

On optimality criteria for age replacement

P. Coolen-Schrijner and F.P.A. Coolen

Department of Mathematical Sciences
University of Durham, Durham, DH1 3LE, UK

Abstract

Age replacement is a well-known topic in the literature of Operational Research and Reliability. Traditionally, the probability distribution of a unit's failure time is assumed to be known, and the cost criterion is derived via the renewal reward theorem, which implicitly assumes that the same preventive replacement strategy will be used over a very long period of time. As an alternative, one can use a one-cycle criterion, aiming at minimisation of costs per unit of time only over the period that one unit is in place. We discuss these two criteria, and we also consider possible alternatives. Recently, we have presented a nonparametric predictive approach to age replacement, which is based on rather minimal assumptions for the failure time distributions, and provides full flexibility to the information from the process. We summarize the main conclusions from this research, where we also considered both the renewal criterion and the one-cycle criterion. We discuss further aspects related to age replacement, highlighting several interesting topics for future research.

Key words: Age replacement, nonparametric predictive inference, optimality criteria

1. Introduction

Age replacement strategies for technical units describe that a unit is replaced preventively upon reaching age T , or correctively upon failure before T , where preventive replacements are typically less expensive than corrective replacements. We call this 'strategy T '. The classical mathematical approach for determining the optimal replacement age is based on the renewal criterion, which implicitly assumes that the same replacement strategy is used over a very long period of time, consisting of many cycles, where one cycle is the period of random length between two consecutive replacements. While this is a reasonable approach from classical Operational Research (OR) perspective, where one typically assumes detailed knowledge of the units' failure time distributions, it is less attractive in cases where one uses the observations from such a process to update information on such distributions, and accordingly may want to adapt the replacement age per cycle. Recently, we have introduced and studied nonparametric predictive inference (NPI) for such failure time distributions, allowing

replacement strategies to adapt fully to available failure and replacement data.

In this paper, we present an overview of recent results. We discuss the sometimes counter-intuitive ways in which optimal replacement strategies can adapt to data from the process, when the renewal criterion is used. In addition, we discuss the use of a one-cycle criterion to determine the optimal replacement strategy, both related to the classical OR framework with assumed failure time distribution and to our NPI approach.

In Section 2 we summarize classical age replacement theory, using the renewal criterion. In Section 3 we consider the one-cycle criterion, still from the more established OR perspective with a known failure time distribution. We compare these two criteria for Weibull failure time distributions in Section 4. In Section 5 we consider the use of a multiple cycles criterion. In Section 6 we report on main outcomes of our recent research project on adaptive age replacement within the statistical framework of nonparametric predictive inference (NPI), of which early results were presented in [1]. This includes comparison of the renewal criterion and the one-cycle criterion from NPI perspective, and we summarize insights gained from simulations studies, including some possibly surprising ways in which optimal strategies can adapt to information from the process. Finally, in Section 7 we briefly discuss several further issues, mainly concerning practical issues and related challenges for research.

2. Renewal criterion

In the classical model for age replacement [2, 3], the failure time of the unit is assumed to be an absolutely continuously distributed random quantity $X \geq 0$ with *known* probability distribution, with cumulative distribution function $F(x) = P(X \leq x)$, probability density function (pdf) $f(x)$, hazard rate $h(x) = f(x)/(1 - F(x))$, and expected value $E(X)$. We assume that $h(x)$ is monotonously strictly increasing, which is often considered to be a natural assumption for situations where age replacement may be cost effective [2, 3]. We use $h(\infty)$ to denote the limiting value (which may be ∞) of $h(x)$ for $x \rightarrow \infty$. To avoid mathematical complications, we assume that $F(0) = 0$, $F(x) > 0$ for all $x > 0$, and $E(X) < \infty$. The costs included in the age replacement model are assumed to be known constants, with c_1 the costs of preventive replacement, and c_2 the costs of corrective replacement, where $c_2 > c_1 > 0$.

The renewal criterion minimises the expected costs per unit of time, where the same replacement strategy is assumed to be used over an infinite period of time, applied to a sequence of units whose failure time random quantities are independent and identically distributed. The renewal reward theorem [3] implies that the cost function, for strategy $T > 0$, equals the expected costs per cycle divided by the expected length of a cycle,

$$C_r(T) = \frac{c_1(1 - F(T)) + c_2F(T)}{\int_0^T (1 - F(x))dx}. \quad (1)$$

If there is a finite optimal strategy T_r corresponding to the renewal criterion, then it is the unique value which satisfies

$$h(T) \int_0^T (1 - F(x))dx - F(T) = \frac{c_1}{c_2 - c_1}. \quad (2)$$

Such a finite T_r exists if $h(\infty) > \frac{c_2}{(c_2 - c_1)E(X)}$, else it is better not to replace the unit preventively. In this latter case, the cost function obtains its minimal value in the limit,

$$C_r(\infty) = \frac{c_2}{E(X)}. \quad (3)$$

3. One-cycle criterion

The one-cycle criterion [4, 5, 6, 7] minimises the expected costs per unit of time for a single cycle. This criterion is more natural than the renewal criterion in situations where one may wish to change the strategy per cycle, for example to take new information into account. Mazzuchi and Soyer [6] proposed this one-cycle criterion for Bayesian adaptive age replacement strategies. The cost function, as function of the random failure time X , for replacement strategy T , is the costs per unit of time during one cycle,

$$C_1(X, T) = \begin{cases} c_2/X & \text{if } X < T, \\ c_1/T & \text{if } X \geq T. \end{cases} \quad (4)$$

To avoid mathematical complexity, we assume that $E(1/X)$ exists, which is a condition on the failure time distribution for X for values close to 0 (which particularly excludes the use of the Exponential distribution close to 0 in what follows). The one-cycle criterion is minimisation of the expected value of $C_1(X, T)$ with regard to the probability distribution for X ,

$$C_1(T) = E(C_1(X, T)) = c_2 \int_0^T \frac{1}{x} f(x) dx + \frac{c_1}{T}(1 - F(T)). \quad (5)$$

The optimal strategy T_1 , corresponding to this criterion, might be infinite. If we assume a monotonously strictly increasing hazard rate, then a finite optimum strategy T_1 exists, and it is the unique solution to

$$Th(T) = \frac{c_1}{c_2 - c_1}. \quad (6)$$

The limiting value of this cost function, for $T \rightarrow \infty$, is

$$C_1(\infty) = c_2 E(1/X). \quad (7)$$

For such age replacement situations, where the unit's failure time has a monotonously strictly increasing hazard rate and where $E(X)$ and $E(1/X)$ are finite, it is easy to show [4, 8] that

$$T_1 < T_r. \quad (8)$$

Of course, T_r can be infinite as discussed in Section 2. Hence, in the classical stochastic setting with known probability distribution for a unit's failure time, if the unit is subject to wearout in the sense of a strictly increasing hazard rate for its failure time, then the optimal age replacement strategy according to the one-cycle criterion leads to earlier preventive replacement than the optimal strategy according to the renewal criterion. We illustrate this for the Weibull distribution in Section 4.

4. Comparison for Weibull distributions

In this section we illustrate the different optimality criteria, discussed in Sections 2 and 3, for Weibull distributions. In addition, we briefly consider estimating parameters of the Weibull distribution using failure times, and the effect on optimal replacement times.

The Weibull distribution, $W(\alpha, \beta)$, with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$, has pdf

$$f(x) = \frac{\alpha}{\beta} \left(\frac{x}{\beta}\right)^{\alpha-1} \exp\left\{-\left(\frac{x}{\beta}\right)^\alpha\right\}, \quad (9)$$

for $x \geq 0$. Its hazard rate is

$$h(x) = \frac{\alpha}{\beta} \left(\frac{x}{\beta}\right)^{\alpha-1}, \quad (10)$$

and its expected value and variance are

$$E(X) = \beta \Gamma\left(1 + \frac{1}{\alpha}\right), \quad (11)$$

$$Var(X) = \beta^2 \left[\Gamma\left(1 + \frac{2}{\alpha}\right) - \left\{ \Gamma\left(1 + \frac{1}{\alpha}\right) \right\}^2 \right]. \quad (12)$$

The expected value and variance are both increasing in β .

We restrict our discussion here to increasing hazard rates, so shape parameter $\alpha > 1$. For age replacement with the one-cycle criterion, the optimal strategy is

$$T_1 = \beta \left(\frac{c_1}{\alpha(c_2 - c_1)} \right)^{1/\alpha}. \quad (13)$$

Clearly, T_1 is increasing in β . If we would use a Weibull distribution with the value of α assumed to be known, but estimating the value of β from failure data, then most commonly used estimators of β , e.g. the moment estimator or the maximum likelihood estimator, will be increasing in the sample mean, hence the T_1 corresponding to such estimates would increase in the sample mean. It is interesting to remark that, for this situation with α assumed to be known, this T_1 is a fixed quantile of this Weibull distribution, as

$$F(T_1) = 1 - \exp \left\{ -\frac{c_1}{\alpha(c_2 - c_1)} \right\} \quad (14)$$

does not depend on β .

For the renewal reward criterion, the optimal strategy T_r is not available analytically for the Weibull distribution, but must be computed numerically via equation (2). Next, we illustrate the differences between these two criteria for the Weibull distribution with scale parameter $\beta = 1$, and we set $c_1 = 1$ and $c_2 = 10$. Table 1 gives the values of T_1 and T_r for several $W(\alpha, 1)$ distributions. Figure 1 shows the cost functions, for both the renewal and one-cycle criteria, for $W(1.5, 1)$ and $W(3, 1)$.

α	T_1	T_r
1.2	0.1377	0.6861
1.5	0.1764	0.3781
2	0.2357	0.3365
3	0.3333	0.3825
5	0.4670	0.4886
10	0.6376	0.6445

Table 1: Optimal replacement times for $W(\alpha, 1)$

It is easy to show that T_1 is increasing in α , which is illustrated in Table 1. The hazard rate is monotonously strictly increasing if the shape parameter α is greater than one.

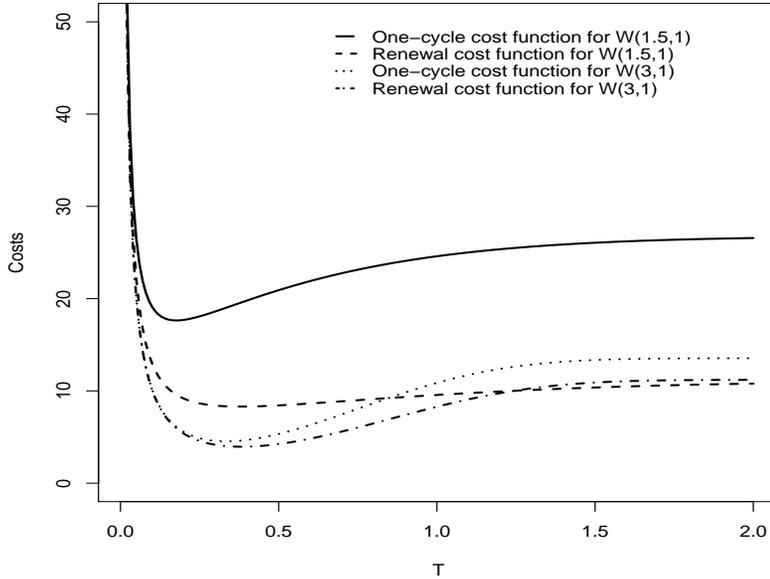


Figure 1: One-cycle and Renewal cost functions for $W(1.5, 1)$ and $W(3, 1)$

We see that T_1 is always less than T_r , confirming (8). The optimal replacement time according to the renewal criterion is first decreasing and thereafter increasing in α , which is in agreement with the fact that the expected value of a $W(\alpha, 1)$ distributed random variable is decreasing in α , for $\alpha \in (0, 2.1645)$ and increasing in α for $\alpha > 2.1645$ (see e.g. [9], p.259). The difference between T_1 and T_r decreases in α , for $\alpha \geq 1$, which can be shown analytically by the fact that, for such values of α , the derivative with respect to α of the left hand side of (6) minus the left hand side of (2) is negative for T in the range containing these values of T_1 and T_r . We also see that the difference between the one-cycle cost function and the renewal cost function decreases in α for $T \rightarrow \infty$. Indeed, by relatively easy analysis, one can prove that the limits, for $T \rightarrow \infty$, of these two cost functions, converge to each other for increasing α .

5. Optimality over multiple cycles

The one-cycle criterion and the renewal criterion, which effectively considered an infinite number of cycles, can be regarded as extreme cases of an optimality criterion over $k \geq 1$ cycles, aiming at minimal costs per unit of time. Let L_i be the length of cycle i , for $i = 1, \dots, k$, and let N_c be the random number of cycles, out of these k , that end with corrective replacement. Then the cost function according to the k -cycle criterion

is, for $T > 0$,

$$C_k(T) = \sum_{l=0}^k E \left(\frac{1}{\sum_{i=1}^k L_i} \mid N_c = l \right) [lc_2 + (k-l)c_1] \binom{k}{l} \{F(T)\}^l \{1-F(T)\}^{k-l}. \quad (15)$$

To calculate this cost function for most failure time distributions, this expected value requires nested integrals to be computed numerically. For example, for $k = 2$ we get

$$C_2(T) = 2c_2 \int_0^T \int_0^T \frac{1}{u+v} f(u)f(v) dudv + 2(c_1 + c_2) \{1-F(T)\} \int_0^T \frac{1}{T+u} f(u) du + \frac{c_1}{T} \{1-F(T)\}^2. \quad (16)$$

If no preventive replacements will be carried out, so using strategy $T = \infty$, then all cycles end with corrective replacement, and the random length L_i of cycle i is identical to the random failure time of the i -th unit, X_i say. Then the limiting costs are equal to the term in (15) with $l = k$, so

$$C_k(\infty) = kc_2 E \left(\frac{1}{\sum_{i=1}^k X_i} \right). \quad (17)$$

In the Appendix we prove that, if the failure times X_i are non-trivial, independent and identically distributed, then $C_k(\infty)$ is decreasing in k . We strongly believe that $C_k(T)$ also decreases in k , at any T , but we have not yet been able to prove this. We have checked several examples numerically, see e.g. Figure 2 where we plot $C_k(T)$ for $k = 1, 2, 3$ and $k = \infty$, which is equal to $C_r(T)$, for the $W(2, 1)$ distribution. We have not found any example where these cost functions are not decreasing in k for all T .

In addition to the result that $T_1 < T_r$ [8], as discussed in Section 3, which holds in case of increasing hazard rates and assuming that $E(X)$ and $E(1/X)$ are finite, we also conjecture that, in such cases, the optimal strategy T_k corresponding to the k -cycles criterion, increases in k . If this were true, than T_1 and T_r would give the two extreme optimal age replacement strategies over all k -cycles criteria, which would in particular be useful in situations where T_1 and T_r are close to each other, as that would indicate that the precise choice of the number of cycles in the optimality criterion had little relevance in such cases. Table 2 gives the optimal strategy T_k corresponding to the k -cycles criterion in case the failure times follow a $W(2, 1)$ distribution. The largest difference is between T_1 and T_2 , for larger k the value of T_k converges quickly to $T_\infty = T_r$. The integrals in the calculation of T_k were calculated with the package 'ADAPT' in the statistical software R ¹. It uses Genz's Fortran ADAPT subroutine

¹<http://www.cran.r-project.org/>

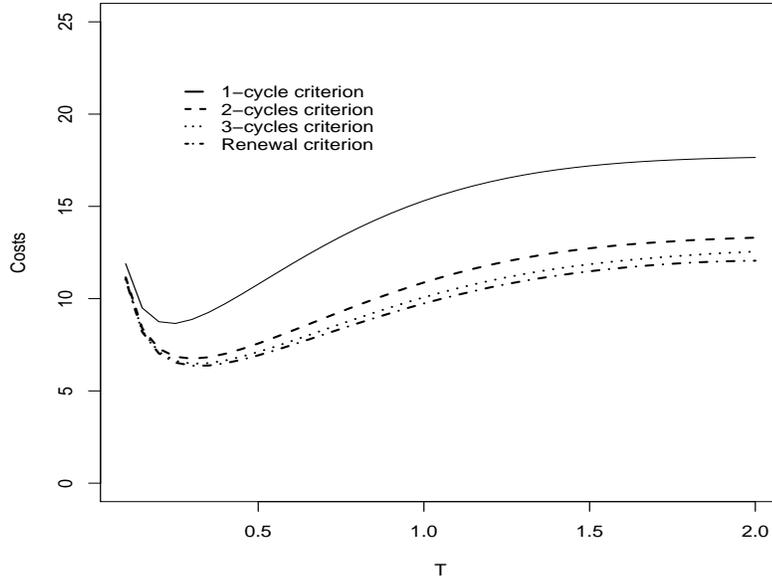


Figure 2: Cost functions for $W(2,1)$ based on k -cycles ($k=1,2,3$) and renewal criterion

to do all the calculations. The Fortran functions have been modified to use double precision. We used these to illustrate that T_k 's are increasing in k and converge to T_∞ , where T_∞ can be calculated numerically via equation (2).

k	T_k
1	0.2357
2	0.2998
3	0.3139
4	0.3200
5	0.3235
6	0.3258
7	0.3274
8	0.3285
∞	0.3365

Table 2: Optimal replacement times T_k for $W(2,1)$

6. NPI-based age replacement

In a recent research project, we have developed and presented methods, based on non-

parametric predictive inference (NPI) [10, 11], for age replacement, using the renewal criterion as well as the one-cycle criterion. In NPI the starting point is not a known failure time distribution, but the wish to make rather minimal assumptions about such a distribution, and hence to work in a nonparametric statistical framework. The information on which to base inferences then comes from assumed observations. Hence, in this work [1, 8, 12], we assumed that n observed failure times were available. NPI is based on Hill's assumption $A_{(n)}$ [13], which in this context states that the random failure time of unit $n + 1$ has equal probability to fall in each of the $n + 1$ intervals on the positive real line as created by the first n observed failure times (assuming, for simplicity, no tied failure times were observed). No further assumptions are made on the probabilities for the failure time of unit $n + 1$ within each of these intervals, which implies that for many inferences we work with bounds for probabilities that are consistent with $A_{(n)}$ [10]. In our replacement settings, we use lower and upper survival functions [11, 14] for the failure time of unit $n + 1$, and these result in lower and upper cost functions, the difference between these reflects the amount of information available (hence decreases with n). The main aim of this research project was to study whether or not such an approach was achievable, where the crucial difference with earlier work is that no assumptions on the hazard rate of the underlying failure time distribution are made, hence providing a very flexible framework.

Our first results were reported in [1], where we presented analytical forms for the lower and upper cost functions for NPI-based age replacement of unit $n + 1$, with the renewal criterion, and we proved that the optimal preventive replacement times corresponding to these two cost functions are in (or 'just before') an observed failure time, if the data suggest that preventive replacement is cost effective. We also showed how to determine whether the data suggest that it is better not to replace preventively. This result is of great benefit for computation of the optimal replacement times. We showed consistency of our NPI-based lower and upper cost functions via comparison with the related theoretical cost function based on an assumed underlying failure time distribution, in the following sense. If the n observations from such an underlying failure time distribution were the equally spaced percentiles for that distribution, whereby the data would give a 'perfect reflection' of the failure time distribution, then the theoretical cost function is everywhere between our lower and upper cost functions. To study the performance of our method, we simulated failure times from different Weibull distributions with shape parameter greater than 1 (so modelling wear-out), and applied our NPI-based methods. This allowed comparison of our resulting preventive replacement strategies

and the theoretical strategies, corresponding to the known failure time distribution. Mostly, our method adapted very well to the data, in the sense that our NPI-based optimal replacement times were close to the theoretical optimal replacement times, already for only $n = 10$ observed failure times. Of course, for larger n the results were increasingly good. Situations where our method suggested quite different replacement times than the theoretical method occurred when the simulated data did not reflect wear-out, which happened occasionally for $n = 10$ but very rarely for larger n . The results reported in [1] show that NPI-based methods for age replacement are useful alternatives to the established approaches, which already work quite well for relatively few observed failure times.

We continued this research by focussing on the ability of our NPI-based methods for age replacement to adapt to further failure time information from the process. So, with n failure times available, we assumed that our optimal strategy is used for unit $n + 1$, and the information about this unit is then taken into account to determine the optimal replacement strategy for unit $n + 2$. The information about unit $n + 1$ consists of either an observed failure, occurring before the optimal preventive replacement time for this unit, or a right-censored observation in case of preventive replacement. The results of this research are presented in [12]. First, we showed how the right-censored observation, in case of preventive replacement of unit $n + 1$, could be taken into account in the lower and upper NPI-based cost functions for unit $n + 2$, using NPI theory for dealing with right-censored data [15]. We derived analytical expressions for these cost functions, and proved that, again, optimal replacement times would coincide with previous observations. We proved theorems on the way in which the optimal replacement times may (not) vary in case of a series of preventively replaced units, which give more insight in the effect of such additional process data. Again, we performed extensive simulation studies into the adaptive behaviour of our NPI-based age replacement strategies, focussing this time on the differences between the optimal replacement times for different units in such a process. Although most conclusions from these simulations confirmed our intuitions, e.g. the difference between the optimal replacement times for units $n + 1$ and $n + 2$ tends to decrease with n , there were several results which, at first, may have been surprising. Detailed study of these results has provided insights into the use of the renewal criterion for age replacement, which had not been discussed before. For example, in some simulated situations, a failure of unit $n + 1$ before its optimal preventive replacement time, led to an increase of the optimal replacement time for unit $n + 2$. We discuss such situations in detail in [12]. Hence, not only have we

shown that our NPI-based approach is feasible for age replacement, but it also provides detailed insights into the optimality criterion used, and its effect if further information from the process is taken into account.

We also studied NPI-based adaptive methods for age replacement using the one-cycle criterion discussed in Section 3, which is attractive for fully adaptive replacement strategies. The results of this research are presented in [8]. For mathematical convenience, in this work we did add a further assumption on the probability distribution of the failure time of unit $n + 1$ within the intervals created by the n observed failure times [8]. We derived explicit expressions for the optimal preventive replacement times, which enable relatively straightforward computation of optimal strategies. We studied the performance of this method, again via simulation studies, both with regard to the way in which it adapts to failure time information, and in how the optimal strategy compares to the optimal replacement time for the renewal criterion. This latter comparison is not trivial, as the theoretical comparison in Section 3 was only proven for failure time distributions with increasing hazard rates [8], while in NPI such an assumption is not used. The main conclusions reported in [8] are that, again, our method adapts well to failure data (in a similar manner as for the renewal criterion). In all our simulations, we always found that the NPI-based optimal preventive replacement time according to the one-cycle criterion, did not exceed the corresponding optimal time according to the renewal criterion. This suggests that the result $T_1 < T_r$ holds quite generally, so not only for failure time distributions with monotonously strictly increasing hazard rates as considered in Section 3. Of course, we do not suggest that no data sets can be found for which this is not the case, but we have not managed to find such counter examples. We have not considered other optimality criteria for age replacement in our NPI-based approach, which is interesting for future research, but such work will mostly require numerical methods as analytical results will become very complex.

7. Concluding remarks

The choice of optimality criterion for age replacement will, in practice, depend on several issues in addition to the planning horizon. For example, if a clear budget is set for replacement activities, over a fixed period, then this may directly lead to earlier or later replacement. An advantage of our NPI-based method is its explicitly predictive nature, through the use of the next unit's random failure time. In particular in combination with the one-cycle criterion, it is therefore easy to check the expected risks

involved with delaying, or bringing forward, preventive replacement of the next unit. More generally, one may wish to use utility functions in a decision theoretic framework [16], to measure the benefits and risks, related to different replacement strategies, to the wider context of the company or society involved. Again, the NPI-based method provides the attractive feature of such utilities being directly in terms of possible events involving the next unit. There are several exciting research opportunities in the practicalities of budgets and utilities for such maintenance and replacement engineers [17]. For example, management may need to set the budget at an early stage, and their utility function is likely to be quite different to that of the maintenance engineer, whose input they may ask in order to set the budget. This could lead to a game-theoretic scenario, where the engineer aims to maximize the budget made available to him, whereas management may wish to minimize this, or to keep it about constant over consecutive planning periods. In the light of this, there may be strategic interests that could affect the way in which subjective beliefs are reported. This raises interesting questions with regard to communication of utilities and uncertainties. Some possible methods towards resolving such issues have been studied in so-called ‘principal agent theory’ [17, 18]. Such ‘distributed decision making’ is further complicated in preventive replacement settings by the almost paradoxical nature of preventive replacement: if one does this perfectly, which would imply that preventive replacement always takes place just before the unit would fail, then it may well give the impression of over-maintenance, because the units would never fail. Ideally, one would be able to find ways to resolve such conflicting situations, e.g. if one can predict how long a preventively replaced unit would still have been in functional order, supported by measurements on such a unit. Extending our NPI-based method to take such possible measurements into account would be an interesting topic for research, and would greatly enhance its practical value. Such research should, ideally, be directly linked to practical applications.

In situations where time between replacements is typically large (say ‘months’ or even ‘years’), one may wish to apply a discount factor [19] to take inflation into account. We have not included discount factors in our analysis. It would be interesting to study how discount factors affect the optimal NPI-based preventive replacement times, in particular of course for the renewal argument. However, if typical time periods considered are so long that such discounting becomes relevant, one should also focus on several other parts of the age replacement models which may vary over time, for example the costs for corrective and preventive replacements, which may not remain constant over a long period. Also, replacement units may have undergone design changes over a longer

period of time, which could affect the failure time distributions.

In practice, good decision making with regard to preventive replacements may well have to take a variety of criteria into account, which one may not be able to combine into one single criterion. This may, for example, involve economic, environmental and safety considerations. In the literature on multi-criteria decision making [20], methods that are suitable for such situations have been presented and analyzed, but as far as we are aware this has not yet been applied to replacement problems in a practically interesting and relevant way. Such methods would also be of interest in combination with NPI-based methods, as criteria could again be directly formulated in relation to the next unit to be used in the process. There are several interesting research topics related to this issue, such topics would benefit from theoretical research which is directly linked to practical applications.

If one aims at optimal preventive replacement strategies over many cycles, and one acknowledges uncertainty about the failure time distributions of the units, then it may be worthwhile to take into account the information one gets from the observation on the current unit or the next unit, about the failure time distributions of future units. Intuitively, it may perhaps be beneficial to delay a unit's preventive replacement. If it were to function without any problems for a longer period of time, this might indicate that, perhaps, one was too pessimistic about the failure time distribution, whereas if it fails soon after the initially determined optimal replacement time, this will lead to higher costs now, but the information may allow better planning in the future. Hence, it is also a matter of intertemporal decision making [16], so one should balance (e.g. via utilities) the possible risks and benefits at different times. This is also a challenging research area that has not yet received any serious attention with regard to replacement problems, as far as we are aware. Again, it seems that the adaptive predictive nature of our NPI-based methods are ideally suited for combination with such practically relevant issues.

Recently, we have studied a variation to the age replacement work which we briefly discussed in Section 6. We considered opportunity-based age replacement [21], where preventive replacement of units can only take place at randomly occurring opportunities. This is, for example, of practical importance in situations where a unit is part of a larger system, and replacement requires the system's functioning to be interrupted, whereas it may be relatively cheap to replace the unit during periods when the system

is not functioning. Theoretically, this work was also important, as we successfully combined NPI-based methods (on the failure time distributions) with classical stochastic processes (on the randomly occurring replacement opportunities), which indicated that such combined methods are possible, both in theory and computationally, and hence this may significantly widen the applicability of NPI-based methods. We will report on this research elsewhere [22], it also leads to interesting topics for further research.

There are several other interesting research topics related to our NPI-based methods for age replacement. For example, one could include small problems with the units in the process, which occur randomly and can be resolved by minimal repairs. This was studied by Sheu, *et al.* [7], in an adaptive Bayesian framework. We have not yet seen this in combination with opportunity-based age replacement, which is perhaps the more natural setting for such minimal repair activities. Development of NPI-based methods for other replacement problems, e.g. block replacement [3], is also of great interest, and would enhance the applicability of our methods.

Appendix

We prove that, if the failure times $X_i > 0$ are independent and identically distributed, then $C_k(\infty)$, as presented in equation (17) in Section 5, is decreasing in k (we assume here further that the X_i are non-trivial, i.e. they do not take on a single value with probability 1, and that all expected values here are finite).

Let $Y_l > 0$, for $l = 1, \dots, k$, with $k \geq 2$, be identically distributed real-valued random quantities (again, we assume the Y_l to be non-trivial, and all expected values to be finite). Let $g(\cdot)$ be a strictly convex positive function. Then

$$g\left(\frac{1}{k} \sum_{l=1}^k Y_l\right) < \frac{1}{k} \sum_{l=1}^k g(Y_l), \quad (18)$$

and the same relation holds for the expected values,

$$E \left\{ g\left(\frac{1}{k} \sum_{l=1}^k Y_l\right) \right\} < \frac{1}{k} \sum_{l=1}^k E\{g(Y_l)\}. \quad (19)$$

The random failure times X_i , for $i = 1, \dots, k$, as introduced above, have average value $\bar{X}_k = \frac{1}{k} \sum_{i=1}^k X_i$, and let

$$\bar{X}_k^l = \frac{1}{k-1} \sum_{i \neq l} X_i. \quad (20)$$

Then

$$\bar{X}_k = \frac{1}{k} \sum_{i=1}^k X_i = \frac{1}{k} \sum_{l=1}^k \bar{X}_k^l. \quad (21)$$

Let $g(x) = 1/x$, for $x > 0$, which is indeed a strictly convex positive function. We now use (19), with $Y_l = \bar{X}_k^l$, which gives

$$E \left\{ \frac{k}{\sum_{l=1}^k \bar{X}_k^l} \right\} < \frac{1}{k} \sum_{l=1}^k E \left\{ \frac{1}{\bar{X}_k^l} \right\} = E \left\{ \frac{1}{\bar{X}_k^k} \right\}, \quad (22)$$

where the equality follows from the fact that the X_i are identically distributed, and therefore so are the \bar{X}_k^l , hence they have the same expected value for all $l = 1, \dots, k$. Using (21) in the left-hand side of (22), gives

$$E \left\{ \frac{k}{\sum_{i=1}^k X_i} \right\} < E \left\{ \frac{k-1}{\sum_{i=1}^{k-1} X_i} \right\}, \quad (23)$$

which completes the proof.

Acknowledgements

This research was supported by the UK Engineering and Physical Research Sciences Council, grant GR/R92530/01. We thank our colleague, Iain MacPhee, for suggesting the idea for the proof presented in the Appendix. We are grateful to two referees for making us aware of relevant literature, and for suggestions to improve the presentation of this paper.

References

- [1] Coolen-Schrijner, P. and Coolen, F.P.A. (2004a). Nonparametric predictive inference for age replacement with a renewal argument. *Quality and Reliability Engineering International*, **20**, 203-215.
- [2] Aven, T. and Jensen, U. (1999). *Stochastic Models in Reliability*. Springer, New York.
- [3] Barlow, R.E. and Proschan, F. (1965). *Mathematical Theory of Reliability*. Wiley, New York.
- [4] Ansell, J., Bendell, A. and Humble, S. (1984). Age replacement under alternative cost criteria. *Management Science*, **30**, 358-367.

- [5] Derman, C. and Sacks, J. (1960). Replacement of periodically inspected equipment. *Naval Research Logistics Quarterly*, **7**, 597-607.
- [6] Mazzuchi, T.A. and Soyer, R. (1996). A Bayesian perspective on some replacement strategies. *Reliability Engineering and System Safety*, **51**, 295-303.
- [7] Sheu, S.H., Yeh, R.H., Lin, Y.B. and Juan, M.G. (1999). A Bayesian perspective on age replacement with minimal repair. *Reliability Engineering and System Safety*, **65**, 55-64.
- [8] Coolen-Schrijner, P. and Coolen, F.P.A. (2006). Nonparametric adaptive age replacement with a one-cycle criterion. *Reliability Engineering and System Safety*, to appear.
- [9] Abramowitz, M. and Stegun, I.A. (1970). *Handbook of mathematical functions*. Ninth ed. 1970. Dover Publications, Inc, New York.
- [10] Augustin, T. and Coolen, F.P.A. (2004). Nonparametric predictive inference and interval probability. *Journal of Statistical Planning and Inference*, **124**, 251-272.
- [11] Coolen, F.P.A., Coolen-Schrijner, P. and Yan, K.J. (2002). Nonparametric predictive inference in reliability. *Reliability Engineering and System Safety*, **78**, 185-193.
- [12] Coolen-Schrijner, P. and Coolen, F.P.A. (2004b). Adaptive age replacement based on nonparametric predictive inference. *Journal of the Operational Research Society*, **55**, 1281-1297.
- [13] Hill, B.M. (1968). Posterior distribution of percentiles: Bayes' theorem for sampling from a population. *Journal of the American Statistical Association*, **63**, 677-691.
- [14] Coolen, F.P.A. (2004). On the use of imprecise probabilities in reliability. *Quality and Reliability Engineering International*, **20**, 193-202.
- [15] Coolen, F.P.A. and Yan, K.J. (2004). Nonparametric predictive inference with right-censored data. *Journal of Statistical Planning and Inference*, **126**, 25-54.
- [16] French, S. and Rios Insua, D. (2000). *Statistical Decision Theory*. Arnold, London.
- [17] Baker, R. (2004). Over-maintenance and the principal-agent problem: a carrot for the engineer. *Proceedings of the 5th IMA International Conference on Modelling*

in Industrial Maintenance and Reliability, Wang, Scarf and Newby (eds), The Institute of Mathematics and its Applications, 33-37.

- [18] Schneeweiss, C. (2003). *Distributed Decision Making (2nd ed.)*. Springer, Berlin.
- [19] Barlow, R.E. (1998). *Engineering Reliability*. SIAM, Philadelphia.
- [20] Belton, V. and Stewart, T.J. (2002). *Multiple Criteria Decision Analysis*. Kluwer Academic Publishers, Boston.
- [21] Dekker, R. and Dijkstra, M.C. (1992). Opportunity-based age replacement: Exponentially distributed times between opportunities. *Naval Research Logistics*, **39**, 175-190.
- [22] Coolen-Schrijner, P., Coolen, F.P.A. and Shaw, S.C. (2006). Nonparametric adaptive opportunity-based age replacement strategies. *Journal of the Operational Research Society*, to appear.