k)$ denote M 's realized utility when the price is x^* and k failures are actually observed; similarly, $U_I(x^*|k)$ for I . Clearly,

$$U_M(x^*|k) = x^* - c(k + 1)$$

and

$$U_I(x^*|k) = (w - x^*) - d(k + 1).$$

Let $U_M(x^*)$ be M 's utility when the price is x^* ; then by the law of total probability

$$U_M(x^*) = \sum_k U_M(x^*|k)\pi(k);$$

similarly,

$$U_I(x^*) = \sum_k U_I(x^*|k)\pi(k).$$

Clearly, $U_M(x^*)$ and $U_I(x^*)$ are the *expected utilities* of M and I when both agree upon x^* as the (fair) equilibrium price.

Verify that $U_M(x^*) = x^* - c - cE(K)$ and that $U_I(x^*) = (w - x^*) - d - dE(K)$, where $E(K) \stackrel{\text{def}}{=} \sum_k k\pi(k)$, is the expected value of K , when K has distribution $\pi(k)$. The above expected utilities will be equal if

$$E(K) = \frac{(w - 2x^*) + (c - d)}{(d - c)}.$$

If $E(K)$ is greater (less) than the right-hand side of the above equation, then $U_M(x^*) > (<) U_I(x^*)$, and thus with K unknown, the altruism of M and I could get compromised.

But x^* , we recall, is a surrogate for $x^*(i)$, and since $x^*(i) = (c + w - d)/2 + i(c - d)/2$ —see Section 3—the condition for altruism given above reduces to the requirement that

$$E(K) = i, \quad i = 0, 1, 2, \dots, k^*.$$

If $E(K) > (<) i$, then M will experience a higher (lesser) expected utility than I .

The realized utilities of M and I will depend on k , the number of failures that actually occur during the period of warranty. With k failures, M 's realized utility is $x^*(i) - c(k + 1)$, where $x^*(i) = (c + w - d)/2 + i(c - d)/2$. Similarly, I 's realized utility will be $(w - x^*(i)) - d(k + 1)$. Verify that the realized utilities will be equal if $k = i$; when $k < (>) i$, M 's realized utility will be bigger (smaller) than I 's. Thus, a large number of failures will be more deleterious to M than to I . This means that M will strive to manufacture the items in such a manner that $\pi_M(k)$ assigns small probabilities to large values of K .

5. Adversarial pricing in the presence of uncertainty

The material of Section 4 is viable if both M and I are equitable (or altruistic), or have been mandated to operate in a regulated environment. The altruism of M and I could be innate, or it could be motivated by a desire to gain market dominance by driving away the competition through a process of the lowest possible pricing. Since the former is a rarity (except in highly socially conscious societies), and the latter could be illegal, the more realistic scenario is the one in which both M and I are adversaries. This means that the utilities of M and I are private to them, and each will act in a manner that maximizes his/her expected utility. With utilities being private, any negotiations about the reasonableness of pricing has to centre around the differences between $\pi_M(k)$ and $\pi_I(k)$, $k = 0, 1, \dots, N$. These, we recall, are the (personal) probability distributions of M and I , respectively, about K , the number of failures during the period of warranty.

As a starting point, we note that under an adversarial scenario; the k^* of Fig. 1 is unknown. This is because the function $M(k)$, $k = 0, 1, \dots$, is private to M , and therefore unknown to I . Similarly, the function $I(k)$ is unknown to M . It is reasonable to suppose that M will not inform I as to what c is, and similarly, I will not inform M as to what d is. It is reasonable to suppose that w is known to M , since w is the price that I levies upon C , but this may only be known retrospective of an agreed upon x , since w must be greater than x . Of course, w could also be outside the control of I , it being dictated by the market-place. To describe the process by which M and I , as adversaries, arrive upon a price x , we start with a discussion of how M arrives upon an initial offering price to I . The initial offering price could be negotiated down by I , and the process by which this is done is articulated in this section.

5.1 M 's initial offering price

As a start, suppose that K were to be known to M as k . Thus, $M(k)$ is known to M , so that for M the initial offering price $\tilde{x}(k)$ has to be $M(k)$ or more. But not knowing $I(k)$, M is not able to know how much more $\tilde{x}(k)$ can be; in the altruistic scenario, M would know that $\tilde{x}(k)$ cannot exceed $I(k)$ because if it did, I would reject M 's offer. Thus, in principle, M 's initial offering price could be as large as M pleases, as long as it is greater than $M(k)$. However, this is unreasonable because I will not agree to any large value of $\tilde{x}(k) \geq I(k)$, and may even be offended by it to the extent that I may seek an alternative source. Thus, M needs to bound his/her $\tilde{x}(k)$. This M can do in two possible

ways. The first is to guess a value of d (by, say, a prior on d) and the second is to assume a concave increasing utility for money; e.g. a logarithmic utility. The latter strategy is the preferable one, since the former could end up dictating M 's pricing policy based on M 's view of I 's costs; it could even result in $\tilde{x}(k)$ being less than $M(k)$.

To arrive upon an initial offering price that is reasonable, M 's utility for money needs to be concave as is illustrated by Fig. 9. If the utility is linear or convex, M would be prompted to make the initial offering price so large that M and I would not be able to arrive upon a workable solution vis-a-vis price.

With a concave utility for money, and were M to know that $K = 0$, M 's initial asking price $\tilde{x}(0)$ would be the smallest price at which M 's utility will be within an ϵ of 1, for some $\epsilon > 0$. Specifically,

$$\tilde{x}(0) = \inf\{x: U_M(x - M(0)) \geq 1 - \epsilon\},$$

where $U_M(x)$ is M 's utility for price x .

The question of what M will choose for ϵ is an intriguing one. Certainly, the choice of ϵ is subjective, depending on the extent to which M desires to do business with I . The most likely scenario is the one in which ϵ is random, with a (subjective) probability distribution $\pi_M(\epsilon)$. But to lay down a foundation for our analysis, it is best to suppose—for now—that ϵ is fixed by M . Then, were M to assume that his/her units are highly reliable, so that $K = 0$, M 's initial offering price to I would be $\tilde{x}(0)$. But K is unknown to M and thus M needs to know how to proceed using $\tilde{x}(0)$ as a standard or a *yardstick* for the initial offering price. Since $M(k) > M(0)$, for $k = 1, 2, \dots$, M must be prepared to accept a sacrifice in utility for every failure in excess of zero. Suppose then, that were $K = k$, M was willing to sacrifice $(k\epsilon)$ units of utility so that $\tilde{x}(k)$, M 's initial offering price, would be given as

$$\tilde{x}(k) = \inf\{x: U_M(x - M(k)) \geq 1 - k\epsilon\}.$$




FIG. 9. M 's concave utility for money and initial asking price.

Figure 10 illustrates the above feature vis-a-vis the standard with $K = 0$. In Fig. 10, we have shown the case of $\tilde{x}(k)$ being greater than $\tilde{x}(0)$. This will always be true so long as M 's utility function is invariant with k ; also, we recall that $M(k) \geq M(0)$ and so $\tilde{x}(k)$ is pre-dispositioned to be greater than (or equal) to $\tilde{x}(0)$.

Using the above scheme, M is able to obtain, besides $\tilde{x}(0)$, the quantities $\tilde{x}(1), \dots, \tilde{x}(N)$. However, K is unknown, but M has a distribution for it, namely $\pi_M(k)$. Which of the above $\tilde{x}(k)$ values should M choose for an *unconditional* (upon K) initial offering price for each unit sold to I ? This boils down to a problem of decision making under uncertainty for which we need to first define some criteria for gain or loss (i.e. a utility for making decisions). For this, we consider the following scenario.

Suppose that M were to choose $\tilde{x}(k)$; this choice would suggest that M 'expects' that the number of failures during the period of warranty will be k . This in turn implies that M is prepared to accept a utility of $(1 - k\epsilon)$. Now, if the actual number of failures turns out to be k , M will not experience a surprise in the actual utility that M will receive. Suppose now that the actual number of failures is j , $j \neq k$. Then, the actual utility experienced by M is $(1 - j\epsilon)$. The *deviation* of the actual utility from what M expects to receive is

$$(1 - j\epsilon) - (1 - k\epsilon) = \epsilon(k - j) \stackrel{\text{def}}{=} D(k, j).$$

If $k > j$, this deviation is positive, if $k < j$, it is negative. Thus, the expected deviation in utility, when M chooses $\tilde{x}(k)$ is $\sum_{j=0}^N \epsilon(k - j)\pi_M(j)$. This could be positive or negative, depending on the nature of $\pi_M(j)$. By the principle of maximization of expected utility, M will choose that $\tilde{x}(k)$ for




FIG. 10. Comparison of M 's initial asking price when $K = 0$ and when $K = k$.




FIG. 11. Decision tree for M 's unconditional initial offering price.

which $\sum_{j=0}^N \epsilon(k-j)\pi_M(j)$ is a maximum, i.e. \tilde{x} is that $\tilde{x}(k)$ which yields

$$\max_{(\tilde{x}(k); k=0,1,\dots,N)} \sum_{j=0}^N \epsilon(k-j)\pi_M(j).$$

Figure 11 illustrates the schematics of the above described decision-making process.

A comparison of Fig. 11 with Fig. 7 of Section 4 reveals many parallels. The main difference is in the consequence terms. In the altruistic scenario, a squared error loss is assumed. This is reasonable because M does not wish to penalize I and vice versa. In the adversarial case, both M and I wish to maximize their expected utilities. Thus for M , a suitable consequence to consider is $D(k, j)$, $k, j = 0, 1, \dots, N$.

This completes our discussion on M 's initial offering price. We next discuss how I will react to M 's initial offer of \tilde{x} monetary units per unit item.

5.2 I 's response to M 's initial offering price

I 's response to M 's initial (unconditional) offering price of \tilde{x} per unit would entail two actions; to accept M 's offer or to reject M 's offer. If I rejects M 's offer, then M has several choices as a comeback and these will be discussed in Section 5.3. Here, we outline the process by which I arrives upon one of the above two decisions. We start by assuming that K were known to I as k .

When I knows k , I also knows $I(k)$, the maximum unit price that I can accept. Consider $U_I(I(k) - x)$, I 's utility for money. Clearly, when $x = I(k)$, I 's utility is zero, since there is no profit to I if I 's income is balanced out by I 's outlay. When x is larger than $I(k)$, I will experience a loss and thus $U_I(I(k) - x)$ will be negative. Thus, $U_I(I(k) - x)$ is a decreasing function of x , commencing with 1 when $x = 0$ and taking negative values once x crosses the threshold of $I(k)$. This feature is best encapsulated by supposing that $U_I(I(k) - x)$ is a convex decreasing function




FIG. 12. I 's utility as a function of price x .

of x ; see Fig. 12. A specific choice for this utility function could be the exponential function which ensures that the utility is bounded both above (by 1) and below by some negative constant. Figure 12 illustrates the behaviour of $U_I(I(k) - x)$ for different values of k , $k = 0, 1, \dots$. Recall (see Fig. 1) that $I(k)$ decreases in k .

With the above in place, suppose now that I accepts M 's initial offering price of \tilde{x} , when I assumes that $K = k$. Then, if $\tilde{x} < I(k)$, I will experience a positive utility of $U_I(I(k) - \tilde{x})$; however, if $\tilde{x} > I(k)$, I will experience a negative utility of $U_I(I(k) - \tilde{x})$. I 's utility will be zero, if $\tilde{x} = I(k)$. An exponential form of the utility function will ensure that the positive and the negative utilities will not be symmetric in $|I(k) - \tilde{x}|$, i.e. for $\tilde{x} < I(k)$, $U_I(I(k) - \tilde{x})$ will be greater than $|U_I(I(k) - \tilde{x})|$ when $\tilde{x} > I(k)$.

What happens if I chooses to reject M 's offer of \tilde{x} , and I assumes that $K = k$? Clearly, if $\tilde{x} < I(k)$, I will have *forgone* an opportunity to gain $U_I(I(k) - \tilde{x})$ units of utility; however, if $\tilde{x} > I(k)$, I will have lost nothing. Thus to summarize, when I chooses to reject M 's offer of \tilde{x} per unit, I 's utility will be

$$U_I^*(k) = \begin{cases} -U_I(I(k) - \tilde{x}), & \text{if } \tilde{x} < I(k), \\ 0, & \text{if } \tilde{x} \geq I(k). \end{cases}$$

The above ingredients constitute the inputs to I 's decision tree of Fig. 13 for acceptance or rejection of M 's offer, when K is unknown.

If $\pi_I(k)$ denotes I 's (subjective) prior distribution for K , then I 's expected utility in accepting M 's offer of \tilde{x} will be

$$\sum_{k=0}^N U_I(I(k) - \tilde{x}) \pi_I(k).$$

Similarly, I 's expected utility in rejecting M 's offer will be

$$\sum_{k=0}^N U_I^*(k) \pi_I(k).$$




FIG. 13. I 's decision tree for accepting or rejecting M 's offer.

Thus, I will decide to accept (reject) M 's initial offer of \tilde{x} per unit of item sold to I by M , if

$$\sum_{k=0}^N U_I(I(k) - \tilde{x})\pi_I(k) > (<) \sum_{k=0}^N U_I^*(k)\pi_I(k).$$

In a truly adversarial scenario, with I having arrived upon a decision to accept M 's offer of \tilde{x} , I may want to have M lower \tilde{x} before committing to accept the initial offer. Should this happen, M may choose to do so by re-doing the analysis of Section 5.1 by increasing the ϵ , or by rationalizing a change in $\pi_M(k)$, or both, so that the resulting \tilde{x} is lower than its previous value. However, the more interesting scenario is the one wherein I rejects M 's initial offer. This is discussed next.

5.3 M 's response in the face of rejection by I

How should M react to a rejection of \tilde{x} by I ? M should start by investigating the causes of rejection. There are three: the first is \tilde{x} , the offering price; the second is the nature of $U_I(\cdot)$ and $U_I^*(\cdot)$, I 's utilities, and the third is $\pi_I(k)$, I 's prior distribution for K . Since M 's offering price \tilde{x} is a function of $U_M(\cdot)$, M 's utility, ϵ , the fraction by which M is prepared to sacrifice utility for every failure during the warranty, and $\pi_M(k)$, M 's prior for K , M may first consider increasing his/her ϵ as a way to lower \tilde{x} . But suppose that M is neither prepared to change his/her $U_M(\cdot)$ nor prepared to change his/her ϵ . Then, the only option that is open to M is tweaking the $\pi_M(k)$. Specifically, M may consider shifting $\pi_M(k)$ so that the smaller values of k are emphasized, and/or decreasing the entropy of $\pi_M(k)$, say ε_M , where

$$\varepsilon_M = - \sum_{k=0}^N \pi_M(k) \log \pi_M(k).$$

Recall (see Fig. 1) that since the equilibrium price in the altruistic case increases with k , $k = 0, 1, 2, \dots$, M 's initial offering price will also tend to increase with k . Thus, our proposal is that $\pi_M(k)$ be tweaked by M to shift the probability mass in $\pi_M(k)$ towards the small values of k .

But suppose that M does the above (or feels that the $\pi_M(k)$ used by M reflects M 's honest opinions about the uncertainty of K), and M 's new initial offering price is still unacceptable to I . Then, M 's only hope is to make I change $\pi_I(k)$ so that the smaller values of k will be emphasized. It is not reasonable to assume that M is able to convince I to change his/her $U_I(\cdot)$ and $U_I^*(\cdot)$. So how can M convince I to change $\pi_I(k)$?

Our proposed strategy is to let M offer I a free sample of say n units from M 's large inventory of N units for testing by I . If M feels confident of his/her $\pi_M(k)$, then M 's hope is that the result of the test will provide I sufficient evidence to change $\pi_I(k)$. M is now faced with the question of how large should n be, including, of course, the possibility that n may turn out to be zero, in which case the relationship between M and I comes to closure. M 's optimal choice for n is discussed below.

5.3.1 Adversarial testing to change I 's mind. Recall that it costs c monetary units for M to produce an item and that the period of warranty is $\tau > 0$. Suppose that M were to offer n units to I (from M 's inventory of N) to test under the stipulation that the cost of the units plus the cost of testing would be borne by M . For the purpose of discussion suppose that N is so large that M 's inventory is not significantly depleted by n . Since it is unreasonable to expect that I will test the n items for

τ units of time— τ could be large— M 's proposal could be to test the n items for some fraction, say f , of τ for $0 < f < 1$. Let r be the number of failures observed by I during the time period $[0, f\tau]$; $r = 0, 1, \dots, n$. M 's hope is that the results of the test experiment (n, f, r) will provide I sufficient evidence to change $\pi_I(k)$ in such a manner that I will reverse the decision to reject M 's offer of \tilde{x} per unit item. M 's problem is to design his/her experiment (n, f) in a manner that achieves M 's objective. Let the cost of testing an item for 1 unit of time be α , irrespective of when an item under test fails. Then, M 's cost of testing n items for a duration of $f\tau$ units of time will be $n(c + \alpha f\tau)$. This is a disutility of M , and a serious one, because the results of the test could even surprise M in the sense that they may re-affirm I 's decision of not accepting M 's offer.

To proceed further, we first discuss I 's decision to accept or reject M 's offer, subsequent to performing the test and observing r failures out of n in time $f\tau$. This is very simple to outline because I will use Bayes' Law to update $\pi_I(k)$ in the light of (n, f, r) as a first step and re-compute his/her expected utilities for accepting or rejecting M 's offer of \tilde{x} . Specifically, I will accept (reject) the offer if

$$\sum_{k=0}^{N-n} U_I(I(k) - \tilde{x}) \pi_I(k|(n, f, r)) > (<) \sum_{k=0}^{N-n} U_I^*(k) \pi_I(k|(n, f, r)),$$

where $\pi_I(k|(n, f, r))$ is I 's updated probability of k failures in time τ , $k = 0, 1, \dots, (N - n)$ given that n items tested for the period $[0, f\tau]$ yield r failures, $r = 0, 1, \dots, n$. Strategies for obtaining $\pi_I(k|(n, f, r))$ will be discussed later in Section 5.3.2, once the general principles are laid out.

For any fixed values of n and f , there will be some values of r that lead to acceptance and the remaining to rejection. In principle, there could be no values of n which could lead to acceptance. This could happen, e.g. if $\pi_I(k)$ is so skewed to the right, that no amount of evidence to the contrary will cause I to reverse the decision to reject. In any case, let $A(n)$ denote the set of values of r that result in I 's acceptance of M 's offer, and $R(n)$ the values of r that do not result in I changing his/her mind about rejection. Values of r that lead to an equality of expected utilities may be assigned to either $A(n)$ or $R(n)$.

We now consider M 's utility in I 's acceptance of M 's offer of \tilde{x} , when $K = k$. For every unit sold to I , M will experience a profit of $\tilde{x} - c(k + 1)$ —see Section 3.1. Thus, the profit to M when $(N - n)$ units are sold is $(N - n)(\tilde{x} - c(k + 1))$. However, M is to bear the cost of testing n units, and this reduces M 's profit by $n(c + \alpha f\tau)$. Thus, M 's net profit, were I to accept M 's offer subsequent to testing, will be $(N - n)(\tilde{x} - c(k + 1)) - n(c + \alpha f\tau)$. Since M 's utility is not linear in monetary units, we need to transform the above net profit by $U_M(\cdot)$, and so M 's utility in I 's acceptance of M 's offer is

$$U_M[(N - n)(\tilde{x} - c(k + 1)) - n(c + \alpha f\tau)].$$

Now M 's expected utility will be determined by averaging the above with respect to $\pi_M(k|(n, f, r))$. However, when M is contemplating the value of n , r is unknown and so M must also average over r . Consequently, M 's expected utility in I 's acceptance of M 's offer is

$$\sum_{k=0}^{N-n} \sum_{r \in A(n)} U_M[(N - n)(\tilde{x} - c(k + 1)) - n(c + \alpha f\tau)] \pi_M(r|n, f) \pi_M(k|n, f, r).$$

We now consider M 's utility in I 's rejection of M 's offer subsequent to testing, when $K = k$. When I rejects, M experiences a loss due to cost of testing and thus M 's disutility is $n(c + \alpha f \tau)$. Using an argument analogous to the one used when I accepts, M 's expected disutility when I rejects is

$$\sum_{k=0}^{N-n} \sum_{r \in R(n)} U_M(n(c + \alpha f \tau)) \pi_M(r|n, f) \pi_M(k|n, f, r).$$

Thus to summarize, M 's total expected utility in offering I an opportunity to test n items for a duration of $f \tau$ time units, and accepting or rejecting the batch of $(N - n)$ units, with a unit cost of \tilde{x} per unit for acceptance is

$$\begin{aligned} & \sum_{k=0}^{N-n} \sum_{r \in A(n)} U_M[(N - n)(\tilde{x} - c(k + 1)) - n(c + \alpha f \tau)] \pi_M(r|n, f) \pi_M(k|n, f, r) \\ & - \sum_{k=0}^{N-n} \sum_{r \in R(n)} U_M(n(c + \alpha f \tau)) \pi_M(r|n, f) \pi_M(k|n, f, r). \end{aligned}$$

M will choose that combination of n and f which maximizes the above as a test plan for I to consider.

As stated before, if the results of the test surprise M about the quality of his/her items and that I continues to reject subsequent to the test, M may want to re-compute an \tilde{x} in the light of $\pi_M(k|n, f, r)$ and see if the new \tilde{x} would lead to I 's acceptance of M 's offer. If I continues to reject M 's offer, then the utilities of M and I , and/or the inherent quality of the items, is such that M and I are unable to complete a transaction. They have agreed to disagree!

5.3.2 Assessing predictive distributions. The material of Section 5.3.1 entails an assessment of the predictive distributions $\pi_i(k|n, f, r)$, $i = I, M$, and $\pi_M(r|n, f)$. The purpose of this section is to discuss such assessments. For this, we first need to recall—see Section 4.1—that a realistic likelihood for K is a binomial distribution with parameters N and p . Thus, for the scenario that is of relevance to us here,

$$\pi(k|N, n, f, p) = \binom{N-n}{k} p^k (1-p)^{N-n-k}, \quad k = 0, 1, \dots, N-n,$$

where p is an item's *propensity* of failures in time $[0, \tau]$. Suppose that $\pi_M(\theta)$ and $\pi_I(\theta)$ denote the prior distributions of M and I , respectively, for θ , where θ is the item's *propensity* of failure in time $f \tau$. Then, it is relatively straightforward to verify that $p = m\theta$, where $m = F(\tau)/F(f \tau)$, and $F(\tau)$ is the distribution function of an item's time to failure, say T , evaluated at τ ; i.e. $P(T \leq \tau) = F(\tau)$ and $P(T \leq f \tau) = F(f \tau)$. Since p and θ are interpreted as propensities, the distribution function F must also be so interpreted. With the above in place, it now follows from the law of total probability that

$$\begin{aligned} \pi_M(r|n, f) &= \int_0^1 \pi_M(r|n, f, \theta) \pi_M(\theta) d\theta \\ &= \int_0^1 \binom{n}{r} \theta^r (1-\theta)^{n-r} \pi_M(\theta) d\theta. \end{aligned}$$

Furthermore, by Bayes' Law

$$\pi_M(\theta|(n, r, f)) \propto \binom{n}{r} \theta^r (1 - \theta)^{n-r} \pi_M(\theta)$$

is M 's posterior distribution for θ based on the results of the test. Similarly, I 's posterior for θ is

$$\pi_I(\theta|(n, r, f)) \propto \binom{n}{r} \theta^r (1 - \theta)^{n-r} \pi_I(\theta).$$

Since $p = m\theta$, the above two posterior distributions can be used to induce M 's and I 's posterior distributions of p , namely $\pi_M(p|(n, r, f))$ and $\pi_I(p|(n, r, f))$. However to do so, both M and I should agree on an m . If not, each can choose his/her m to obtain their induced distributions. Once this is done, the law of total probability can be invoked again to obtain, for $i = M, I$,

$$\pi_i(k|(n, r, f)) = \int_0^1 \binom{N-n}{k} p^k (1-p)^{N-n-k} \pi_i(p|(n, r, f)) dp.$$

This completes our discussion on the matter of adversarial pricing under uncertainty. The development here is possible because of the sequential nature of the underlying process. Namely, M first makes an offer and then I reacts to it; this is then followed by M 's counter-offer and so on, until M and I agree on a price or agree to disagree.

6. Summary, conclusions and perspective

The topic of warranties is intriguing. It is at the interface of the economic, mathematical and moral sciences. An overview is in Blishke & Murthy (1996). Our interest in this topic is more from a mathematical than an economic or sociological (to include the moral and the psychological) perspective. The mathematical content of the warranty problem has many facets; these include an interpretation of probability with the objective, the subjective and the logical coming into play simultaneously, probabilistic models with multiple scales, aspects of games and decision theory and optimization. Singpurwalla & Wilson (1993) give a perspective on the statistical and game theoretic aspects of warranties. Neither has the issue of optimal pricing under warranties been addressed therein nor has this matter been addressed in the other references mentioned below. The focus in Singpurwalla & Wilson (1993) has been the specification of two dimensional warranties, one based on age and one based on usage. This scenario has spawned the development of failure models with multiple scales, one indexing time and the other indexing usage. The scales bear a relationship to each other. Several plausible probability models which reflect the essence of this feature have been introduced by Singpurwalla & Wilson (1998).

Statistical aspects of the warranty problem entail many facets ranging from an analysis of failure data from warranty claims (cf. Lawless, 1998) to a time series analysis of warranty claim costs. Such data consist of inter-related time series and their analysis via raster-scan methods; see Chen *et al.* (1995) who use Markov-mesh models for filtering and forecasting censored warranty claims data.

There has been much debate in the literature about the nature of probability, between proponents of the objective, the subjective and the logical views. Some have maintained that all the three views are valid and necessary, each to suit a particular scenario. Interestingly, the topic of warranties brings

into play all three views and it does so simultaneously; see Singpurwalla (2000) for an explication of the above.

We have said before that the focus of this paper is on pricing under warranties. This is a topic on which little, if any, has been written, and especially so from a subjectivistic Bayesian perspective. Such a perspective introduces some realism into the problem. However, to develop the material it has been necessary to proceed on a step-by-step basis. Specifically, we first develop a strategy for equitable behaviour under certainty and then relax it by introducing uncertainty. We then go to the adversarial case. The strategy used here in going from Section 3 to Section 5 is akin to that used in the law of total probability and hierarchical modelling. We first condition on the unknowns and then average out with respect to the unknown. Our strategy also entails the use of recent developments in Bayesian statistics, namely, information fusion and adversarial testing. Putting these two together in the context of the warranty problem is a noteworthy feature of this paper.

The purely game theoretic scenario of pricing under warranties has not been considered by us, and this could be a topic for future work. The game theoretic scenario arises if pricing is to be under a sealed contract, i.e. there are several M 's and one I , and each M was to offer his/her offering price \tilde{x} in a sealed envelope. I would then choose one (or none) of these for further negotiations or outright acceptance. Under such circumstances, each M would have to contemplate the actions of the other M 's and this could lead to the problem of an infinite regress. The strategy used by us is a sequential one wherein M initiates an action and I reacts to it, followed by M 's reaction to I 's action and so on. Sequential strategies avoid the problem of infinite regress.

Finally, since this paper has emphasized on general principles and overall plans, one may wish to consider specific cases and specialized examples. To avoid the routine nature of work which is characteristic of such endeavours, we have refrained from embarking on such activities. They certainly would expand the paper but will contribute little, if any, to its academic and intellectual content. Cutting new territory is a more attractive option, but as of now, we see little opportunity of doing so. Perhaps the legal profession can come up with other issues that need to be addressed using the tools of probability and risk analysis.

Acknowledgements

NDS acknowledges with deep gratitude the hospitality and the excellent environment provided by *The Santa Fe Institute*, Santa Fe, NM, during his tenure there as a visiting professor, Fall 2001, when much of the work reported here was done. Also acknowledged are the comments by a referee, that caused us to change the title of this paper. A suggestion by Professor Colin Aitken to include material that would appeal to lawyers has caused us to write Section 1. This research was supported by Grants DAAD19-02-1-0195 and DAAD19-01-0502 from the US Army Research Office.

REFERENCES

- ARROW, K. J. (1951) *Social Choice and Individual Values*. New York: John Wiley & Sons.
- BLISHKE, W. R. & MURTHY, D. N. P. (1996) *Product Warranty Handbook*. New York: Marcel Dekker.
- CHEN, J., LYNN, N. J. & SINGPURWALLA, N. D. (1995) Markov mesh models for filtering and forecasting with leading indicators. *Analysis of Censored Data*. Institute for Mathematical Statistics, Lecture Notes—Monograph Series, vol. 27. Hayward, CA: Institute for Mathematical Statistics, pp. 39–54.

- CHEN, J., LYNN, N. J. & SINGPURWALLA, N. D. (1996) Forecasting warranty claims. *Product Warranty Handbook* (W. R. Blishke & P. Murthy eds). New York: Marcel Dekker, pp. 803–816.
- COOKE, R. M. (1991) *Experts in Uncertainty: Opinion and Subjective Probability in Science*. New York: Oxford University Press.
- DEGROOT, M. H. (1974) Reaching a consensus. *J. Am. Stat. Assoc.*, **69**, 118–121.
- GENEST, C. & ZIDEK, J. V. (1986) Combining probability distributions: a critique and an annotated bibliography. *Stat. Sci.*, **1**, 114–135.
- KEENEY, R. L. & RAIFFA, H. (1976) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley.
- LAWLESS, J. F. (1998) Statistical analysis of product warranty data. *Int. Stat. Rev.*, **66**, 41–60.
- LINDLEY, D. V. (1983) Reconciliation of probability distributions. *Oper. Res.*, **31**, 866–880.
- SINGPURWALLA, N. D. (2000) Warranty contracts and equilibrium probabilities. *Statistical Science in the Courtroom* (J. L. Gastwirth ed.). New York: Springer, pp. 363–377.
- SINGPURWALLA, N. D. (2003) Warranty: a surrogate of reliability. *Mathematical Reliability: An Expository Perspective* (R. Soyer, T. Mazzuchi & N. D. Singpurwalla eds). Norwell, MA: Kluwer, pp. 317–333.
- SINGPURWALLA, N. D. & WILSON, S. P. (1992) Warranties. *Bayesian Statistics 4* (J. Berger, J. Bernardo, P. Dawid & A. Smith eds). Oxford: Clarendon Press, pp. 435–446.
- SINGPURWALLA, N. D. & WILSON, S. P. (1993) Warranties: their statistical and game theoretic aspects. *SIAM Rev.*, **35**, 17–42.
- SINGPURWALLA, N. D. & WILSON, S. P. (1998) Failure models indexed by two scales. *Adv. Appl. Probab.*, **30**, 1058–1072.