Pricing a motor extended warranty with limited time and usage cover

By Fidelis T Musakwa

Presented at the Actuarial Society of South Africa's 2013 Convention
31 October–1 November 2013, Sandton Convention Centre

ABSTRACT
Providers of motor extended warranties with limited usage often face difficulty evaluating the impact of usage limits on warranty price because of incomplete usage data. To address this problem, this paper employs a non-parametric interval-censored survival model of time to accumulate a specific usage. This is used to develop an estimator of the probability that a provider is on risk at a specific time in service. The resulting pricing model is applied to a truck extended warranty case study. The case study demonstrates that interval-censored survival models are ideal for use in pricing motor extended warranties with limited usage cover. The results also suggest that employing a usage rate distribution to forecast the number of vehicles on risk can be misleading, especially on an extended warranty with a relatively high usage limit.

KEYWORDS:
Motor extended warranty; two-dimensional warranty; risk premium; interval censoring; warranty cover

CONTACT DETAILS
Fidelis T Musakwa, Wits Business School, University of the Witwatersrand, PO Box 98, Wits 2050; 2 St David's Place, Parktown, Johannesburg, 2193, South Africa
Email: fidelismusakwa@yahoo.com

1. INTRODUCTION

Vehicle warranties compensate customers on covered parts that fail during a covered period (Wu, 2012). Normally, a base warranty is tied to a new vehicle sale (Murthy, 1992). A motor base warranty's cover period is often set on two parameters (1) age, and (2) usage. Such a cover period is commonly referred to as being two-dimensional (Murthy & Djamaludin, 2002). For automobiles, usage refers to accumulated distance travelled while for other vehicles, such as earthmoving equipment, usage refers to accumulated operating hours. Base warranties expire on reaching the age or usage limit, whichever occurs first. For example, a base warranty with a cover period of '24 months/200,000 kilometres' expires on the earlier of reaching age 24 months or accumulating a usage of 200,000 kilometres.

An extended warranty provides cover after the base warranty expires. Exceptions, however, exist on motor extended warranties that provide benefits excluded from the base warranty during the base warranty cover period (Hayne, 2007). Unlike a base warranty, a customer has a choice on whether to buy a motor extended warranty (Murthy & Djamaludin, 2002). Motor extended warranty customers are buyers of new and pre-owned vehicles. Providers of motor extended warranties include vehicle manufacturers, banks, insurers, and motor dealers (Li et al., 2012; Musakwa, 2012). Similar to a base warranty, time and/or usage are used to define a motor extended warranty's cover period. This paper focuses on motor extended warranties with cover period set on time and usage.

The main factors determining the cost of providing a warranty are: (1) cover period; (2) benefits provided; (3) claim frequency; and (4) claim severity (Rai & Singh, 2005). Quantifying the influence of these four factors on warranty cost can be difficult. For example, if the cover period is set on time and usage, then an extended warranty provider's exposure to risk at a specific time is unknown because the provider has partial knowledge of accumulated usage on covered vehicles (Cheng, 2002). Such challenges have so far been mostly addressed through base warranty studies and few extended warranty studies (Jack & Murthy, 2007). This is despite some unique features associated with motor extended warranty providers: for example, contract terms and conditions; and data available for use in pricing (Shafiee et al., 2011).

Key questions on pricing motor extended warranties remain unanswered. If cover period is set on time and usage, how can a warranty provider estimate the probability of being on risk at a specific time in service? How does variation of vehicle age and accumulated usage, at point of extended warranty sale, influence the provider's exposure to risk? Can a usage rate distribution be reliably used to forecast the number of vehicles on risk? Above all, how do answers to the foregoing questions influence the 'fair' price of a motor extended warranty? Motivated by the need to answer these questions, this study develops a motor extended warranty risk premium model. Here, risk premium means the undiscounted expected claims cost emerging during the cover period. Pricing factors ignored include tax, commission, investment income, contingency loading, profit loading and expenses.

This study adds to the body of knowledge on pricing motor extended warranties in four ways. First, the study develops an estimator of the probability that a provider is on risk at a specific time in service on a covered vehicle. Doing so enhances clarity on assessing the impact of base warranty and extended warranty usage limits on the extended warranty's cost. Second, the study develops a method to estimate claim severity that employs past claimed amount data. This better captures the effect of extended warranty deductibles and limits of liability on the risk premium. Third, a unique design of employing a non-parametric interval-censored survival model is utilised to directly measure the probability distribution of time to accumulate a specific usage. The study shows how to structure incomplete usage data to estimate such a survival function. Additionally, the study demonstrates that given incomplete usage data, an interval-censored survival model provides knowledge on the distribution of time to accumulate a specific usage without relying on usage rate assumptions. Such a survival model is beneficial because it considers variability of usage (1) within an individual vehicle, and (2) across a population of vehicles under extended warranty cover. Fourth, case study results indicate that employing a usage rate distribution to forecast the number of vehicles on risk can be misleading, especially on an extended warranty with a relatively high usage limit. This is despite observing that some positively skewed statistical distributions fit well to usage rate data.

The paper proceeds as follows. Section 2 reviews previous research on pricing warranties. Section 3 develops a risk premium model for a motor extended warranty with limited usage cover. Section 4 applies the model to a case study of a truck extended warranty. Finally, section 5 concludes.

2. PREVIOUS RESEARCH

This section reviews literature on two fundamental factors influencing a warranty's risk premium: namely, (1) claim severity; and (2) exposure at risk. Analysing claim severity involves projecting costs of rectifying failures. Exposure at risk provides a unit of measuring risk (Antonio et al., 2010). Expressing claim severity per exposure unit provides a way of calculating a warranty's risk premium for a given cover period.

2.1 Claim Severity

Most extended warranty studies estimate claim severity using claims incurred data; that is, paid and outstanding authorised claims (Hayne, 2007; Cheng 2002; Weltmann & Muhonen, 2002). Using claims incurred to quantify claim severity is particularly appropriate for (1) setting reserves and (2) reviewing premiums on extended warranties whose terms and conditions remain unaltered going forward. However, claims incurred data may be problematic if pricing an extended warranty with different terms and conditions from contracts underlying the claims incurred data. Extended warranty terms and conditions that influence claim severity include: (1) the level of deductibles and limits on covered components; (2) the set of covered components; (3) the causes of failure covered; and (4) the method to rectify failures (e.g., replacement or minimal repair).

To avoid potential flaws of using claims incurred when pricing extended warranties, this study employs invoiced claim amount data to estimate claim severity. The invoiced claim amount is the claimed amount, on a component, invoiced when a claim is reported. Utilising invoiced claim amount data has multiple benefits. Firstly, it is free from the effect of deductibles and limits. This provides a good understanding of how applying various levels of limits and/or deductibles impacts claim severity. Secondly, it can be used to assess claim severity in instances where the provider's liability is conditional on the cause of failure. For example, a provider may be liable to a constant fraction of a claim stemming from damages caused by normal use. The use of invoiced claim amount enables one to assess the sensitivity of claim severity to changes in this constant fraction.

Wu (2012) points out that the predictive importance of past claims data decreases with a backward movement in time. Therefore, it is sensible to assign less weight to relatively older observations. Only recently has such a weighting method been applied in the warranty literature. For example, Wu and Akbarov (2011) apply such weighting to forecast the number of warranty claims. But the same weighting principle has yet to be applied to model claim severity. In the spirit of Wu and Akbarov (2011), this study assigns weights that decrease with a backward movement in time to forecast cost per exposure unit.

2.2 Modelling Vehicle Exposure at Risk

Vehicle warranty exposure is unitised on either time or usage. Kerper and Bowron (2007) unitise exposure on usage, which they define as distance travelled. Kalbfleisch et al. (1991) and Lawless (1998) use time a vehicle is on risk as the exposure unit. The appropriateness of a particular exposure unit depends on how the warranty cover period is specified. This is straightforward for warranties with a cover period set on only time or usage: unitise exposure in time (usage) if cover period is set on time (usage). The suitable exposure unit is unclear on warranties whose cover period is set on both time and usage. In such circumstances, Kerper and Bowron (2007) argue that usage is ideal because claim occurrence closely matches usage more than age. In contrast, Majeske (2007) recommends measuring exposure on the time dimension because it is relatively simple and a provider can track vehicle age regardless of whether usage is observed on that vehicle.

Most studies projecting vehicle population at risk over time allow for warranties expiring because of vehicles exceeding the usage limit (Su & Shen, 2012). These projections often rely on a usage rate statistical distribution, for example, Weibull (Jung & Bai, 2007), Lognormal (Alam & Suzuki, 2009; Rai & Singh, 2005) and Gamma distribution (Su & Shen, 2012; Majeske, 2007). Other studies discretise the usage rate distribution; for example classifying drivers into low, medium and high usage rate categories (Shahanaghi et al., 2013; Cheng & Bruce, 1993). The vehicle population at risk at a specific time in service is then elicited from the usage rate distribution assuming that usage rates are constant on a vehicle but vary across vehicles (Wu, 2012).

Such a premise has the advantage of simplifying the modelling process. However, several factors undermine the validity of assuming a constant usage rate on a vehicle. Examples include change in vehicle ownership and application. A vehicle can also be idle for some period. Overall, the questionable validity of assuming a constant usage rate implies that it remains largely unknown whether usage rate distributions are fit for the purpose of forecasting the distribution of time to accumulate a certain usage. This paper is a step towards addressing this knowledge gap by directly modelling time to accumulate a specific usage.

3. THE MODEL

This section formulates a risk premium model of a motor extended warranty with cover set on both time and usage. It starts by developing an estimator of the probability that a provider is on risk at a specific time in service. This is followed by a discussion of how interval-censored survival models contribute towards estimating the exposure probability. Next, section 3.3 discusses how exposure probabilities are determined in the special case of a warranty extended on only usage. Section 3.4 discretises output from an interval-censored survival model of time to a specific usage. This is important for risk premium models defined in discrete time. Section 3.5 estimates claim severity from past invoiced claim amount data. Finally, section 3.6 calculates the risk premium.

3.1 Estimating an Extended Warranty Provider's Exposure Probability

To develop an estimator of a warranty provider's probability of being on risk at a specific time in service, this section separately assesses the provider's exposure probability on one cover dimension while ignoring the other cover dimension. That is, the provider's exposure probability is first calculated assuming that cover period is set only on time. Next, the provider's exposure probability is calculated assuming that cover is set on only usage. Finally, the two sets of exposure probabilities are combined to obtain the provider's exposure probability at a specific time in service. The logic of doing so is first demonstrated graphically and then expressed mathematically. In what follows, usage is characterised as accumulated distance travelled. The concepts nonetheless apply to other definitions of accumulated usage.

Consider Figure 1, which shows a scenario of a base warranty and extended warranty cover period set on time and usage. If we ignore the base warranty's limit on usage, then the extended warranty provider's exposure probability with time is deterministic. That is, extended warranty cover begins on the extended warranty start date and ends on the extended warranty end date. The dashed line on the interval $[t^*, \hat{t}]$ reflects the cumulative density function of time to attain the base warranty usage limit. From time $\hat{t}$ onwards, all vehicles on risk are projected to have an accumulated usage that exceeds the base warranty usage limit. The dashed line from time $\tilde{t}$ onwards reflects the survival time to attain the extended warranty usage limit. In essence, the base warranty usage limit increases the extended warranty provider's exposure probability on the interval [t^*, extended warranty start date]. In contrast, the extended warranty usage limit reduces the extended warranty exposure probability on the interval [$\tilde{t}$, extended warranty end date]. The resulting extended warranty provider's exposure probability is that it increases from time t^* and becomes one on the extended warranty start date. The exposure probability remains equal to one until time $\tilde{t}$. Thereafter, exposure probability decreases as some vehicles exceed the extended warranty usage limit. After the extended warranty expiry date, the extended warranty provider's exposure probability reduces to zero.

Formally, let $p(t)^{Time}$ denote the extended warranty provider's exposure probability at time t, ignoring coverage on usage. That is:

$$p(t)^{Time} = \begin{cases} 1, \text{ Extended warranty start date} \leq t \leq \text{Extended warranty end date} \\ 0, \text{ Otherwise} \end{cases} \tag{1}$$

Let $p(t)^{Usage}$ denote the extended warranty provider's exposure probability at time t, ignoring coverage set on the time dimension. Suppose we can observe how usage accumulates over time on vehicles $j = 1,2,\ldots, N$. Further, assuming that cover is set only on usage, let $I_j(t)$ be an indicator variable for the status of whether the accumulated usage on vehicle j, at time t, is within extended warranty usage cover. Following Su and Shen (2012), $p(t)^{Usage}$ is calculated as:

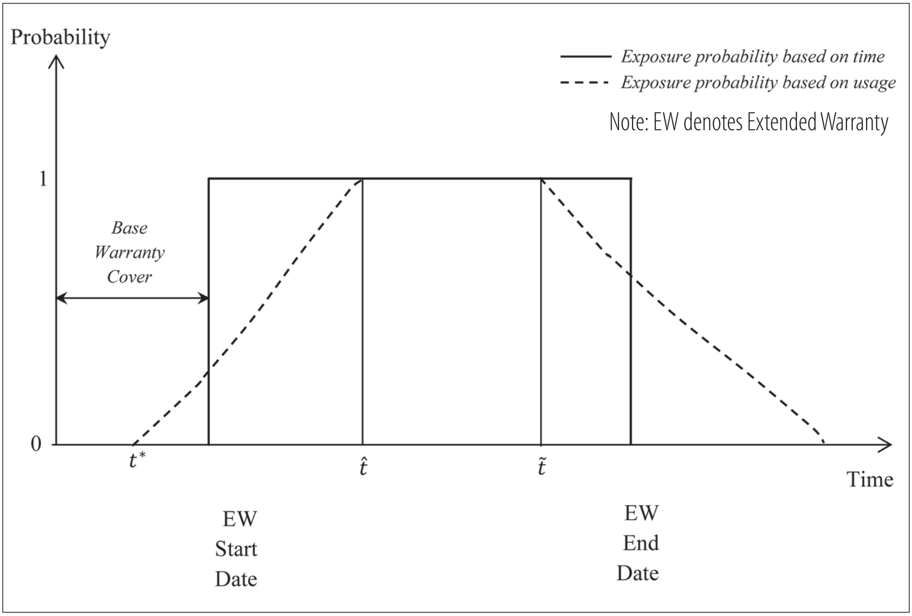

FIGURE 1 Example of base warranty and extended warranty usage limits with an impact on the extended warranty risk premium

$$p(t)^{Usage} = \frac{\sum_{j=1}^{N} I_j(t)}{N} \tag{2}$$

In Figure 1, $p(t)^{Usage}$ is represented by the dashed line: $p(t)^{Usage}$ increases on the interval $[\,t^*, \hat{t}\,]$ as vehicles come on risk due to exceeding the base warranty usage limit; $p(t)^{Usage}$ equals one on the interval $[\,\hat{t}, \tilde{t}\,]$ as all vehicles are projected to be within the extended warranty usage cover; and $p(t)^{Usage}$ starts decreasing from $\tilde{t}$ onwards until it reaches zero as vehicles exceed the extended warranty usage limit.

Let $p(t)$ denote the motor extended warranty provider's probability of being on risk on a covered vehicle at time t. Let t^{BW} denote the base warranty expiry time, which coincides with the starting time of the motor extended warranty. In Figure 1, t^{BW} is depicted as EW start date. Let t^{EW} denote the extended warranty expiry time. In Figure 1, t^{EW} is depicted as EW end date. Intuitively, the estimator $p(t)$ can be deduced from $p(t)^{Time}$ and $p(t)^{Usage}$ as follows:

$$p(t) = \begin{cases} p(t)^{Usage}, & t \le t^{BW} \\ 1, & (t^{BW} < t \le t^{EW}) \text{ and } (t < \hat{t}) \\ \min\left(p(t)^{Usage}, p(t)^{Time}\right), & \left(t^{BW} < t \le t^{EW}\right) \text{ and } \left(t \ge \hat{t}\right) \\ 0, & \text{Otherwise} \end{cases} \tag{3}$$

3.2 Modelling the Survival Time to Accumulate a Specific Usage

Section 3.1 used $p(t)^{Usage}$ to derive the estimator of the probability that a motor extended warranty provider is on risk at time in service t. This section proceeds to estimate $p(t)^{Usage}$ by modelling a vehicle's time to accumulate a specific usage using a non-parametric interval-censored survival model. A distinct property of interval-censored survival models is that they estimate the survival function using data on intervals containing the time to event of interest. This applies when exact survival times are unobserved. Motor extended warranty providers are in such a scenario regarding the survival time of a vehicle on risk to reach the extended warranty's usage limit.

Motor extended warranty providers often observe accumulated usage when warranty holders submit claims. Other usage data sources include: point of sale, cancellation of warranty (Kerper & Bowron, 2007) and follow-up studies (Karim & Suzuki, 2005). As a whole, the time that extended warranty providers observe accumulated usage is random. These observation times are censoring times. To illustrate, consider the following example. Suppose a warranty provider is interested in knowing the time in service that a vehicle under warranty accumulates a usage of 200,000 kilometres. From available usage records, it may be impossible to observe the exact age that a vehicle reaches 200,000 kilometres. Instead, the data may indicate the age interval that the vehicle accumulated 200,000 kilometres. In sum, motor extended warranty providers have data on time to accumulate a specific usage that is interval-censored with random censoring times.

To formulate $p(t)^{Usage}$ as an interval-censored survival model, consider the following experimental design. Let T_U denote the age that a vehicle accumulates a usage of U; and U_t denotes the accumulated usage at age t. Note, U is a constant but U_t is a random variable. The survival function of T_U is:

$$S(t) = \Pr(T_U > t) = \Pr(U_t < U) \tag{4}$$

Thus, the cumulative density function (CDF) of T_U is $F(t) = 1 - S(t) = \Pr(U_t \geq U)$. Instead of observing T_U, a motor extended warranty provider observes the age interval containing T_U. In continuous time, it is impossible to observe the exact age that a vehicle accumulates a usage of U. Following Adamic et al. (2010) let the age interval containing T_U for vehicle i be $(L_i, R_i]$; that is, $\{T_U : L_i < T_U \leq R_i\}$. In words, L_i, the left censoring time, is the furthest observed age when accumulated usage was less than U. In contrast, R_i, the right censoring time, is the earliest observed age when accumulated usage exceeded U. Figure 2 presents examples of a vehicle's age intervals containing time to accumulate 400,000 kilometres, 600,000 kilometres and 800,000 kilometres.

To derive the likelihood function of $S(t)$, let B denote the set of distinct ordered elements from the union of $\{L_i : i = 1,...,n\}$ and $\{R_i : i = 1,...,n\}$. From B, one can obtain a set of disjoint intervals, $\{(p_1, q_1], (p_2, q_2], ..., (p_m, q_m]\}$ where $0 < p_1 < q_1 < p_2 < q_2 < ... < q_m < \infty$. These disjoint intervals are sometimes referred to as Turnbull's innermost intervals, in tribute to Turnbull (1976), or places of maximal cliques in graph theory (Gentleman & Vandal, 2001). For a non-parametric approach to interval-censored survival analysis, weights are assigned to the set of Turnbull's innermost intervals (Dehghan & Duchesne, 2011). Let $w_j \in [0,1)$ denote the weight assigned to Turnbull's innermost interval j, subject to the condition:

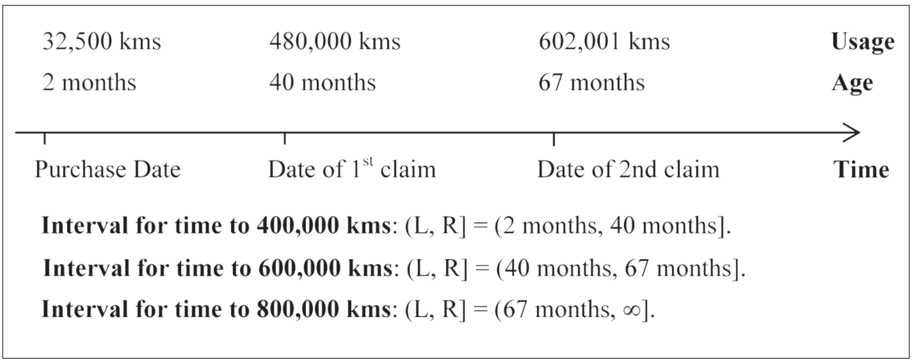

Figure 2 Examples of intervals containing time to reach 400,000 kms, 600,000 kms and 800,000 kms

$$\sum_{j=1}^{m} w_j = 1 \tag{5}$$

To determine subject *i*'s input to the weight of Turnbull's innermost interval *j*, let α_{ij} be an indicator variable defined as follows:

$$\alpha_{ij} = \begin{cases} 1, & \left(p_j, q_j\right] \subseteq \left(L_i, R_i\right] \\ 0, & \text{Otherwise} \end{cases} \tag{6}$$

Assuming that: (1) censoring is non-informative; and (2) the *n* observations are independent and identically distributed (IID), the likelihood of the survival function is proportional to (Zhang & Sun, 2010):

$$\text{Likelihood} = \prod_{i=1}^{n} \left[S\left(L_i\right) - S\left(R_i\right) \right] = \prod_{i=1}^{n} \left[\sum_{j=1}^{m} \alpha_{ij} w_j \right] \tag{7}$$

Non-informative censoring means that the distribution of T_U is independent of the time that usage is observed. The IID assumption is intuitive since vehicles accumulate usage independently of each other. Assuming an identical distribution of T_U across vehicles eligible to a specific motor extended warranty provides the mathematical convenience of characterising variability of T_U across vehicles using a single function, $S(t)$.

The Non-Parametric Maximum Likelihood Estimator (NPMLE) is the solution that maximises the likelihood. The corresponding NPMLE survival function is (Sun, 2001):

$$\hat{S}\left(t\right) = \begin{cases} 1, & t < p_1 \\ 1 - \sum_{k=1}^{j} w_k, & q_j < t < p_{j+1} \qquad \left(1 \le j \le m-1\right) \\ 0, & t > q_m \end{cases} \tag{8}$$

Generally, a survival function is defined for all $t \geq 0$. The NPMLE, however, is indeterminate on the innermost intervals, $\left\{\left(p_1, q_1\right], \left(p_2, q_2\right], \ldots, \left(p_m, q_m\right]\right\}$. This non-uniqueness implies that the NPMLE is indifferent to how probability mass is distributed on the innermost intervals (Maathuis & Hudgens, 2011). An arbitrary method to distribute mass on innermost intervals is nonetheless necessary to uniquely define the NPMLE for all $t \geq 0$.

Many NPMLE methods exist. These methods are mostly iterative since a closed form NPMLE solution is non-existent. Examples of iterative NPMLE methods are: the self-consistent (SC) algorithm (Turnbull, 1976); the expectation-maximization (EM) algorithm (Dempster et al., 1977); the Iterative Convex Minorant (ICM) algorithm (Groeneboom & Wellner, 1992); the hybrid EM-ICM algorithm (Wellner & Zhan, 1997); and the self-reduction algorithm (Groeneboom et al., 2008). The SC algorithm is relatively simple. However, it tends to converge at a relatively slow rate (Gómez et al., 2009). Moreover, there can be multiple SC solutions, thereby failing to characterise the NPMLE solution (Gentleman & Geyer, 1994). In contrast, the EM-ICM solution is the global NPMLE (Wellner & Zhan, 1997). The EM-ICM algorithm iterates relatively fewer times to converge. Despite iterating fewer times, computing time may still be substantial and comparable to the EM algorithm (Gómez et al., 2009).

Imputing interval-censored data is an example of a non-iterative NPMLE method. This method involves employing the Kaplan–Meier product limit estimator on a dataset where interval-censored data is converted to complete and right censored data of time to an event of interest (Hsu et al., 2007). Such imputing methods produce reasonable estimates of $S(t)$ on data with mostly short intervals. However, imputing can result in biased estimates of $S(t)$ if lifetime data has mostly wide intervals. Classifying intervals into short or long partly depends on the context, thereby increasing the risk of bias. Law and Brookmeyer (1992) considered censoring intervals of about 24 months and less as short in a study of AIDS. The same lifetime interval may, however, be considered long in other lifetime studies. As a result, imputing methods are potentially misleading.

3.3 Exposure Probabilities when only extending Usage Cover

Formulating $p(t)^{Usage}$ with a survival model helps determine exposure probabilities at a specific time in service when extending a base warranty only by usage. An example is extending a base warranty from a cover of '36 months/120,000 kilometres' to '36 months/200,000 kilometres'. To this end, let U^{BW} and U^{EW} denote the accumulated usage limit for the base warranty and extended warranty, respectively. For the usage only extension example above, $U^{BW} = 120{,}000$ kilometres and $U^{EW} = 200{,}000$ kilometres. The probability that a vehicle's accumulated usage is within U^{BW} is $S_{BW}(t) = \Pr\left(U_t < U^{BW}\right)$. Likewise, the probability that a vehicle's accumulated usage is within U^{EW} is $S_{EW}(t) = \Pr\left(U_t < U^{EW}\right)$. The probability of a motor extended warranty provider being on risk at time t is thus:

$$p(t) = p(t)^{Usage} = S_{EW}(t) - S_{BW}(t) \tag{9}$$

Equation (9) ensures that a warranty extended on only usage can be priced by a model where exposure is indexed on time.

3.4 Discretising the Survival Function of Time to accumulate a Specific Usage

For consistency, both exposure time and survival function of time to attain a specific usage should be in either discrete or continuous time. Up to this point, the survival function of time to reach a specific usage has been modelled in continuous time. Conversely, exposure has been modelled in discrete time. To address this inconsistency, this section uses the continuous time survival function to obtain its discrete time counterpart.

Let $q(k)$ denote the probability of attaining a specific usage in the kth time interval. This probability is obtained from the continuous time survival function as follows:

$$q(k) = \hat{S}(k-1) - \hat{S}(k) \qquad k = 1,2,... \tag{10}$$

Equation (10) assumes that:

$$\hat{S}(0) = 1 \tag{11}$$

The survival function in discrete time, $\hat{S}^{Discrete}(k)$, is thus:

$$\hat{S}^{Discrete}(k) = 1 - \sum_{j=1}^{k-1} q(j) \tag{12}$$

3.5 Modelling Claim Severity

As mentioned earlier, this study uses claimed amount to estimate claim severity. To this end, consider the following notation:

C_r : Claim severity for component *r*.
Y_r : Claimed amount for component *r*.
$h(Y_r)$: Probability density function for Y_r.
$H(Y_r)$: Cumulative density function for Y_r.
L_r : Limit of cover for component *r*.

Y_r can be easily adjusted for past inflation. The claim severity estimate for component *r* is subsequently estimated as follows:

$$C_r = \begin{cases} Y_r, & Y_r < L_r \\ L_r, & Otherwise \end{cases} \tag{13}$$

The corresponding expected claim severity for component *r* is:

$$E(C_r) = \underbrace{\int_0^{L_r} y_r h(y_r) dy_r}_{\text{Expected claim severity of claims below limit}} + \underbrace{L_r \{1 - H(L_r)\}}_{\text{Expected claim severity of claims above limit}} \tag{14}$$

This has the advantage of reviewing the impact of changing the limits on claim severity. A similar concept to Equation (14) can be applied when analysing the impact of deductibles on claim severity.

Claim severity statistics can be obtained through simulation. This involves randomly selecting an observation from the claimed amount distribution. Next, Equation (13) is used to obtain the corresponding claim severity. This process is repeated many times, resulting in a simulated claim severity distribution, where statistics can be extracted.

3.5 Determining the Motor Extended Warranty Risk Premium

Cheng and Bruce (1993) point out that the undiscounted risk premium is the sum of incremental costs per exposure over the cover period. Unlike Cheng and Bruce (1993), this study takes account of the fact that more recent observations may be more predictive than earlier observations (Wu & Akbarov, 2011) when estimating an incremental cost per exposure month. Let *RP* denote an extended warranty's undiscounted risk premium; *t* be an index for exposure month; *j* be an index for calendar period; $C_{t,j}$ be the average claim severity for exposure month *t* based on data from calendar period *j*; $E_{t,j}$ be the total exposed to risk for exposure month *t* based on data from calendar period *j*; $\lambda_j \in (0,1)$ be the weight assigned to calendar period *j*. Then an extended warranty's risk premium can be calculated as follows:

$$RP = \sum_t \underbrace{\left(\sum_j \lambda_j \frac{C_{t,j}}{E_{t,j}} \right)}_{\text{Weighted Cost per Exposure}} \times \underbrace{p(t)}_{\text{Chance of Exposure}}, \qquad \sum_j \lambda_j = 1 \tag{15}$$

The weights, λ_j, decrease with a backward movement in calendar period.

Stochastic methods can capture the effect of extended eligibility, which results in uncertainty on vehicle age and corresponding cumulated usage at point of sale. Extended eligibility means that an extended warranty designed for new vehicles can be bought on a used vehicle provided the vehicle's base warranty cover is still in force (Hayne, 2007). Shafiee et al. (2011) propose a bivariate function, such as the Beta-Stacy distribution, to model the joint uncertainty of a vehicle's age and accumulated usage at point of sale. While a bivariate distribution is mathematically convenient, there may be instances where the extended warranty being priced lacks past data. In such instances, a sensible approach is to generate a bivariate distribution guided by the sales profile anticipated by those responsible for implementing the warranty sales and marketing strategy. The bivariate distribution is then used to randomly select a vehicle age and accumulated usage at point of sale. Next, the corresponding risk premium is calculated using Equation (15). This process is repeated several times to produce a distribution of the extended warranty's risk premium.

4. CASE STUDY

The case study had three objectives.

(1) To present a practical situation where a provider's exposure probabilities were calculated to price a motor extended warranty whose cover period was defined in terms of time and usage. The section omits details on how cost per exposure month, risk premium and gross premium were calculated. See Cheng and Bruce (1993) for a numerical example on deriving the gross premium from the risk premium.
(2) To compare survival times from a survival model conditional on a usage rate distribution against those from a non-parametric interval-censored survival model.
(3) To check if usage rate data fits well to some commonly used statistical distributions.

4.1 Data

The case study employed proprietary data from an insurer located in South Africa. The data to model usage and exposure probabilities was from 857 heavy commercial vehicles involved in medium and long-haul operations. A heavy commercial vehicle was defined as a commercial truck with a gross vehicle mass ranging from 8,501 kilograms to 16,500 kilograms. From the data, it was impossible to discern whether a truck was utilised for medium or long-haul operations. Usage data was obtained from four sources: sales, claims, withdrawals and maintenance records.

4.2 Pricing Objectives

The case study objective was to price extensions of various manufacturer warranty covers to '84 months/800,000 kilometres', inclusive of the manufacturer warranty. The manufacturer warranty options considered were: '24 months/200,000 kilometres', '36 months/400,000 kilometres', '36 months/600,000 kilometres' and '60 months/ 800,000 kilometres'.

4.3 Exposure Probabilities

Figure 3 presents the NPMLE of the interval-censored survival function of time to accumulate usage of 200,000 kilometres, 400,000 kilometres, 600,000 kilometres and 800,000 kilometres. These survival functions were estimated in R using a package called 'interval'. This package estimates the survival function using the expectation-

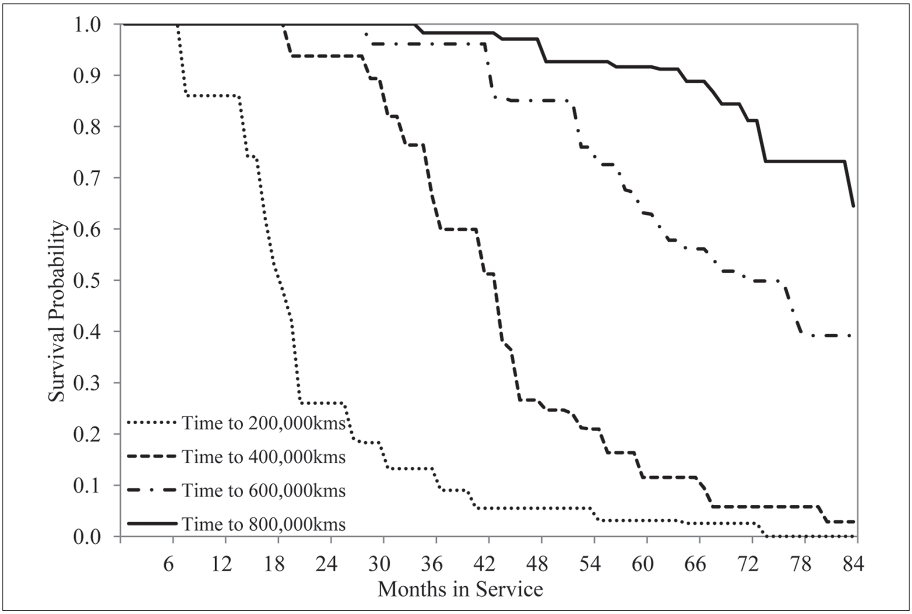

FIGURE 3 NPMLE of time to reach a specific usage on medium to long haul trucks

maximization algorithm by Dempster et al. (1977). The 'interval' package also checks if the NPMLE solution satisfies Kuhn–Tucker conditions necessary for global convergence (Fay & Shaw, 2010).

Unlike the conventional stepped survival curve of the product-limit estimator, an interval-censored survival function has periods where the survival function is indeterminate. These are the non-horizontal parts of survival curves in Figure 3. To thus fully specify the survival function over month in service, Figure 3 linearly interpolates the survival curve where it is indeterminate.

Figure 4 presents the estimated probability that an extended warranty provider is on risk when extending various manufacturer warranty terms to '84 months/800,000 kilometres'. For instance, the '24 months/200,000 kilometres' line depicts the probability that an extended warranty provider is on risk when extending a base warranty of '24 months/200,000 kilometres' to '84 months/800,000 kilometres', inclusive of the base warranty. For each warranty extension, the line becomes vertical at the manufacturer warranty expiry date. All the warranty extensions have their exposure probability line converging to a line matching the survival time to 800,000 kilometres. Exposure before the base warranty expiry date results from vehicles attaining the base warranty usage cover limit before the base warranty expiry date. For extending the '60 months / 800,000 kilometres' base warranty, only a fraction of vehicles come on risk on the base warranty expiry date. This is because some vehicles are anticipated to have reached the extended warranty usage limit by the base warranty expiry date.

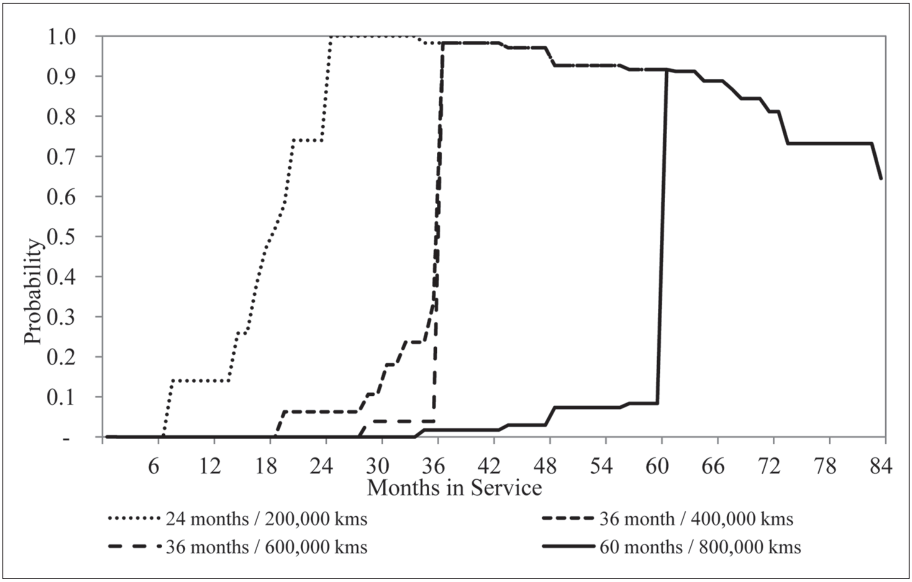

FIGURE 4 Probability of extended warranty provider being on risk

Note. The warranty extension increases cover of the illustrated base warranties to 84 months/800,000 kilometres.

4.4 Comparison of using NPMLE versus Monthly Usage Rates

As mentioned before, motor warranty literature mostly models the survival time to accumulate a specific usage conditional upon a usage rate distribution. This section denotes such a survival function by $S(t|g(u))$, where g(u) is the usage rate probability density function. Following Shafiee et al. (2011), Jack et al. (2009) and Rai and Singh (2005), this study calculated usage rate using Equation (16).

$$\text{Usage rate} = \frac{\text{Accumulated usage}}{\text{Age}} \tag{16}$$

The survival time to accumulate a specific usage is then linearly interpolated from the usage rate distribution. For each vehicle, only the furthest observed accumulated usage and the corresponding age are used to calculate its usage rate.

An alternative method to estimate the survival time to accumulate a specific usage is the non-parametric interval-censored survival model, denoted by $S_{NPMLE}(t)$. This paper posits that $S_{NPMLE}(t)$ reflects the true survival time since it directly employs the interval truly known to contain the survival time of interest. The data source used to estimate $S(t|g(u))$ and $S_{NPMLE}(t)$ is the same: for usage and corresponding age data of a vehicle, $S(t|g(u))$ extracts a usage rate, while $S_{NPMLE}(t)$ extracts an interval containing the time to event of interest.

This section investigates the statistical significance of the difference in output from $S_{NPMLE}(t)$ and $S(t|g(u))$. That is, this section tests the following hypothesis:

$$H_0 : S_{NPMLE}(t) = S(t|g(u)) \quad \forall t$$
$$H_{alt} : S_{NPMLE}(t) \neq S(t|g(u)) \quad \forall t$$

This hypothesis is tested using three weighted log-rank tests; namely, Sun (1996) score test, Finkelstein (1986) score test and generalised Wilcoxon–Mann–Whitney test by Fay and Shaw (2010). These tests are chosen because they can handle interval-censored survival models. They are also distribution-free, implying that they rely on less restrictive assumptions about the data compared to parametric tests. Moreover, weighted log-rank tests are robust (Fay & Shaw, 2010). The weighted log-rank tests are conducted on survival time to accumulate usage of 200,000 kilometres, 400,000 kilometres, 600,000 kilometres and 800,000 kilometres. Doing so provides insight on how $S_{NPMLE}(t)$ and $S(t|g(u))$ differs with respect to age to accumulate various levels of usage.

Table 1 Comparison of NPMLE against monthly usage rates survival curves

	Sun (1996) Score Test		Finkelstein (1986) Score Test		Fay and Shaw (2010) Monte Carlo Wilcoxon Test		
	Score Statistic	*p*-value	Score Statistic	*p*-value	Score Statistic	*p*-value	*p*-value 99% CI
Age at 200,000 kms	7.795	0.455	8.468	0.423	10.171	0.094	[0.093, 0.096]
Age at 400,000 kms	11.125	0.293	11.259	0.291	11.625	0.062	[0.061, 0.063]
Age at 600,000 kms	13.362	0.197	13.528	0.195	11.990	0.041	[0.040, 0.041]
Age at 800,000 kms	21.403	0.021	21.415	0.022	15.818	0.008	[0.007, 0.008]

Note: The 99% confidence intervals of the Fay and Shaw (2010) Monte Carlo Wilcoxon test are drawn from 100,000 Monte Carlo resamples.

Table 1 presents the weighted log-rank test results. The difference between $S_{NPMLE}(t)$ and $S(t|g(u))$ is statistically not significant at the 5% level for age at 200,000 kilometres, 400,000 kilometres and 600,000 kilometres. However, difference between $S_{NPMLE}(t)$ and $S(t|g(u))$ is statistically significant at the 2.5% level for age at 800,000 kilometres. Notably, each log-rank test's p-value decreases with age at a higher accumulated usage. A plausible explanation of these results is that usage rates vary with time, thereby violating the constant usage rate assumption underpinning the $S(t|g(u))$ approach. Although the findings are from a small sample relative to the entire population of trucks, the results nonetheless suggest that the reliability of $S(t|g(u))$ decreases with age at a higher accumulated usage.

Figure 5 plots $S_{NPMLE}(t)$ and $S(t|g(u))$ for age at 800,000 kilometres. The median time to accumulate 800,000 kilometres was 80 months under $S(t|g(u))$, but was 96 months under $S_{NPMLE}(t)$. Overall, the survival time to attain 800,000 kilometres inferred from using a usage rate distribution was lower than that deduced from a non-parametric interval-censored survival model, particularly after 60 months in service. This suggests that when considering relatively high usage levels, such as 800,000 kilometres on truck warranties, $S(t|g(u))$ is likely to underestimate the number of vehicles on risk at a specific month in service.

4.5 Goodness of Fit Test of Usage Rate Distributions

The Anderson–Darling test is used to test if the Lognormal, Gamma and Weibull distribution fit well to usage rate data. These three distributions are selected because they are commonly used to model usage rates. The method of moments estimator is used to estimate distribution parameters. Attractive properties of the Anderson–Darling test include: (1) it is non-parametric, thereby applicable to testing goodness of fit of various statistical distributions; and (2) unlike the Kolmogorov–Smirnov test, the Anderson–Darling test can be used when parameters of a theoretical distribution are estimated from sample data.

TABLE 2 Anderson–Darling goodness of fit test results

Distribution	Parameter Estimates	Anderson–Darling Test Statistic	*p*-value
Lognormal (μ, σ)	$\mu = 8.96$; $\sigma = 0.44$	0.38	0.39
Gamma (α, β)	$\alpha = 8.24$; $\beta = 0.98 \times 10^{-3}$	0.27	0.66
Weibull (α, β)	$\alpha = 3.22$; $\beta = 9{,}614.29$	0.25	0.73

Note: For Weibull and Gamma distribution, alpha is the shape parameter and beta is the scale parameter.

Table 2 presents the Anderson–Darling test results, which indicate that the Lognormal, Gamma and Weibull distribution fit well to usage rate data (p-values > 30%). This reaffirms findings from other studies showing that positively-skewed statistical distributions are suitable to model usage rates (Shahanaghi et al., 2013; Su & Shen, 2012; Jung & Bai, 2007; Kerper & Bowron, 2007; Majeske, 2007; Rai & Singh, 2004). Section 4.4 nonetheless showed that the reliability of usage rate modelling on estimating an extended warranty provider's exposure probability decreases as the usage cover considered increases. This highlights that a good fit of a statistical distribution to usage rate data is neither a necessary nor sufficient condition for knowledge about the survival time to accumulate a specific usage.

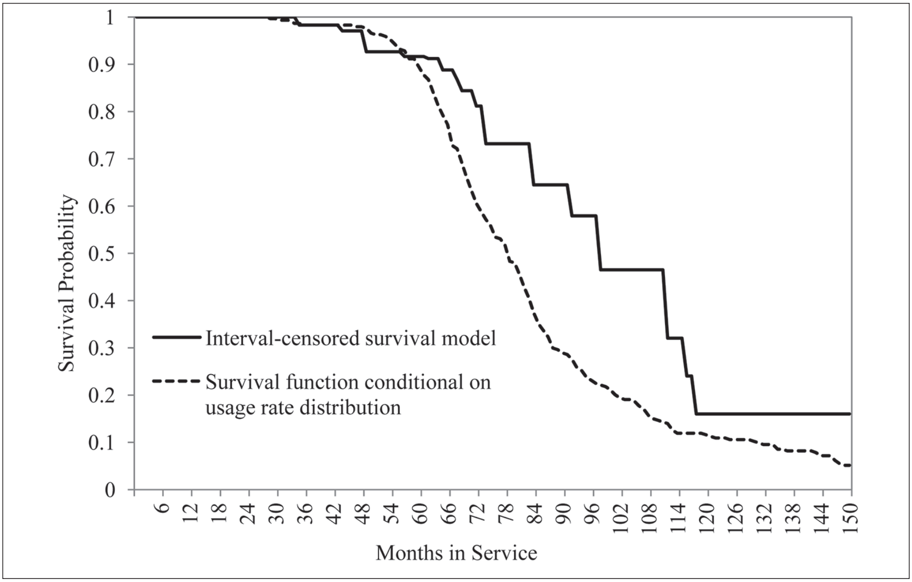

FIGURE 5 NPMLE versus usage rate survival function for age at 800,000 kms

5. CONCLUDING REMARKS AND FUTURE RESEARCH

This paper sought to estimate the risk premium of a motor extended warranty whose cover is limited by time and usage. The effect of limiting usage on an extended warranty's risk premium was captured by determining the provider's probability of being on risk at a specific time in service. The study demonstrates that interval-censored survival models can suitably be employed to estimate such exposure probabilities, especially given that extended warranty providers often have incomplete data on how usage accumulates with time. Case study findings show that the reliability of employing usage rate distributions to elicit the vehicle population at risk at a specific time in service decreases with the level of usage cover considered. If usage cover limits are relatively high, then estimating exposure probabilities by employing a usage rate distribution tends to increase the chance of underestimating the risk premium.

The study can be extended in various ways. A useful area of research is to investigate if the distribution of time to attain a specific usage is stable over time. This provides insight on the reliability of employing such a distribution, estimated from past data, to estimate the risk premium on motor extended warranties sold to new customers. Another worthwhile area of research would be to include other factors when estimating the provider's probability of being on risk at a specific time in service. Examples of such factors include: withdrawal, theft and accident. A shortcoming of this study is that it does not incorporate factors explaining variations in the time to attain a specific usage. To address this weakness, it is worthwhile to consider an interval-censored survival function with covariates. This can be achieved by using, for example, an interval-censored proportional hazard model.

ACKNOWLEDGEMENTS

This paper benefitted from discussions and comments from Sias Esterhuizen.

REFERENCES

Adamic, P, Dixon, S & Gillis, D (2010). Multiple decrement modeling in the presence of interval censoring and masking. *Scandinavian Actuarial Journal*, 2010(4):312–27

Alam, MM & Suzuki, K (2009). Lifetime estimation using only failure information from warranty database. *IEEE Transactions on Reliability*, **58**(4):573–82

Antonio, K, Frees, EW & Valdez, EA (2010). A multilevel analysis of intercompany claim counts. *ASTIN Bulletin*, **40**(1):151–77

Cheng, J (2002). Unearned premiums and deferred policy acquisition expenses in automobile extended warranty insurance. *Journal of Actuarial Practice*, **10**:63–95

Cheng, JS & Bruce, SJ (1993). A pricing model for new vehicle extended warranties. *Casualty Actuarial Society Forum*, *Special Edition*, 1–24

Dehghan, M & Duchesne, T (2011). A generalization of Turnbull's estimator for nonparametric estimation of the conditional survival function with interval-censored data. *Lifetime Data Analysis*, **17**(2):234–55

Dempster, AP, Laird, NM & Rubin, DB (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B (Methodological)*, **39**(1):1–38

Fay, MP & Shaw, PA (2010). Exact and asymptotic weighted logrank tests for interval-censored data: The interval R package. *Journal of Statistical Software*, **36**(2):1–34

Finkelstein, DM (1986). A proportional hazards model for interval-censored failure time data. *Biometrics*, **42**(4):845–54

Gentleman, R & Geyer, CJ (1994). Maximum likelihood for interval-censored data: Consistency and computation. *Biometrika*, **81**(3):618–23

Gentleman, R & Vandal, AC (2001). Computational algorithms for censored-data problems using intersection graphs. *Journal of Computational and Graphical Statistics*, **10**(3):403–21

Gómez, G, Calle, ML, Oller, R & Langohr, K (2009). Tutorial on methods for interval-censored data and their implementation in R. *Statistical Modelling*, **9**(4):259–97

Groeneboom, P & Wellner, JA (1992). *Information bounds and non-parametric maximum likelihood estimation*. Birkhäuser. Basel

Groeneboom, P, Jongbloed, G & Wellner, JA (2008). The support reduction algorithm for computing non-parametric function estimates in mixture models. *Scandinavian Journal of Statistics,* **35**(3):385–99

Hayne, RM (2007). Extended service contracts, an overview. *Variance*, **1**(1):8–39

Hsu C-H, Taylor JMG, Murray S & Commenges, D (2007). Multiple imputation for interval-censored data with auxiliary variables. *Statistics in Medicine*, **26**(4):769–81

Jack, N, Iskandar, BP & Murthy, DNP (2009). A repair–replace strategy based on usage rate for items sold with a two-dimensional warranty. *Reliability Engineering and System Safety*, **94**(2):611–17

Jack, N & Murthy, DNP (2007). A flexible extended warranty and related optimal strategies. *The Journal of the Operational Research Society*, **58**(12):1612–20

Jung, M & Bai, DS (2007). Analysis of field data under two-dimensional warranty. *Reliability Engineering and System Safety*, **92**(2):135–43

Kalbfleisch, JD, Lawless, JF & Robinson, JA (1991). Methods for the analysis and prediction of warranty claims. *Technometrics*, **33**(3):273–85

Karim, MR & Suzuki, K (2005). Analysis of warranty claim data: a literature review. *International Journal of Quality and Reliability Management*, **22**(7):667–86

Kerper, J & Bowron, L (2007). An exposure based approach to automobile warranty ratemaking and reserving, *Casualty Actuarial Society Forum*, *Spring*, 29–43

Law, CG & Brookmeyer, R (1992). Effects of mid-point imputation on the analysis of doubly censored data. *Statistics in Medicine*, **11**(12):1569–78

Lawless, JF (1998). Statistical analysis of product warranty data. *International Statistical Review*, **66**(1):41–60

Li, K, Mallik, S & Chhajed, D (2012). Design of extended warranties in supply chains under additive demand. *Production and Operations Management*, **21**(4):730–46

Maathuis, MH & Hudgens, MG (2011). Nonparametric inference for competing risks current status data with continuous, discrete or grouped observation times. *Biometrika*, **98**(2):325–40

Majeske, KD (2007). A non-homogeneous Poisson process predictive model for automobile warranty claims. *Reliability Engineering and System Safety*, **92**(2):243–51

Murthy, DNP (1992). A usage dependent model for warranty costing. *European Journal of Operational Research*, **57**(1):89–99

Murthy, DNP & Djamaludin, I (2002). New product warranty: A literature review. *International Journal of Production Economics*, **79**(3):231–60

Musakwa, FT (2012). Valuing the unexpired risk of motor service plans. Conference proceedings of the Actuarial Society of South Africa 2012 Convention, 16–17 October 2012, Cape Town International Convention Centre, Cape Town, 147–66

Rai, B & Singh, N (2004). Modeling and analysis of automobile warranty data in the presence of bias due to customer-rush near warranty expiration limit. *Reliability Engineering and System Safety*, **86**(1):83–94

Rai, B & Singh, N (2005). A modeling framework for assessing the impact of new time/mileage warranty limits on the number and cost of automotive warranty claims. *Reliability Engineering and System Safety*, **88**(2):157–69

Shafiee, M, Chukova, S, Saidia-Mehrabad, M & Niaki, STA (2011). Two-dimensional warranty cost analysis for second-hand products. *Communications in Statistics – Theory and Methods*, **40**(4):684–701

Shahanaghi, K, Noorossana, R, Jalali-Naini, SG & Heydari, M (2013). Failure modeling and optimizing preventive maintenance strategy during two-dimensional extended warranty contracts. *Engineering Failure Analysis*, 28:90–102

Su, C & Shen, J (2012). Analysis of extended warranty policies with different repair options. *Engineering Failure Analysis*, 25:49–62

Sun, J (1996). A non-parametric test for interval-censored failure time data with application to AIDS studies. *Statistics in Medicine*, **15**(13):1378–95

Sun, J (2001). Variance estimation of a survival function for interval-censored survival data. *Statistics in Medicine*, **20**(8):1249–57

Turnbull, BW (1976). The empirical distribution function with arbitrarily grouped, censored and truncated data. *Journal of the Royal Statistical Society, Series B (Methodological)*, **38**(3):290–95

Wellner, JA & Zhan, Y (1997). A hybrid algorithm for computation of the nonparametric maximum likelihood estimator from censored data. *Journal of the American Statistical Association*, **92**(439):945–59

Weltmann, LN & Muhonen, D (2002). Extended warranty ratemaking. *Casualty Actuarial Society*, *Winter*, 187–216

Wu, S (2012). Warranty data analysis: A review. *Quality and Reliability Engineering International*, **28**(8):795–805

Wu, S & Akbarov, A (2011). Support vector regression for warranty claim forecasting. *European Journal of Operational Research*, **213**(1):196–204

Zhang, Z & Sun, J (2010). Interval censoring. *Statistical Methods in Medical Research*, **19**(1):53–70