Proceedings of Machine Learning Research 215:222–229, 2023 ISIPTA 2023

A Comparison Between a Frequentist, Bayesian and Imprecise Bayesian Approach to Delay Time Maintenance

Marc Fischer MARC.FISCHER@EMSE.FR
Mines Saint-Etienne, Univ Lyon, CNRS, UMR 5307 LGF, Centre SPIN, F - 42023 Saint-Etienne France

Abstract

Delay time models are stochastic maintenance decision aid tools that divide the failure time of a system into the appearance of a defect and its evolution towards a breakdown. In this study, an imprecise Bayesian approach to delay time modelling has been developed and compared with the frequentist and precise Bayesian approach based on a virtual maintenance problem. The conditional failure rate was the unknown parameter that had to be estimated via eight samples of failure times of increasing size. The goal was to minimise two loss functions related to the downtime and cost. The frequentist and precise Bayesian methods converge towards the optimal decision as the sample size grows but are strongly sub-optimal when no or only few data are available. The imprecise Bayesian approach based on E-admissibility returns large decision intervals in the lack of data, thereby straightforwardly representing the crucial difference between knowledge and ignorance.

Keywords: stochastic maintenance, robust Bayesianism, decision rules, frequentism

1. Introduction

In the stochastic modelling of maintenance, delay-time analysis (DTA) explicitly makes a difference between the appearance of a defect (at a time U) and the system breakdown caused by this defect at time $U + H$, where H is the delay-time Christer (1999); Wang (2012a); Scarf et al. (2019). It can be seen as being a kind of periodic off-line condition monitoring. Christer and Waller (1984) developed the first detailed delay-time model back in 1984 in order to optimise either the downtime per unit time $D(T)$ or the cost per unit time $C(T)$ as an enterprise can be primarily interested in either of these two quantities. Delay time models are conceptually simple (thereby reducing the risk of over-fitting) and they can successfully be applied to a wide range of industrial problems Wang (2012a). In this work, we shall consider the simplest situation, namely that of a system with a single failure mode where a single defect appears at a random time and progressively evolves into a full-blown failure after a random and stochastically independent delay time. Such a model can be used to simulate multi-component systems where the failure of one component has no impact on the proper function of the other components and it was applied to the maintenance of three infusion-pump components Baker and Wang (1993) and to a sample of about 100 infusion pumps Baker and Wang (1991). DTA generally assumes that the time of appearance of a defect U and the delay time between the arrival of that defect and the breakdown of the system H follow two exponential or Weibull distributions. The parameters of these distributions are most often estimated via Maximum Likelihood Estimation (MLE) Baker and Wang (1991) but sometimes also through a classical Bayesian approach Wang and Jia (2007); Wang (2012b). There are not, however, yet any studies that represent our prior ignorance through a family of priors rather than through a single one, which is problematic as the reliance on a single prior can all too easily lead one to mistake ignorance for specific knowledge Norton (2010); Fischer (2019); Fischer and Vignes (2021).

The present work thus sets out to compare a frequentist, a classical Bayesian and an imprecise Bayesian approach to the problem of DTA parameter estimation and inspection optimisation based on samples of failure times virtually generated via Monte-Carlo simulations. In Section 2, the stochastic model is laid out along with the samples that will be used for the parameter estimation. The frequentist, classical Bayesian and imprecise Bayesian approach are considered in Section 3, 4 and 5, respectively. The conclusions and an outlook are given in Section 6.

2. Delay Time Modelling and Problem Setting

2.1. Stochastic Decision Model

We consider a system that can at most have only one defect at the same time before a failure. We consider that both the appearance time of the defect U and the delay time before the failure H are exponentially distributed and we shall call their parameters k_f (the appearance rate of defects in defects/month) and λ_h (the conditional failure rate in failures/defect/month), respectively. It can be easily shown that the failure time $Y = U + H$ has a probability density function given by

$$f_Y(y) = k_f \lambda_h e^{-k_f y} \int_{h=0}^{y} e^{(k_f - \lambda_h)h} dh$$

$$= \frac{k_f \lambda_h e^{-k_f y}}{k_f - \lambda_h}\left(e^{(k_f-\lambda_h)y} - 1\right) \tag{1}$$

when $\lambda_h \neq k_f$.

Let us now suppose that an inspection is carried out every T units of *operating time*, i.e. time during which the machine is operational and isn't being repaired or inspected. $T = +\infty$ means that there is no inspection. Let d_{br}, d_i, d_{defect} be the downtime due to a breakdown, an inspection and an inspection repair (when a defect is detected), respectively. It can be interesting for an enterprise to minimise the ratio between the downtime and the operating time $D_1(T)$ given by $D_1(T) = \frac{E(Downtime)}{E(Operating\ Time)}$ where $Operating\ Time$ is the length of a renewal cycle and $Downtime$ is the downtime associated to that renewal cycle. Indeed, according to the renewal-reward theorem Wang (2008), that quantity is the long-term ratio between the cumulated downtime and the cumulated operating time given that the inspection period is equal to T.

Since $Operating\ Time = min(Y, T)$, by virtue of the law of total expectation, we have

$$\begin{aligned} E(Operating\ Time) =& E(Y|Y \leq T)F_Y(T) + \\ & T\left(1 - F_Y(T)\right) \end{aligned}$$

with F_Y being the cumulative distribution function (cdf) of the failure time Y.

Let X_{intact}, X_{defect} and $X_{failure}$ be binary random variables equal to 1 when there is no defect during the renewal cycle, one defect that does not lead to a failure and a failure, respectively. We have

$$\begin{aligned} Downtime =& X_{intact} d_i \\ & + X_{defect}(d_i + d_{defect}) + X_{failure} d_{br}, \end{aligned} \tag{2}$$

so that

$$\begin{aligned} E(Downtime) =& p_{intact} d_i + \\ & p_{defect}(d_i + d_{defect}) \\ & + p_{failure} d_{br} \end{aligned} \tag{3}$$

Furthermore,

$$p_{intact} = 1 - F_U(T) \tag{4}$$

$$p_{defect} = \left(1 - p(Y \leq T|U \leq T)\right)F_U(T) \tag{5}$$

and

$$p_{failure} = F_Y(T). \tag{6}$$

Likewise, let c_{br}, c_i, c_{defect} be the cost due to a breakdown, an inspection and an inspection repair (when a defect is detected), respectively. It is interesting to minimise

$$C_1(T) = \frac{E(Cost)}{E(Operating\ Time)} \tag{7}$$

with

$$\begin{aligned} E(Cost) =& p_{intact} c_i + \\ & p_{defect}(c_i + c_{defect}) + \\ & p_{failure} c_{br} \end{aligned} \tag{8}$$

An enterprise can be interested in minimising the cost by unit of operating time $C_1(T)$ but also the downtime by unit of operating time $D_1(T)$ (when too much downtime could lead to the loss of market shares, for example). Often times, a compromise between these two goals has to be made.

Let $T_{D_1,opt,S}$ be the optimal inspection period minimising $D_1(T)$ based on the information contained in a sample S of failure times (when $S = \emptyset$, this would be a decision made in the absence of information other than the parameter bounds). $T_{D_1,opt,\lambda_{h,true}}$ is the real optimal inspection period based on the real value of λ_h. Let us consider that the enterprise will use this type of pumps in 6 factories for a period equal to $t = 4$ years of operational time (i.e. not including all the potential downtimes) [1]. It is then interesting to compute the quantity

$$\begin{aligned} &\Delta Downtime = \\ &6\left(D_1(T_{D_1,opt,S}, \lambda_{h,true}) - D_1(T_{D_1,opt,\lambda_{h,true}}, \lambda_{h,true})\right)t \end{aligned} \tag{9}$$

which represents the *sub-optimality* of the decision made on the basis of sample S, i.e. the amount of downtime that would have been avoided if we had known the real parameter value perfectly. $D_1(T_{D_1,opt,S}, \lambda_{h,true})$ is the expected downtime based on the estimated optimal inspection period obtained thanks to sample S whereas $D_1(T_{D_1,opt,\lambda_{h,true}}, \lambda_{h,true})$ is the expected downtime based on the true optimal inspection period determined with the true value of λ_h.

Likewise, we can define for the cost

$$\begin{aligned} &\Delta Cost = \\ &6\left(C_1(T_{C_1,opt,S}, \lambda_{h,true}) - C_1(T_{C_1,opt,\lambda_{h,true}}, \lambda_{h,true})\right)t \end{aligned} \tag{10}$$

which is the sub-optimality of the decision made on the basis of sample S with respect to $C_1(T)$.

[1] Of course, another number of factories can be considered according to the industrial situation.

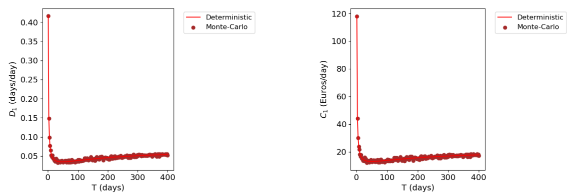

(*a*) D1 (*b*) C1

Figure 1: Profile of $D_1(T)$ and $C_1(T)$ for the real parameter value $\lambda_{h,true}$ = 0.09 failures/defect/month

.

2.2. Virtual Case Study

We consider a pump that can have at most one defect which progressively deteriorates into a failure. The enterprise knows the defect arrival rate k_f = 1/4 defects/month. The conditional failure rate λ_h is equal to 9/100 failures/defect/month, however the enterprise only knows that $\lambda_h \in [0.01; 0.1]$ failures/defect/month.

We suppose that the cost and downtime parameters are $d_i = 0.4$ days, $d_{defect} = 2$ days, $d_{br} = 28$ days, $c_i = 110$ €, $c_{defect} = 1000$ €, and $c_{br} = 9000$ €.

$D_1(T)$ and $C_1(T)$ can be seen in Figure 1.

It is worth noting that the results of the analytical model are very close to those from a Monte-Carlo simulation. The optimal inspection period for minimising D_1 is $T_{D_1,opt,S}$ = 47.75 days and that for minimising C_1 is $T_{C_1,opt,S}$ = 49.93 days and it is these values that the probabilistic approaches ought to approximately retrieve.

1000 failure times were then generated for k_f = 1/4 defects/month and $\lambda_h = \lambda_{h,true}$ = 9/100 defects/failure/month. Eight samples of failure times *measured in the absence of inspections* $S_1 \subset S_2 \subset S_3 \subset S_4 \subset S_5 \subset S_6 \subset S_7 \subset S_8$ whose lengths are equal to 2, 5, 10, 20, 60, 200, 600, and 1000, respectively, were then created [2]. In this way, we can simulate the situation of a maintenance engineer who must make an optimal decision in the face of incomplete information (or no information at all).

3. Frequentist Approach

In a frequentist framework, given a sample of failure times S, the parameters (here λ_h) are typically estimated by maximising the likelihood function which is given by Equation 11.

$$L(\lambda_h|y_1, y_2, \ldots, y_n) = \prod_{i=1}^{n} f_Y(y_i)$$

$$= \prod_{i=1}^{n} \frac{k_{f,0}\lambda_h}{k_{f,0} - \lambda_h}\left(e^{-\lambda_h y_i} - e^{-k_{f,0} y_i}\right) \tag{11}$$

There does not appear to be an analytical expression for the MLE. Instead, the likelihood function must be numerically maximised. Generally, the maximum likelihood estimate of the parameter(s) is used to minimise the loss functions C_1 and D_1 Baker and Wang (1991); Christer et al. (1995). Accordingly, the MLE corresponding to the different samples were used to compute $T_{D_1,opt,S}$, $T_{C_1,opt,S}$, $\Delta Downtime$, and $\Delta Cost$ (see Equation 9 and 10). The results can be seen in Table 1 (where the last line corresponds to the results obtained with the true parameter value).

Sample Size	$\lambda_{h,MLE}$	$T_{D_1,opt,S}$	$T_{C_1,opt,S}$	Δ *Downtime*	Δ *Cost* (€)
2	0.0662	60.3735	67.7731	3.2699	1428.5680
5	0.0690	58.3667	64.6834	2.4081	1033.6236
10	0.0960	45.6374	47.2540	0.1278	49.4582
20	0.0946	46.1062	47.8410	0.0766	29.7284
60	0.0860	49.3381	51.9885	0.0654	26.0579
200	0.0862	49.2551	51.8798	0.0589	23.4314
600	0.0880	48.5267	50.9301	0.0159	6.2836
1000	0.0902	47.6784	49.8361	0.0002	0.0605
+∞	0.0900	47.7537	49.9327	0.0000	0.0000

Table 1: Optimal inspection corresponding to the different samples. The last line corresponds to the true parameter value $\lambda_{h,true}$= 0.09 failures/defect/month. The inspection periods and downtimes are given in days.

We can see that λ_h converges towards $\lambda_{h,true}$ while $\Delta Downtime$ (days) and $\Delta Cost$ (€) go to zero. However, for a very small sample size, the differences and sub-optimality are quite large.

4. Precise Bayesian Approach

In a situation where the only pieces of information we have about a parameter are its lower and upper bound, many precise Bayesians think that we ought to represent our ignorance through a uniform prior, which stems from the principle of maximum entropy Jaynes (1982). Consequently, we considered a uniform prior over the values of $\lambda_h \in [0.01; 0.1]$ failures/defect/month we shall call $f_{1,0}(\lambda_h)$, wherein 0 stands for *prior* and 1 stands for the index of the probability distribution (see Section 5). The posterior distribution given a sample S is obtained through a straightforward application of Bayes' theorem:

$$f(\lambda_h|S) = \frac{L(S|\lambda_h) f_{1,0}(\lambda_h)}{\int_{\lambda_h=0.01}^{0.1} L(S|\lambda_h) f_{1,0}(\lambda_h) d\lambda_h} \tag{12}$$

Given a sample S, the value of λ_h is estimated by

$$E(\lambda_h|S) = \int_{\lambda_h=0.01}^{0.1} \lambda_h f(\lambda_h|S) \tag{13}$$

[2]They can be obtained by sending an email to the author.

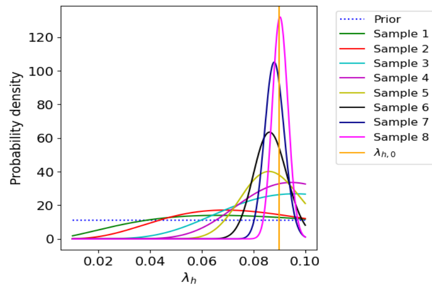

Figure 2: Uniform prior and posteriors for λ_h.

wherein $f(\lambda_h|S)$ is the prior when $S = \emptyset$. For a given inspection period T and a sample S, the expected value of D_1 with respect to λ_h is given by Eq. 14.

$$E_{\lambda_h}\Big(D_1(T,\lambda_h)\Big) = \int_{\lambda_h=0.01}^{0.1} D_1(T,\lambda_h) f(\lambda_h|S) d\lambda_h \quad (14)$$

Likewise, the expected value of C_1 with respect to λ_h is given by Eq. 15.

$$E_{\lambda_h}\Big(C_1(T,\lambda_h)\Big) = \int_{\lambda_h=0.01}^{0.1} C_1(T,\lambda_h) f(\lambda_h|S) d\lambda_h \quad (15)$$

These two quantities are the *posterior expected losses* with respect to the downtime and to the cost, respectively. Minimising such an expected loss function with respect to T is a standard decision criterion for Bayesian decision making Robert et al. (2007). The prior and posteriors can be seen in Figure 2 and the posterior estimates of λ_h can be found in Table 2. As could be expected, we can see that the posterior and expected value of λ_h converge towards the real value $\lambda_{h,true} = 0.09$ failures/defect/month.

Sample Size	$E(\lambda_h|S)$	$T_{D_1,opt,S}$	$T_{C_1,opt,S}$
0	0.0550	72.1056	88.7999
2	0.0613	64.9950	75.5063
5	0.0664	60.5123	68.0865
10	0.0785	52.9380	56.8024
20	0.0829	50.7739	53.8654
60	0.0843	50.0010	52.9380
200	0.0863	49.2281	51.8559
600	0.0881	48.4552	50.9285
1000	0.0903	47.6823	49.8464
$+\infty$	0.0900	47.7537	49.9327

Table 2: Optimal inspection periods corresponding to the different samples. The last line corresponds to the true parameter value $\lambda_{h,true}$= 0.09 failures/defect/month.

Table 2 also contains the optimal inspection periods for $D_1(T)$ and $C_1(T)$ which unsurprisingly converge towards those obtained for $\lambda_{h,true}$. In Table 3, we can see the differences between the optimal loss based on a prior (or a sample) and the optimal loss based on the true parameter value $\lambda_{h,true}$ for the precise Bayesian and frequentist approach. We can see that the differences are quite large in the absence of empirical data or for small sample sizes such as $n_Y = 2$ or $n_Y = 5$ failures. Interestingly enough, the frequentist decisions grounded on $\lambda_{h,MLE}$ turn out to be superior to the Bayesian ones based on the uniform priors for $n_Y \leq 60$ failures but a set of Monte-Carlo simulations would be required in order to prove the superiority of the first approach.

n_Y	$\Delta Downtime_B$	$\Delta Downtime_F$	$\Delta Cost_B$	$\Delta Cost_F$
2	5.5803	3.2699	2568.2310	1428.5680
5	3.3308	2.4081	1469.9209	1033.6236
10	0.6447	0.1278	261.8602	49.4582
20	0.2298	0.0766	91.3934	29.7284
60	0.1295	0.0654	54.4745	26.0579
200	0.0568	0.0589	22.8585	23.4314
600	0.0131	0.0159	6.2591	6.2836
1000	0.0001	0.0002	0.0481	0.0605
$+\infty$	0.0000	0.0000	0.0000	0.0000

Table 3: Sub-optimality of the decisions based on the prior/samples. The last line corresponds to the true parameter value $\lambda_{h,true}$= 0.09 failures/defect/month.

5. Imprecise Bayesian Approach

From a precise Bayesian standpoint, since we only know, in the absence of data, that $\lambda_h \in [0.01; 0.1]$ failures/defect/month, we assume that all values in that interval are equally likely so that we represent our initial ignorance through a uniform prior $f_{1,0}(\lambda_h)$ on this interval.

However, we are, logically speaking, equally ignorant about the values of any variable $Z = fun(\lambda_h)$ where fun is a deterministic function (that is not too extreme). We can, with equal justification, apply the principle of indifference to

$$Z \in [min_{\lambda_h \in [0.01;0.1]}\Big(fun(\lambda_h)\Big), \quad max_{\lambda_h \in [0.01;0.1]}\Big(fun(\lambda_h)\Big)]$$

and then deduce the cumulative probability distribution (cdf) for Z and thereupon the corresponding cdf for λ_h, the derivation of which leads to a prior pdf for the values of λ_h that is uniform, i.e., indifferent, with respect to the values of Z.

Let us now consider the mean delay time between the appearance of a defect and the breakdown of the pump $E(H)$. Since $H \sim exp(\lambda_h)$, we have $E(H) = \frac{1}{\lambda_h} \in [10; 100]$ months. Since we are equally ignorant about $E(H) \in [10; 100]$ months, we are, from a precise Bayesian point of view, also entitled to representing our ignorance through a uniform prior over $E(H)$ whose probability density as a function of λ_h is given by Eq. 16.

$$f_{2,0}(\lambda_h) = \frac{1}{\lambda_h^2} \frac{1}{\frac{1}{\lambda_{h,min}} - \frac{1}{\lambda_{h,max}}} \tag{16}$$

Given a very large number of pumps having a defect which appeared at time 0, let $prop_{Safe}(t, \lambda_h) = e^{-\lambda_h t}$ be the expected proportion of pumps still working at time t. For a given t, our maintenance engineer only knows that

$$prop_{Safe}(t, \lambda_h) \in [e^{-0.1t}, e^{-0.01t}] \tag{17}$$

If he or she decides to represent his/her ignorance through a uniform pdf with respect to $prop_{Safe}(t, \lambda_h)$, the pdf of λ_h is given by Eq. 18.

$$f_{3,0}(\lambda_h) = \frac{te^{-\lambda_h t}}{e^{-0.01t} - e^{-0.1t}} \tag{18}$$

If we are ignorant about the value of $prop_{safe}(t, \lambda_h)$, we are logically equally ignorant about the value of $\frac{1}{prop_{safe}(t,\lambda_h)}$ which belongs to the interval $[e^{0.01t}; e^{0.1t}]$. If the engineer decides to express his/her initial ignorance via a uniform prior over $\frac{1}{prop_{safe}(t,\lambda_h)}$, the corresponding pdf with respect to λ_h is given by Eq. 19.

$$f_{4,0}(\lambda_h) = t\frac{e^{\lambda_h t}}{e^{0.1t} - e^{0.01t}} \tag{19}$$

From a fundamental point of view, our engineer only knows that $\lambda_h \in [0.01; 0.1]$, which is equivalent to $E(H) \in [10; 100]$,

$$prop_{Safe}(t, \lambda_h) \in [e^{-0.1t}, e^{-0.01t}] \tag{20}$$

and

$$\frac{1}{prop_{safe}(t, \lambda_h)} \in [e^{0.01t}; e^{0.1t}]. \tag{21}$$

According to objective precise Bayesianism, he or she ought to represent his/her initial ignorance through a flat prior over the variable he/she is ignorant about. Unluckily, such a flat prior for, say, $E(H)$, would be non-uniform with respect to the three other Bayesian random variables and thus express a *specific knowledge* about their distributions even though we are supposed to only know their bounds Norton (2008). Moreover, it would be very hard and contrived to argue that we are actually ignorant about only one of these variables (such as λ_h) and that we would be *irrational* if we were to feel equally ignorant about the other variables. Instead, in such a situation our genuine ignorance can only be realistically expressed through a family of prior probability distributions Walley (1990, 2000).

Since it seems to be extremely hard (if not downright impossible) to build up a prior/posterior distribution conjugate to the likelihood function given by Eq.11, our engineer decides to represent his/her initial ignorance through a discrete set of four priors, namely $f_{1,0}$, $f_{2,0}$ and $f_{3,0}$ with t = 50 months and $f_{4,0}$ with t = 50 months (other values of t or even a whole interval of values are possible based on the engineer's subjective beliefs). Epistemically, this means that our engineer feels that she has no reasons to favour any (range of) values of λ_h (as the precise Bayesian in Section 4 did) but also no reasons to favour any values of $E(H)$, of the proportion of pumps (with an initial defect) still working 50 months later and of the multiplicative inverse of that proportion. The four priors can be visualised in Figure 3. For each of the four priors, as in Section 4, the posteriors based on the eight samples and the optimal inspection periods obtained by minimising the posterior expected loss with respect to D_1 and C_1 were computed. The posteriors can be seen in Figure 4-5. The lower and upper predictions for $E_i(\lambda_h|S)$ (i being the index of the prior) are compared with the frequentist and precise Bayesian estimate of λ_h in Table 4. It can be seen that the interval

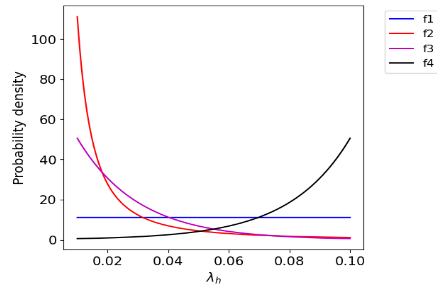

Figure 3: The four priors of λ_h.

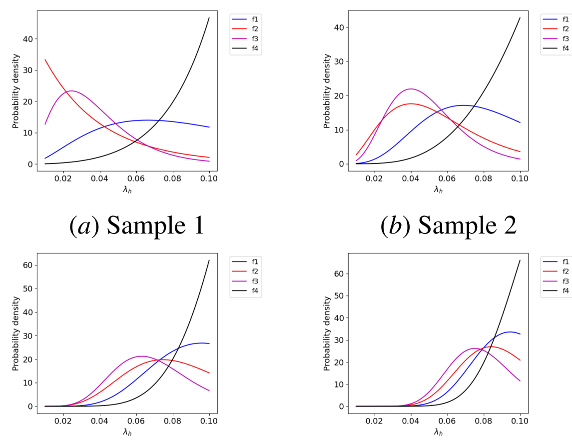

(*a*) Sample 1

(*b*) Sample 2

(*c*) Sample 3

(*d*) Sample 4

Figure 4: The four posteriors for Sample 1, 2, 3, and 4 based on the uniform prior.

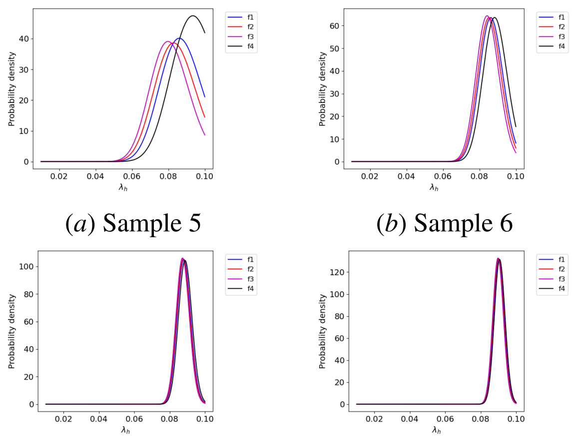

(a) Sample 5 (b) Sample 6

(c) Sample 7 (d) Sample 8

Figure 5: The four posteriors for Sample 5, 6, 7, and 8 based on the uniform prior.

n_Y	$\lambda_{h,MLE}$	$E_1(\lambda_h)$	$min_i\left(E_i(\lambda_h)\right)$	$max_i\left(E_i(\lambda_h)\right)$
0		0.0550	0.0256	0.0810
2	0.0662	0.0613	0.0364	0.0812
5	0.0690	0.0664	0.0475	0.0816
10	0.0960	0.0785	0.0655	0.0867
20	0.0946	0.0829	0.0748	0.0883
60	0.0860	0.0843	0.0800	0.0879
200	0.0862	0.0863	0.0845	0.0881
600	0.0880	0.0881	0.0874	0.0889
1000	0.0902	0.0903	0.0899	0.0908

Table 4: Expectation intervals for λ_h.

becomes smaller and smaller as the sample size n_Y grows but that it is quite large in the absence of data ([0.0256 ; 0.0810] failures/defect/month) and for a sample of only two failure times ([0.0364 ; 0.0812] failures/defect/month).

n_Y	$min(T_{D_1,opt})$	$max(T_{D_1,opt})$	$\Delta Downtime_B$	$\Delta Downtime_F$
0	51.7014	$+\infty$	9.8058	
2	51.5468	118.6333	5.5803	3.2699
5	51.3922	83.0805	3.3308	2.4081
10	49.0735	61.1306	0.6447	0.1278
20	48.3007	54.9475	0.2298	0.0766
60	48.6098	52.0105	0.1295	0.0654
200	48.4552	50.0010	0.0568	0.0589
600	48.1461	48.7644	0.0131	0.0159
1000	47.5278	47.8369	0.0001	0.0002

Table 5: Imprecise optimal inspection for D_1.

According to the criterion of E-admissibility Troffaes (2007), each inspection period which is optimal with respect to one of the four probability distributions representing the agent's uncertainty is admissible. The lower and upper posterior optimal inspection periods for D_1 and C_1 alongside the suboptimality of the frequentist and precise Bayesian

n_Y	$min(T_{C_1,opt})$	$max(T_{C_1,opt})$	$\Delta Cost_B$	$\Delta Cost_F$
0	55.1021	$+\infty$	4837.4707	
2	54.9475	$+\infty$	2568.2310	1428.5680
5	54.6383	114.7689	1469.9209	1033.6236
10	51.5468	69.0140	261.8602	49.4582
20	50.7739	59.5848	91.3934	29.7284
60	50.9285	55.5658	54.4745	26.0579
200	50.9285	52.7834	22.8585	23.4314
600	50.4647	51.2376	6.2591	6.2836
1000	49.5373	50.0010	0.0481	0.0605

Table 6: Imprecise optimal inspection for C_1.

decisions are shown in Table 5 and 6. For large sample sizes (such as $n_Y = 1000$ and $n_Y = 600$), the intervals for $T_{D_1,opt,S}$ and $T_{C_1,opt,S}$ are very narrow and the suboptimality of the frequentist and precise Bayesian decision (represented by $\Delta Downtime_F$, $\Delta Downtime_B$, $\Delta Cost_F$, and $\Delta Cost_B$) is quite small.

Nevertheless, in the absence of any data, the optimal inspection period for minimising D_1 belongs to the interval [51.7014 ; $+\infty$] days. This means that the optimal decision is extremely uncertain, which corresponds to the fact that the precise Bayesian decision would lead to a sub-optimality of $\Delta Downtime_B = 9.8058$ days of downtime in comparison to the true optimal inspection period that could be calculated if we knew the real value of λ_h. Likewise, for $n_Y = 2$ failure times, the optimal inspection periods for minimising C_1 belongs to the interval [54.9475 ; $+\infty$] days. This extremely strong imprecision corresponds to the fact that the frequentist and precise Bayesian decisions would be highly sub-optimal in that $\Delta Cost_B = 2568.2310$ € and $\Delta Cost_F = 1428.5680$ €.

6. Discussion and Conclusions

In this work, we have developed and applied an imprecise Bayesian approach to delay time modelling for the choice of optimal inspection periods in industrial maintenance. Eight samples of virtual data have been generated and the frequentist and precise Bayesian approaches were compared with the newly developed imprecise Bayesian approach. While the three approaches give very similar results for large sample sizes ("the priors wash out" Hawthorne (1994)) , they strongly differ for small sample sizes or in the absence of any data. In contrast to other imprecise decision criteria such as Γ-maximax and Γ-maximin Troffaes (2007) that would have returned a single decision value for the inspection period, E-admissibility returns a set of inspection periods. Far from being a defect, it allows the engineer to realise that when the interval returned by the imprecise probabilistic approach is very large (or even infinite), the evidential basis at his/her disposal is so thin that it is not possible to tell which inspection period would be optimal. A practical consequence of this result is that the enterprise needs to collect more data in order to make an informed decision.

Opponents of imprecise probabilities often argue that sharp, single-valued probabilities are all we need to represent all forms of uncertainty. For instance, Lindley wrote that

> Whatever way uncertainty is approached, probability is the only sound way to think about it

Lindley (2013). The results obtained through this study undermine that notion. The precise Bayesian method always returns a single optimal inspection period without any regard for its reliability and the level of information it is based upon. To be sure, a compentent Bayesian statistician would be able to see that the evidential basis is much stronger for $n_Y = 1000$ than for $n_Y = 2$ because the posterior pdf of λ_h is much flatter in the latter case. However, the shape of the probability distribution of λ_h has no bearing upon the precise Bayesian optimal decision that is always based on minimising the posterior *expected* loss without considering the variance of the loss Bradley (2019). What is more, experience shows that time and time again many classical Bayesians draw very strong conclusions that are entirely grounded on the prior probability distribution and not on any empirical data Benetreau-Dupin (2015); Norton (2010). A precise Bayesian who wants to avoid these pitfalls must resort to *ad-hoc* reasoning, which contradicts the aspiration of precise Bayesianism to be a universal holistic framework for dealing with every sort of uncertainty as expressed by Lindley (2013).

In contrast, an imprecise Bayesian approach relying on E-admissibility allows one to make a distinction between different degrees of knowledge and ignorance, as Sturgeon elegantly expressed it:

> Evidence and attitude aptly based on it must match in character. When evidence is essentially sharp, it warrants sharp or exact attitude; when evidence is essentially fuzzy -as it is most of the time- it warrants at best a fuzzy attitude.

Sturgeon (2008) It is worth noting that the present study also undermines the use of frequentist methods that are often defended on the grounds that they do not rely on a prior that is not capable of expressing genuine ignorance. Like the precise Bayesian approach, the frequentist approach also returns a single number that does not indicate how strongly the empirical evidence truly supports the decision.

Finally, while we considered a finite set of priors in this study as an easy way to demonstrate how the inclusion of imprecision can reveal the difference between ignorance, knowledge and degrees of ignorance, it might be more epistemically realistic to consider a *continuous* family of priors that are to be updated. This could, for example, be realised by considering the function sets

$$f_{0,t}(\lambda_h) = t\frac{e^{-\lambda_h t}}{e^{-0.01t} - e^{-0.1t}} \tag{22}$$

and

$$g_{0,t}(\lambda_h) = t\frac{e^{\lambda_h t}}{e^{0.1t} - e^{0.01t}} \tag{23}$$

and defining the prior set

$$\mathfrak{I}_0 = \{f_{0,t}, g_{0,t}, t \in [t_{min}, t_{max}]\} \tag{24}$$

with $\mathfrak{I}_0$ becoming the vacuous prior set as $t_{max} \to +\infty$. Developing this thought would, however, demand much further work.

Finally, it is worth noting that the imprecise Bayesian approach developed and tested throughout this study can and should also be applied to more complex delay time models such as those involving multi-component systems where several defects can accumulate before a breakdown and where the probability distributions of the defect appearance time and the delay time are best approximated through two Weibull distributions Wang (2012a).

Acknowledgments

I am most thankful to Prof. Matthias Troffaes and Dr. Erik Quaeghebeur for their useful advice and to Prof. Enrique Miranda for his technical help with Latex!

Author Contributions

The entire work was done by Dr. Marc Fischer.

References

RD Baker and W Wang. Estimating the delay-time distribution of faults in repairable machinery from failure data. *IMA Journal of Management Mathematics*, 3(4): 259–281, 1991.

RD Baker and W Wang. Developing and testing the delay-time model. *Journal of the Operational Research Society*, 44(4):361–374, 1993.

Y Benetreau-Dupin. The Bayesian who knew too much. *Synthese*, 2015.

Seamus Bradley. Imprecise Probabilities. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2019 edition, 2019.

A H Christer and W M Waller. Delay time models of industrial inspection maintenance problems. *Journal of the Operational Research Society*, 35(5):401–406, 1984.

AH Christer, Wenbin Wang, RD Baker, and J Sharp. Modelling maintenance practice of production plant using the delay-time concept. *IMA Journal of Management Mathematics*, 6(1):67–83, 1995.

Anthony H Christer. Developments in delay time analysis for modelling plant maintenance. *Journal of the Operational Research Society*, 50(11):1120–1137, 1999.

Marc Fischer. On the usefulness of imprecise Bayesianism in chemical kinetics. In *International Symposium on Imprecise Probabilities: Theories and Applications*, pages 203–215, 2019.

Marc Fischer and Alexis Vignes. An imprecise bayesian approach to thermal runaway probability. In *International Symposium on Imprecise Probability: Theories and Applications*, pages 150–160. PMLR, 2021.

James Hawthorne. On the nature of Bayesian convergence. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, number 1, pages 241–249. Philosophy of Science Association, 1994.

Edwin T Jaynes. On the rationale of maximum-entropy methods. *Proceedings of the IEEE*, 70(9):939–952, 1982.

Dennis V Lindley. *Understanding uncertainty*. John Wiley & Sons, 2013.

J D Norton. Ignorance and Indifference. *Philosophy of Science*, 75:45–68, 2008.

J D Norton. Cosmic Confusions: Not Supporting versus Supporting Not. *Philosophy of Science*, 77(4):501–523, 2010.

Christian P Robert et al. *The Bayesian choice: from decision-theoretic foundations to computational implementation*, volume 2. Springer, 2007.

Philip A Scarf, Cristiano AV Cavalcante, and Rodrigo Sampaio Lopes. Delay-time modelling of a critical system subject to random inspections. *European Journal of Operational Research*, 278(3):772–782, 2019.

Scott Sturgeon. Reason and the grain of belief. *Noûs*, 42 (1):139–165, 2008.

Matthias CM Troffaes. Decision making under uncertainty using imprecise probabilities. *International journal of approximate reasoning*, 45(1):17–29, 2007.

P Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall CRC Monographs on Statistics & Applied Probability 42, 1990.

P Walley. Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning*, 24: 125–148, 2000.

W Wang and X Jia. An empirical bayesian based approach to delay time inspection model parameters estimation using both subjective and objective data. *Quality and Reliability Engineering International*, 23(1):95–105, 2007.

Wenbin Wang. Delay time modelling. In *Complex system maintenance handbook*, pages 345–370. Springer, 2008.

Wenbin Wang. An overview of the recent advances in delay-time-based maintenance modelling. *Reliability Engineering & System Safety*, 106:165–178, 2012a.

Wenbin Wang. A simulation-based multivariate bayesian control chart for real time condition-based maintenance of complex systems. *European Journal of Operational Research*, 218(3):726–734, 2012b.