AIMS Mathematics

AIMS Mathematics, 11(9): 28502–28546.
DOI: 10.3934/math.20261136
Received: 26 May 2026
Revised: 12 July 2026
Accepted: 12 August 2026
Published: 07 September 2026

https://www.aimspress.com/journal/Math

Research article

Semi-parametric distributional inference for mean residual life with latent random inspection

Thamer Manshi*

Department of Statistics and Operations Research, College of Science, King Saud University, P.O. Box 2455, Riyadh 11451, Saudi Arabia

* **Correspondence:** Email: tmunshi@ksu.edu.sa; Tel: +966507466147.

Abstract: In this paper, a semi-parametric distributional inference framework was developed for mean residual life analysis when the inspection age is latent, random, or only indirectly specified through a stochastic monitoring mechanism. Classical mean residual life methods evaluate the expected remaining lifetime at a fixed and fully observed age, an assumption that may be restrictive in reliability, survival, and predictive-maintenance studies where inspection times are irregular, missing, or governed by external monitoring policies. For a lifetime variable T and an independent random inspection age X, the proposed framework targeted the conditional residual-life functional $(T - X \mid T > X)$, thereby incorporating uncertainty in the observation process directly into lifetime-distribution analysis. The lifetime distribution was estimated nonparametrically, whereas the latent inspection-age distribution was modeled parametrically, leading to a flexible and implementable inference procedure for complete and right-censored lifetime data. Under the exponential inspection model, explicit finite-sum representations were derived, avoiding numerical double integration and enabling direct computation from ordered lifetime observations and Kaplan-Meier survival estimates. The theoretical development established strong consistency and derived the asymptotic normality of the complete-data estimator. A plug-in estimator of the asymptotic variance was also introduced, yielding analytically justified standard errors and large-sample confidence intervals. Simulation studies under Weibull and gamma lifetime distributions demonstrated coherent finite-sample performance across aging regimes, inspection intensities, sample sizes, and censoring proportions. A fixed-age benchmark comparison further showed that replacing the random inspection age by its mean may conceal important distributional effects of monitoring uncertainty. Real-data analyses using kidney catheter survival data, aircraft air-conditioning failure times, and NASA C-MAPSS turbofan degradation lifetimes illustrated the practical value of the proposed methodology in biomedical, engineering, and predictive-maintenance applications. The proposed approach therefore contributes to statistical distribution theory and its applications by extending classical residual-life inference to modern lifetime-data settings in which observation ages are unavailable, irregularly recorded, or operationally uncertain.

Keywords: mean residual life; semi-parametric distributional inference; latent random inspection; lifetime distributions; right-censored data; Kaplan-Meier estimator; reliability data analysis; predictive maintenance
Mathematics Subject Classification: 62N05, 62G05, 62F10, 62E05

1. Introduction

Statistical distribution theory plays a central role in the analysis of lifetime phenomena by providing the probabilistic foundations for modeling survival behavior, quantifying aging patterns, and developing inference procedures for complex data-generating mechanisms. In reliability engineering, biomedical survival studies, and predictive maintenance, lifetime distributions are often investigated not only through their survival and failure-rate functions, but also through residual-life functionals that provide interpretable summaries of remaining useful life. Among these functionals, the mean residual life occupies a particularly important position, as it measures the expected future lifetime of a unit conditional on survival beyond a given age. Classical mean residual life analysis, however, typically assumes that the conditioning age is fixed, fully observed, and externally specified. This assumption can be unrealistic in many modern lifetime-data applications, where observation schemes may involve irregular monitoring, incomplete follow-up, right censoring, or inspection ages that are missing, latent, or indirectly determined by a stochastic monitoring process. These features create a methodological need for distributional inference tools that preserve theoretical rigor while accommodating realistic observation mechanisms and remaining practically implementable. Motivated by this perspective, a semi-parametric distributional inference framework is developed for mean residual life analysis under latent random inspection, with direct applicability to complete and right-censored lifetime data. Let T be a non-negative lifetime random variable with cumulative distribution function $F(t)$ and survival function $\overline{F}(t)=1-F(t)$, where $\overline{F}(t)>0$. The classical mean residual life (MRL) function at age t is defined by

$$m(t)=\frac{\int_t^{\infty}\overline{F}(u)du}{\overline{F}(t)},\overline{F}(t)>0 .$$

This function provides a dynamic summary of future lifetime behavior that is not captured by the marginal lifetime distribution alone. It is therefore widely used in reliability assessment, preventive maintenance, biomedical survival analysis, risk evaluation, and lifetime prediction. Theoretical and inferential developments for $m(t)$ and related residual-life quantities include studies of mixture and system lifetimes, censored regression, and ageing comparisons [1–3]. Further regression developments include varying-coefficient and change-point models [4,5]. Artificial neural networks have also been used to model exponential-type and Weibull lifetime distributions [6–8]. Other applications concern software reliability [9], preventive maintenance of robotic systems [10], predictive maintenance under alternative inspection strategies [11], and reliability of stochastic shipping services [12]. Beyond lifetime applications, robust covariance estimation has been investigated for financial risk management [13]. These studies illustrate the broader role of statistical modeling in reliability and risk assessment.

A common feature of the classical formulation is that the age t at which residual life is evaluated is assumed to be fixed, deterministic, and fully observed. This assumption is natural in controlled laboratory experiments or planned life-testing studies, but it can be restrictive in many observational and field-based applications. In engineering reliability, biomedical follow-up, environmental surveillance, and predictive-maintenance systems, the effective inspection age may be irregular, incompletely recorded, indirectly specified, or generated by an external monitoring process. Thus, the time at which a surviving unit is assessed may not be a controlled deterministic input. Instead, it may be better represented as a random inspection age. Ignoring this additional source of observation uncertainty may lead to an incomplete description of residual lifetime behavior, particularly when the inspection mechanism is distinct from the underlying failure mechanism.

In this paper, mean residual life analysis is considered with a stochastic observation age. Let X be a non-negative random inspection age, independent of T, with distribution function G. The residual lifetime observed at a random inspection age is then defined by $(T-X \mid T>X)$ and its expectation represents the mean remaining lifetime of a unit that is inspected at age X and is found to have survived beyond that age. This functional extends the classical fixed-age mean residual life function by averaging residual-life behavior over a random inspection mechanism. It is particularly relevant when inspection ages are latent, model-based, irregularly recorded, or not jointly observed with the lifetime data.

Residual life at a random age has been studied in relation to variability measures [14] and preservation of generalized ageing classes [15]. Further developments address variability of mean remaining lifetime [16] and generalized orderings for residual life and inactivity time at random times [17]. These contributions show that random-age residual life may reveal aspects of lifetime behavior that are not captured by fixed-age analysis alone. However, statistical inference for this functional remains comparatively limited, especially when the inspection process is not observed together with the lifetime variable. This situation is common in practice: One may observe failure times, censored survival times, or run-to-failure degradation lifetimes, while the corresponding inspection ages are missing, indirectly documented, or known only through a monitoring policy. A fully parametric lifetime model may then be unnecessarily restrictive, whereas completely ignoring the inspection mechanism may obscure the effect of observation uncertainty on residual-life conclusions.

To address this problem, a semi-parametric estimation framework is developed for mean residual life with latent random inspection. The lifetime distribution is estimated nonparametrically from the available lifetime data, whereas the inspection-age distribution is specified parametrically to represent a plausible stochastic monitoring mechanism. This construction preserves flexibility in modeling the lifetime distribution while explicitly incorporating information about the latent observation process. The resulting estimator is applicable to complete lifetime data and is extended to right-censored data by replacing the empirical survival estimator with the Kaplan-Meier estimator.

A practically important special case is obtained when the latent inspection age follows an exponential distribution. This model corresponds to a Poisson-type monitoring mechanism and provides a natural baseline for stochastic inspection. Under this specification, the proposed estimator admits an explicit finite-sum representation in terms of ordered lifetime observations. For right-censored data, an analogous computable form is obtained using the step-function structure of the Kaplan–Meier estimator. These representations avoid numerical double integration and make the method directly implementable in applied reliability and survival studies. Related studies consider mean survival time for time-to-event data [18] and nonparametric mean residual life estimation under random censorship [19]. The limiting behavior of mean residual life has also been investigated [20]. Other contributions use mean residual life to summarize censored survival data [21], compare two mean residual life functions [22], and develop nonparametric estimators [23].

The principal contribution of this paper is the development of a fully implementable semi-parametric inference framework for mean residual life analysis when the inspection age is latent, random, and unobserved rather than fixed and directly recorded. While studies on residual life at random ages have primarily focused on distributional properties, aging classes, and stochastic-ordering characterizations, in this work, the fundamentally different problem of statistical inference from lifetime data is addressed when inspection ages are unavailable or only indirectly represented through a stochastic monitoring mechanism. Specifically, the proposed methodology estimates the conditional residual-life functional $(T-X \mid T > X)$ by combining nonparametric estimation of the lifetime distribution with a parsimonious parametric model for the inspection process. The resulting framework is applicable to complete and right-censored lifetime data, admits explicit finite-sum computational representations under exponential inspection, and naturally supports sensitivity analysis with respect to plausible latent inspection mechanisms. Consequently, the proposed methodology extends random-age mean residual life from a predominantly theoretical concept to a practically implementable inferential framework for estimation, uncertainty quantification, and large-sample analysis in settings involving incomplete observation and latent monitoring.

Compared with classical fixed-age mean residual life methods, the proposed framework explicitly accounts for uncertainty in the inspection age instead of conditioning on a single observed inspection time. In contrast to fully parametric approaches, it leaves the lifetime distribution unspecified and estimates it nonparametrically, thereby improving robustness against lifetime-model misspecification while retaining a transparent parametric representation of the inspection mechanism. Relative to nonparametric mean residual life estimators, its distinguishing feature is the explicit incorporation of a latent inspection process, enabling valid inference even when inspection ages are unavailable. Furthermore, the methodology has several important limitations. The inspection-age distribution must be specified through a parametric family, the baseline theoretical development assumes independence between the lifetime and inspection processes, and the right-censored analysis is restricted to an identifiable finite observation horizon. Accordingly, the proposed approach should be viewed as a flexible semi-parametric sensitivity-analysis framework that complements, rather than replaces, conventional fixed-age and fully parametric residual-life methods.

More broadly, the proposed framework aligns with the growing demand for statistically interpretable prognostic methodologies in reliability engineering and predictive maintenance, where monitoring schedules are frequently irregular, incomplete, or only partially observed. Moreover, developments in prognostics and predictive maintenance have likewise highlighted the importance of reliable remaining-life prediction under uncertain monitoring and incomplete degradation information, further reinforcing the practical motivation for flexible inferential procedures capable of accommodating realistic inspection mechanisms without sacrificing statistical interpretability [11].

The principal contributions of the proposed methodology are summarized below:

(1) A latent random-inspection mean residual life functional is introduced for settings in which the inspection age is stochastic, unobserved, or only indirectly described through an external monitoring mechanism.

(2) A semi-parametric estimation framework is developed by combining nonparametric estimation of the lifetime survival function with a parametric specification of the latent inspection-age distribution.

(3) Explicit finite-sum representations are derived under the exponential inspection model, yielding directly computable formulas based on ordered lifetime observations and avoiding numerical double integration.
(4) The methodology is extended to right-censored lifetime data through the Kaplan-Meier estimator, with a finite-horizon formulation that avoids extrapolation beyond the identifiable support.
(5) Strong consistency is established under mild regularity conditions, ensuring the large-sample validity of the proposed estimator.
(6) A central limit theorem is derived for the complete-data estimator, providing asymptotic normality and quantifying its first-order sampling variability.
(7) A plug-in estimator of the asymptotic variance is proposed, leading to analytically justified standard errors and large-sample confidence intervals for the latent random-inspection mean residual life.
(8) Practical guidance is provided for interpreting and selecting the inspection-intensity parameter in the exponential monitoring model, including its use in sensitivity analysis when inspection ages are unavailable.
(9) Comprehensive simulation studies under Weibull and gamma lifetime models assess finite-sample performance across aging regimes, sample sizes, inspection intensities, and censoring proportions.
(10) Real-data analyses using biomedical survival data, classical engineering failure-time data, and NASA C-MAPSS turbofan degradation lifetimes demonstrate the practical value of the proposed framework for latent-monitoring and incomplete-observation settings.

The empirical analyses are intentionally interpreted as sensitivity analyses rather than validation studies based on observed inspection times. The three datasets provide survival times, failure times, or degradation-based lifetimes, but they do not include directly recorded random inspection ages. As a result, the proposed framework is used to study how estimated residual life changes across plausible values of the latent inspection intensity. This interpretation is central to the paper: The purpose is not to reconstruct unobserved inspection histories, but to provide a coherent residual-life analysis when the observation process is latent, incomplete, or only indirectly specified.

From an applied perspective, the proposed framework is particularly useful in situations where reliable lifetime information is available but the ages at which systems are inspected are not directly observed or are only partially documented. Such situations frequently arise in reliability engineering, biomedical follow-up studies, and predictive maintenance, where inspection schedules may depend on operational constraints, maintenance policies, or administrative decisions rather than fixed experimental designs. Classical mean residual life methods implicitly assume that the evaluation age is known and deterministic, whereas the proposed framework explicitly incorporates uncertainty in the inspection process through a latent random inspection model. Consequently, practitioners can assess the sensitivity of residual-life estimates to plausible monitoring mechanisms without imposing a fully parametric model for the lifetime distribution, thereby providing a more realistic basis for maintenance planning, survival assessment, and reliability decision making. An important practical advantage of the proposed framework is that it does not require the latent inspection process to be directly observed. Instead, investigators can assess how residual-life conclusions change under scientifically plausible monitoring mechanisms while preserving a nonparametric description of the lifetime distribution. This makes the methodology particularly attractive for retrospective reliability studies, biomedical registries, maintenance databases, and condition-monitoring systems, where event times are routinely recorded but inspection histories are often incomplete or unavailable.

The remainder of the paper is organized as follows: In Section 2, the proposed semi-parametric estimator is introduced and its theoretical properties are established. The exponential inspection model and the extension to right-censored data are also developed in this section. In Section 3, simulation studies are presented to evaluate finite-sample performance under different lifetime distributions, inspection intensities, sample sizes, and censoring proportions. In Section 4, real-data sensitivity analyses are provided using biomedical survival, engineering reliability, and degradation-based predictive-maintenance datasets. In Section 5, practical guidance, robustness, limitations, and future research directions are discussed. Section 6 concludes the paper.

Table 1 summarizes the notation used throughout the paper, including the lifetime and latent inspection variables, their distribution and survival functions, the mean residual life functionals and estimators, and the quantities required for the right-censored formulation.

Table 1. Summary of notation used throughout the paper.

Symbol	Description	Remarks
T	Lifetime random variable	Non-negative
X	Latent random inspection age	Independent of T
F	CDF of T	
$\overline{F}$	Survival function of T	
G	CDF of X	Specified parametrically
$\overline{G}$	Survival function of X	
$m(t)$	Classical mean residual life	$E(T-t \mid T>t)$
mr	Random-inspection mean residual life	$E(T-X \mid T>X)$
mr_n	Semi-parametric estimator of mr	
$\widehat{mr}_{KM}$	KM-based estimator of mr	Right-censored data
λ	Inspection intensity	Exponential inspection
n	Sample size	
τ	Finite analysis horizon	Censored case
Y_i	Observed time	$\min(T_i, C_i)$
δ_i	Failure indicator	$I(T_i<C_i)$
C_i	Censoring time	Independent of T

2. Semi-parametric estimation under latent random inspection

In this section, the proposed semi-parametric inference framework is developed for mean residual life under latent random inspection. The general formulation is first introduced, its asymptotic properties are then established, and the methodology is finally extended to exponentially distributed inspection ages and right-censored lifetime data. Let T be a non-negative lifetime variable with distribution function F and survival function $\overline{F}$. Let X be a non-negative random inspection age with distribution function G and survival function $\overline{G}=1-G$. Throughout this section, T and X are assumed to be independent. This assumption represents a monitoring structure in which the failure mechanism and the observation mechanism are driven by separate sources of randomness. It is natural in many reliability and survival studies where inspections are generated by external monitoring protocols, random visits, administrative observation rules, or operational schedules rather than by the unobserved degradation state of the unit.

A central feature of the present formulation is that the inspection variable X is treated as a latent or model-based inspection age. Thus, X is not required to be observed jointly with the lifetime data. Instead, its distribution is specified through a stochastic monitoring model that represents plausible information about the inspection process. This perspective is particularly useful when lifetime or failure data are available, but the corresponding inspection ages are missing, indirectly specified, or not recorded in the original data collection protocol.

The residual lifetime at a random inspection age is defined by $(T-X \mid T>X)$. The corresponding mean residual life under latent random inspection is $mr=E\left(T-X \mid T>X\right)$. By independence of T and X, this functional admits the representation

$$mr=\frac{\int_0^\infty \int_0^\infty \bar{F}(x+y)dydG(x)}{\int_0^\infty \bar{F}(x)dG(x)}, \tag{2.1}$$

whenever the numerator is finite and the denominator is strictly positive. The denominator in (2.1) is $P(T>X)$, the probability that the unit is still operating at the latent inspection age, while the numerator aggregates the remaining lifetime over the joint survival-inspection structure. Hence, *mr* extends the classical fixed-age mean residual life function by incorporating uncertainty in the age at which the unit is inspected.

The overall estimation and inference procedure is summarized in Figure 1.

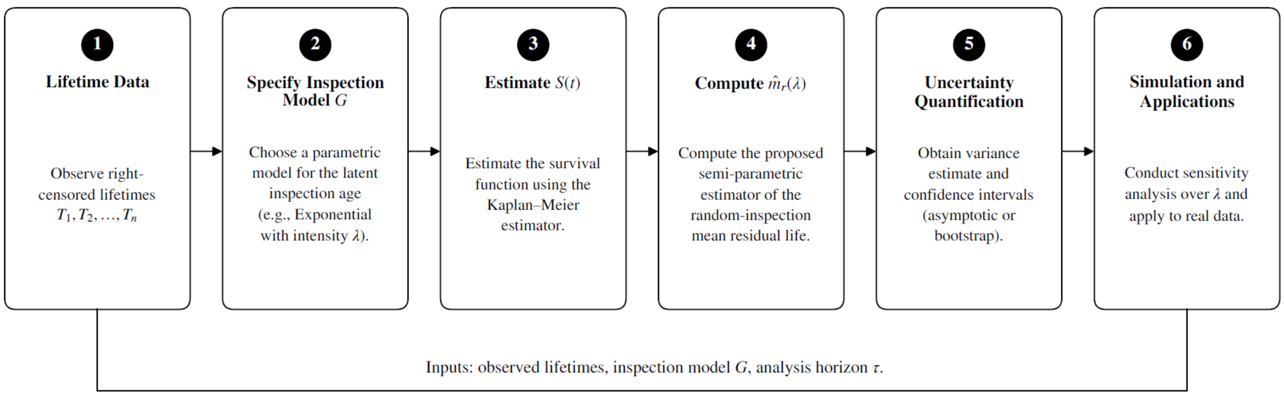

Figure 1. Flowchart of the proposed semi-parametric estimation and inference procedure under latent random inspection.

2.1. Semi-parametric plug-in estimation

Suppose first that complete lifetime observations $T_1,\ldots,T_n$ are available from F. The lifetime distribution is left unspecified and is estimated nonparametrically through the empirical survival function

$$\bar{F}_n\left(t\right)=\frac{1}{n}\sum_{i=1}^{n} I\left(T_i>t\right), \tag{2.2}$$

where $I(\cdot)$ denotes the indicator function. The strict or non-strict inequality at jump points is immaterial for the integral-based functionals considered below. The inspection-age distribution G, in contrast, is specified parametrically according to the assumed stochastic monitoring mechanism. The proposed semi-parametric estimator is obtained by replacing $\bar{F}$ in (2.1) with $\bar{F}_n$:

$$mr_n = \frac{\int_0^\infty \int_0^\infty \bar{F}_n(x+y)dydG(x)}{\int_0^\infty \bar{F}_n(x)dG(x)}, \tag{2.3}$$

The estimator in (2.3) is well defined whenever the denominator is positive. In finite samples, this denominator represents the empirical analogue of the probability that a unit survives beyond the latent inspection age. Throughout the asymptotic development, the inspection distribution G is treated as fixed and specified independently of the sample used to estimate the lifetime distribution. This reflects the intended semi-parametric structure of the method: Uncertainty in the lifetime law is handled nonparametrically, whereas the latent inspection mechanism is incorporated through a prescribed stochastic monitoring model. The numerator averages the remaining lifetime over the latent inspection mechanism, whereas the denominator represents the probability that a unit is functioning when inspected. Consequently, the estimator can be viewed as a weighted empirical analogue of the proposed random-inspection mean residual life functional.

Throughout the asymptotic analysis, the following regularity conditions are assumed. Conditions (C1)–(C4) are imposed for the complete-data analysis, whereas Conditions (C5)–(C7) are also required for the right-censored extension:

(C1) $T_1,\ldots,T_n$ are independent observations from the lifetime distribution F.

(C2) X is a non-negative inspection age with distribution function G, independent of T.

(C3) $E(T)<\infty$ and $D=P(T>X)>0$.

(C4) $\mathrm{E}(T^2) < \infty$.

(C5**)** The censoring variables $C_1,\ldots,C_n$ are independent and identically distributed and are independent of the lifetime variables $T_1,\ldots,T_n$.

(C6**)** The analysis horizon $\tau<\infty$ is fixed and lies within the identifiable support, and the population denominator of the truncated functional is strictly positive.

(C7**)** The standard support and regularity conditions ensuring uniform consistency of the Kaplan-Meier estimator on $[0, \tau]$ are satisfied.

The following theorem establishes the basic large-sample validity of the estimator.

Theorem 2.1. Suppose Conditions (C1)–(C3) hold. Then the semi-parametric estimator mr_n defined in (2.3) satisfies $mr_n \to mr$ almost surely as $n\to\infty$.

Proof. Denote

$$N=\int_0^\infty \int_0^\infty \bar{F}(x+y)dydG(x), \qquad D=\int_0^\infty \bar{F}(x)dG(x).$$

Then, $mr = N/D$, where $D=P(T>X)>0$. Similarly, the plug-in estimator can be written as $mr_n = N_n / D_n$, where

$$N_n = \int_0^\infty \int_0^\infty \bar{F}_n(x+y)\, dy dG(x), \qquad D_n = \int_0^\infty \bar{F}_n(x)\, dG(x).$$

Using $\bar{F}_n(t) = \frac{1}{n}\sum_{i=1}^{n} I(T_i > t)$, the following expression is obtained

$$D_n = \frac{1}{n}\sum_{i=1}^{n}\int_0^\infty I(T_i > x)\, dG(x).$$

Hence,

$$D_n = \frac{1}{n}\sum_{i=1}^{n} G(T_i -),$$

where $G(T_i -) = P(X < T_i)$. If G is continuous, this expression reduces to $G(T_i)$. Since $0 \le G(T_i -) \le 1$, the strong law of large numbers gives

$$D_n \to E\{G(T-)\} = P(X < T) = P(T > X) = D \text{ almost surely.}$$

For the numerator,

$$N_n = \frac{1}{n}\sum_{i=1}^{n}\int_0^\infty \int_0^\infty I(T_i > x+y)\, dy dG(x).$$

For each fixed T_i,

$$\int_0^\infty I(T_i > x+y)\, dy = (T_i - x)_+ .$$

Therefore,

$$N_n = \frac{1}{n}\sum_{i=1}^{n}\int_0^\infty (T_i - x)_+\, dG(x).$$

Define

$$H_G(t) = \int_0^\infty (t-x)_+\, dG(x) = \int_0^t (t-x) dG(x),\ t \ge 0 .$$

Then

$$N_n = \frac{1}{n}\sum_{i=1}^{n} H_G(T_i).$$

Moreover, $0 \le H_G(t) \le t$, so $E\{H_G(T)\} \le E(T) < \infty$. By the strong law of large numbers,

$$N_n \to E\{H_G(T)\} \text{ almost surely.}$$

Using Tonelli's theorem,

$$E\{H_G(T)\} = E\left[\int_0^\infty (T-x)_+ dG(x)\right] = \int_0^\infty \int_0^\infty \overline{F}(x+y)\,dy\,dG(x) = N .$$

Thus, $N_n \to N$ and $D_n \to D > 0$ almost surely. Therefore, $D_n > 0$ eventually with probability one, and by the continuous mapping theorem,

$$mr_n = \frac{N_n}{D_n} \to \frac{N}{D} = mr \text{ , almost surely.}$$

This completes the proof. □

Remark 2.1. *The proof shows that the consistency result follows from the strong law of large numbers applied to two empirical averages. No parametric assumption is imposed on the lifetime distribution F. The condition* $E(T) < \infty$ *controls the residual-life numerator, while* $P(T > X) > 0$ *ensures that the limiting denominator is nonzero. When the inspection distribution G is continuous, as in the exponential inspection model used below,* $G(T-)$ *may be replaced by* $G(T)$.

To complete the inferential theory of the proposed semi-parametric framework, the asymptotic distribution of mr_n is next derived for complete lifetime data. Unlike strong consistency, this result quantifies the estimator's first-order sampling variability. Under a mild second-moment condition, the resulting central limit theorem provides the basis for asymptotic standard errors, confidence intervals, and rigorous large-sample inference for the latent random-inspection mean residual life.

Theorem 2.2. Suppose Conditions (C1)–(C4) hold. Define $Q_G(t) = G(t-) = P(X < t)$, and

$$H_G(t) = \int_0^\infty (t-x)_+ dG(x) = \int_0^t (t-x)\,dG(x), \; t \geq 0 .$$

Then the estimator mr_n defined in (2.3) satisfies

$$\sqrt{n}\{mr_n - mr\} \xrightarrow{d} N\left(0, \sigma_G^2\right),$$

where

$$\sigma_G^2 = \frac{1}{D^2} \mathrm{Var}\{H_G(T) - mrQ_G(T)\} .$$

If $\sigma_G^2 = 0$, the limiting distribution is understood in the degenerate sense.

Proof. From the representation used in the proof of Theorem 2.1, the numerator and denominator of $m_{r,n}$ can be written as

$$N_n = \frac{1}{n}\sum_{i=1}^{n} H_G(T_i), \qquad D_n = \frac{1}{n}\sum_{i=1}^{n} Q_G(T_i),$$

where

$$H_G(t) = \int_0^{\infty} (t-x)_+ \, dG(x), \qquad Q_G(t) = G(t-).$$

The corresponding population quantities are

$$N = E\{H_G(T)\}, \qquad D = E\{Q_G(T)\} = P(T > X),$$

and therefore $mr = N / D$. Since $0 \le Q_G(T) \le 1$, $Q_G(T)$ has a finite second moment. Additionally, $0 \le H_G(T) \le T$, and hence $E(T^2) < \infty$ implies $E\{H_G^2(T)\} < \infty$. Thus, by the multivariate central limit theorem,

$$\sqrt{n}\begin{pmatrix} N_n - N \\ D_n - D \end{pmatrix} \xrightarrow{d} N_2\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_G\right],$$

where Σ_G is the covariance matrix of $(H_G(T), Q_G(T))^{\top}$. Now consider $\phi(a,b) = \frac{a}{b}$, $b > 0$. Since $D > 0$, ϕ is differentiable at (*N, D*), and

$$\nabla\phi(N,D) = \left(\frac{1}{D}, -\frac{N}{D^2}\right)^{\top} = \left(\frac{1}{D}, -\frac{mr}{D}\right)^{\top}.$$

By the delta method,

$$\sqrt{n}\left\{\frac{N_n}{D_n} - \frac{N}{D}\right\} \xrightarrow{d} N\left(0, \sigma_G^2\right),$$

where

$$\sigma_G^2 = \nabla\phi(N,D)^{\top}\Sigma_G \nabla\phi(N,D).$$

This variance can be written equivalently as

$$\sigma_G^2 = \frac{1}{D^2}\operatorname{Var}\{H_G(T) - mrQ_G(T)\}.$$

Since $mr_n = N_n / D_n$, and $mr = N / D$, the desired result follows. □

Theorem 2.2 shows that the proposed estimator is asymptotically Gaussian after standard normalization. Consequently, its sampling uncertainty can be quantified analytically, providing the theoretical basis for confidence intervals and large-sample statistical inference.

Remark 2.2. *Theorem 2.2 provides a basis for large-sample inference on the latent random inspection mean residual life. In particular, it permits the construction of asymptotic standard errors and confidence intervals without relying exclusively on resampling methods. This elevates the proposed methodology from a consistent estimator to a practically useful inferential procedure through analytically justified uncertainty quantification. A detailed derivation of the asymptotic variance and its plug-in estimator is provided in Appendix A.*

The asymptotic normality result in Theorem 2.2 provides a direct basis for estimating the sampling variability of mr_n in the complete-data case. Let

$$D_n = \frac{1}{n}\sum_{i=1}^{n} Q_G(T_i),$$

where

$$Q_G(T_i) = G(T_i -), \qquad H_G(T_i) = \int_0^{\infty} (T_i - x)_+ \, dG(x).$$

A natural plug-in estimator of the asymptotic variance σ_G^2 is

$$\hat{\sigma}_G^2 = \frac{1}{D_n^2}\left[\frac{1}{n}\sum_{i=1}^{n}\left\{H_G(T_i) - mr_n Q_G(T_i)\right\}^2\right].$$

Hence, an approximate $100(1-\alpha)\%$ confidence interval for m_r is given by

$$mr_n \pm z_{1-\alpha/2}\frac{\hat{\sigma}_G}{\sqrt{n}},$$

where $z_{1-\alpha/2}$ denotes the $(1-\alpha/2)$-quantile of the standard normal distribution. This interval provides an analytically justified large-sample inference procedure for the latent random-inspection mean residual life and complements the strong consistency result by quantifying the estimator's sampling uncertainty.

In the exponential inspection model considered next, $Q_G(t)$ and $H_G(t)$ admit closed-form expressions, so the estimator, its estimated asymptotic variance, and the resulting confidence interval are all directly computable from the observed complete lifetime sample.

2.2. Exponential latent inspection model

A particularly important and practically relevant special case arises when the latent inspection age follows an exponential distribution with rate $\lambda > 0$:

$$G(x) = 1 - e^{-\lambda x}, x \geq 0.$$

This specification corresponds to a Poisson-type monitoring mechanism. Larger values of λ represent more frequent inspections, whereas smaller values correspond to less frequent inspections. In this formulation, λ is a monitoring-intensity parameter rather than a parameter of the lifetime distribution. Changing λ therefore modifies the assumed inspection mechanism, while the lifetime distribution continues to be estimated nonparametrically from the observed data. In applications, the value of λ should be interpreted through the expected inspection age, since $E(X)=1/\lambda$ under the exponential inspection model. When information about the monitoring policy is available, λ may be chosen to match the average inspection frequency. For example, if units are typically inspected after an average time α, then $\lambda=1/\alpha$ provides a natural baseline choice. When such information is unavailable, a sensitivity analysis over a plausible range of λ-values is recommended. This enables the analyst to examine how the estimated latent random-inspection mean residual life changes under different monitoring intensities, without imposing a single unverifiable inspection mechanism.

In practical applications, the inspection-age distribution should be chosen according to the characteristics of the monitoring or inspection process. When inspections occur approximately at random with a constant inspection rate, the exponential distribution provides a natural and convenient model because of its memoryless property. If historical inspection records are available, the parameters of the inspection-age distribution may be estimated directly from these data.

Although the exponential inspection model is adopted for computational convenience and interpretability, the general semi-parametric estimator introduced in Section 2.1 is not restricted to this specification. For an alternative inspection-age distribution G_η, indexed by a parameter vector η, the same plug-in principle remains applicable provided that the denominator is strictly positive and the corresponding integrability and moment conditions are satisfied. The consistency argument therefore extends to other inspection models under analogous regularity conditions. However, the numerical stability of the resulting estimator may depend on the tail behavior and concentration of G_η because these features determine the weights assigned to early and late portions of the estimated survival curve. Heavy-tailed inspection distributions may place greater weight on regions where survival estimation is less precise, particularly under censoring. Accordingly, when the inspection mechanism is uncertain, the estimator should be evaluated under several scientifically plausible distributions, such as Weibull, gamma, or lognormal inspection models, and the resulting estimates should be compared as a sensitivity analysis. Stability across these specifications provides evidence that the substantive conclusions are not driven by the assumed inspection model, whereas substantial variation indicates sensitivity that should be reported explicitly.

Since $dG(x)=\lambda e^{-\lambda x}dx$, the target functional in (2.1) becomes

$$mr(\lambda)=\frac{\int_0^\infty\int_0^\infty \lambda e^{-\lambda x}\bar{F}(x+y)\,dydx}{\int_0^\infty \lambda e^{-\lambda x}\bar{F}(x)dx}. \tag{2.5}$$

The corresponding estimator is

$$mr_n(\lambda)=\frac{\int_0^\infty\int_0^\infty \lambda e^{-\lambda x}\bar{F}_n(x+y)\,dydx}{\int_0^\infty \lambda e^{-\lambda x}\bar{F}_n(x)dx}. \tag{2.6}$$

Let $0=T_{(0)}<T_{(1)}\le\cdots\le T_{(n)}$ denote the ordered lifetime observations. Since $\bar{F}_n(t)=0$ for $t\ge T_{(n)}$, only the intervals $\left[T_j,T_{j+1}\right), j=0,\ldots,n-1$ contribute to the integrals. On each such interval,

$$\bar{F}_n(t) = 1 - \frac{j}{n}, \qquad T_{(j)} \le t < T_{(j+1)}.$$

Equivalently,

$$\bar{F}_n(t) = \sum_{j=0}^{n-1}\left(1 - \frac{j}{n}\right) I\left\{T_{(j)} \le t < T_{(j+1)}\right\}. \tag{2.7}$$

To see the origin of the finite-sum form, note that by the change of variable $u = x + y$, the numerator in (2.6) can be written as

$$\int_0^\infty \hat{F}_n(u)\left\{1 - e^{-\lambda u}\right\} du.$$

Thus, on each interval $\left[T_j, T_{j+1}\right)$, the empirical survival level $1 - j/n$ is multiplied by the integrated exponential-inspection weight $\left\{1 - e^{-\lambda u}\right\}$. Substituting (2.7) into (2.6) yields the finite-sum representation

$$mr_n(\lambda) = \frac{\sum_{j=0}^{n-1}\left(1 - \frac{j}{n}\right)\left[\left(T_{(j+1)} - T_{(j)}\right) + \frac{1}{\lambda}\left(e^{-\lambda T_{(j+1)}} - e^{-\lambda T_{(j)}}\right)\right]}{\sum_{j=0}^{n-1}\left(1 - \frac{j}{n}\right)\left(e^{-\lambda T_{(j)}} - e^{-\lambda T_{(j+1)}}\right)}. \tag{2.8}$$

From a computational perspective, after the lifetime observations have been sorted, the finite-sum representation in (2.8) requires only $O(n)$ operations. Thus, the proposed estimator is not only analytically explicit but also computationally direct and scalable for applied reliability and survival studies. The numerator is a weighted empirical residual-life integral, while the denominator is the empirical probability of survival beyond the latent exponential inspection age. This representation avoids numerical double integration and clarifies how the inspection intensity λ changes the contribution of early and late observed lifetimes.

The next result links the proposed random-inspection functional to the classical mean residual life function.

Lemma 2.1. Under the exponential inspection model,

$$mr(\lambda) = \frac{\int_0^\infty m(x)\bar{F}(x)\lambda e^{-\lambda x} dx}{\int_0^\infty \bar{F}(x)\lambda e^{-\lambda x} dx}. \tag{2.9}$$

Proof. The numerator in (2.5) can be written as

$$\int_0^\infty \lambda e^{-\lambda x}\left\{\int_0^\infty \bar{F}(x+y)dy\right\} dx.$$

Using the change of variable $u = x + y$,

$$\int_0^\infty \bar{F}(x+y)\,dy = \int_x^\infty \bar{F}(u)\,du = m(x)\,\bar{F}(x).$$

Substitution into (2.5) gives (2.9), which completes the proof. □

Lemma 2.1 gives a clear interpretation of $mr(\lambda)$: It is a weighted average of the classical mean residual life function $m(x)$, with weights proportional to $\bar{F}(x)\lambda e^{-\lambda x}$. Thus, the proposed functional preserves the meaning of classical residual life analysis while incorporating uncertainty in the observation age. Increasing λ shifts the weight toward earlier ages, whereas smaller values of λ enable later ages to have greater influence.

Proposition 2.1. If $m(x)$ is decreasing in x, then

$$mr(\lambda) \le m(0) = E(T). \tag{2.10}$$

If $m(x)$ is increasing in x, then

$$mr(\lambda) \ge m(0) = E(T). \tag{2.11}$$

Proof. By Lemma 2.1, $mr(\lambda)$ is a weighted average of $m(x)$ with nonnegative weights proportional to $\bar{F}(x)\lambda e^{-\lambda x}$. If $m(x)$ is decreasing, then $m(x) \le m(0)$ for all $x \ge 0$. Hence,

$$\int_0^\infty m(x)\bar{F}(x)\lambda e^{-\lambda x}dx \le m(0)\int_0^\infty \bar{F}(x)\lambda e^{-\lambda x}dx.$$

Dividing by the positive denominator gives (2.10), and the increasing case follows similarly. □

Corollary 2.1. If T has an increasing failure rate, then

$$mr(\lambda) \le E(T). \tag{2.12}$$

Proof. The increasing failure rate property implies decreasing mean residual life. The result follows immediately from Proposition 2.1. □

This corollary is consistent with the physical interpretation of aging systems. Under progressive degradation, the expected remaining lifetime after a random inspection is bounded above by the initial mean lifetime. Hence, the proposed functional behaves coherently with standard aging concepts in reliability theory. Thus, the proposed latent random-inspection functional preserves one of the most fundamental qualitative implications of increasing-failure-rate aging.

2.3. Right-censored lifetime data

In many survival and reliability studies, complete failure times are not fully observed. Instead, observations may be right-censored because the study terminates before failure, the unit is withdrawn, or the monitoring period ends. Let

$$Y_i = \min(T_i, C_i), \qquad \delta_i = I(T_i \le C_i), \qquad i = 1, \ldots, n,$$

where C_i is the censoring time and δ_i indicates whether the failure time is observed. Independent censoring is assumed, so that C_i is independent of T_i. This assumption ensures that censoring is non-informative with respect to the lifetime process. Let $t_1^o < t_2^o < \cdots < t_k^o$ be the distinct observed failure times. At t_i^o, let d_i denote the number of failures and r_i the number at risk immediately before t_i^o. The Kaplan-Meier estimator of the survival function is

$$\hat{\bar{F}}_{KM}(t) = \prod_{t_i^o \le t} \left(1 - \frac{d_i}{r_i}\right). \tag{2.13}$$

The censored-data estimator of $mr(\lambda)$ is obtained by replacing $\bar{F}$ in (2.6) with $\hat{\bar{F}}_{KM}$:

$$\widehat{mr}_{KM}(\lambda) = \frac{\int_0^\infty \int_0^\infty \lambda e^{-\lambda x} \hat{\bar{F}}_{KM}(x+y)\,dy\,dx}{\int_0^\infty \lambda e^{-\lambda x} \hat{\bar{F}}_{KM}(x)\,dx}. \tag{2.14}$$

A practical issue arises because the Kaplan-Meier estimator may have a nonzero tail beyond the largest observed failure time if the largest observed time is censored. To avoid extrapolating beyond the identifiable support, one may either assume that the largest observed time is a failure, or define the estimator on a finite analysis horizon τ within the identifiable support. For clarity, the formulas below are written on a finite horizon τ; when the upper tail is fully identifiable, τ may be taken as the terminal support point.

Set $t_0^o = 0$ and $t_{k+1}^o = \tau$. Define

$$S_i = \hat{\bar{F}}_{KM}\left(t_i^o\right),\ i = 0, \dots, k,$$

with $S_0 = 1$. Since the Kaplan-Meier estimator is a step function, it can be written on $[0, \tau]$ as

$$\hat{\bar{F}}_{KM}(t) = \sum_{i=0}^{k} S_i I\left\{t_i^o \le t < t_{i+1}^o\right\}. \tag{2.15}$$

As in the complete-data case, the step-function structure of the survival estimator converts the double integral into a finite sum over adjacent time intervals. The only difference is that the empirical survival levels are replaced by the Kaplan-Meier levels S_i, which account for censoring through the risk sets and observed failures. Substituting (2.15) into (2.14) yields

$$\widehat{mr}_{KM}(\lambda;\tau) = \frac{\sum_{i=0}^{k} S_i \left[\left(t_{i+1}^o - t_i^o\right) + \frac{1}{\lambda}\left(e^{-\lambda t_{i+1}^o} - e^{-\lambda t_i^o}\right)\right]}{\sum_{i=0}^{k} S_i \left(e^{-\lambda t_i^o} - e^{-\lambda t_{i+1}^o}\right)}. \tag{2.16}$$

The role of τ is thus not merely computational; it defines the portion of the lifetime support over which the survival function is estimable from the censored data without imposing additional tail assumptions. In applications, τ may be chosen as the largest observed failure time or another prespecified time point within the reliable follow-up range. Formula (2.16) is the censored-data analogue of (2.8). Thus, the censored-data estimator retains the same computational structure as the complete-data estimator while avoiding extrapolation beyond the identifiable follow-up horizon. This is particularly important for residual-life functionals, which are sensitive to the upper tail of the survival distribution.

The truncation horizon τ is introduced to restrict inference to the identifiable region of the survival distribution under right censoring, where the Kaplan-Meier estimator remains statistically reliable. Since the tail of the survival curve is increasingly supported by fewer individuals at risk, estimation near the largest observed follow-up time may become unstable and highly variable. Accordingly, τ should be chosen within the identifiable support of the Kaplan-Meier estimator, where a sufficient number of uncensored observations remain to ensure stable estimation and meaningful statistical inference.

From a practical perspective, τ represents a scientifically interpretable analysis horizon rather than an arbitrary truncation point. A natural default choice is the largest observed uncensored failure time, beyond which the Kaplan-Meier estimator is typically unreliable because of extensive censoring and limited information. Alternatively, when inference is intended for a prespecified clinical follow-up period or an operational maintenance horizon, τ may be selected to reflect that application-specific objective. Consequently, the proposed estimator should be interpreted as the mean residual life over the identifiable observation window $[0, \tau]$, providing a principled balance between estimation stability, statistical efficiency, and practical relevance while avoiding unreliable extrapolation into poorly supported regions of the survival distribution.

Corollary 2.2. Assume independent right censoring and the standard support conditions under which the Kaplan-Meier estimator is uniformly consistent on a fixed finite analysis horizon $[0, \tau]$. Let $\lambda > 0$ be fixed, and suppose that

$$D_\tau(\lambda) = \int_0^\tau \lambda e^{-\lambda x} \overline{F}(x) dx > 0 .$$

Then

$$\widehat{mr}_{KM}\left(\lambda;\tau\right) \to mr\left(\lambda;\tau\right) \text{ almost surely as } n \to \infty . \tag{2.17}$$

If τ covers the full identifiable support, or if the tail contribution beyond τ is negligible or otherwise identifiable, then the truncated functional $mr\,(\lambda; \tau)$ agrees with the unrestricted target functional $mr\,(\lambda)$.

Proof. For a fixed finite horizon τ, define the truncated population numerator and denominator by

$$N_\tau\left(\lambda\right) = \int_0^\tau \int_0^{\tau-x} \lambda e^{-\lambda x} \overline{F}\left(x+y\right) dydx,$$
$$D_\tau\left(\lambda\right) = \int_0^\tau \lambda e^{-\lambda x} \overline{F}\left(x\right) dx.$$

Thus,

$$mr\left(\lambda;\tau\right) = \frac{N_\tau\left(\lambda\right)}{D_\tau\left(\lambda\right)},$$

provided that $D_\tau(\lambda) > 0$. The Kaplan-Meier plug-in estimator is similarly written as

$$\widehat{mr}_{KM}(\lambda;\tau) = \frac{\hat{N}_\tau(\lambda)}{\hat{D}_\tau(\lambda)},$$

where

$$\hat{N}_\tau(\lambda) = \int_0^\tau \int_0^{\tau - x} \lambda e^{-\lambda x} \hat{\bar{F}}_{KM}(x+y)\,dydx,$$
$$\hat{D}_\tau(\lambda) = \int_0^\tau \lambda e^{-\lambda x} \hat{\bar{F}}_{KM}(x)\,dx.$$

By the uniform consistency of the Kaplan-Meier estimator on [0, τ],

$$\sup_{0\le t\le\tau} \left| \hat{\bar{F}}_{KM}(t) - \hat{\bar{F}}(t) \right| \to 0 \quad \text{almost surely.}$$

For the denominator, it follows that

$$\begin{aligned} \left| \hat{D}_\tau(\lambda) - D_\tau(\lambda) \right| &\le \int_0^\tau \lambda e^{-\lambda x} \left| \hat{\bar{F}}_{KM}(x) - \bar{F}(x) \right| dx \\ &\le \left(\int_0^\tau \lambda e^{-\lambda x} dx \right) \sup_{0\le t\le\tau} \left| \hat{\bar{F}}_{KM}(t) - \bar{F}(t) \right|, \end{aligned}$$

and therefore

$$\hat{D}_\tau(\lambda) \to D_\tau(\lambda) \quad \text{almost surely.}$$

Similarly, for the numerator,

$$\left| \hat{N}_\tau(\lambda) - N_\tau(\lambda) \right| \le \int_0^\tau \int_0^{\tau - x} \lambda e^{-\lambda x} \left| \hat{\bar{F}}_{KM}(x+y) - \bar{F}(x+y) \right| dydx.$$

Since $0 \le x + y \le \tau$ over the domain of integration,

$$\left| \hat{N}_\tau(\lambda) - N_\tau(\lambda) \right| \le \left[\int_0^\tau \int_0^{\tau - x} \lambda e^{-\lambda x} dydx \right] \sup_{0\le t\le\tau} \left| \hat{\bar{F}}_{KM}(t) - \bar{F}(t) \right|.$$

The double integral is finite for every fixed $\tau < \infty$. Hence,

$$\hat{N}_\tau(\lambda) \to N_\tau(\lambda) \quad \text{almost surely.}$$

Since $D_\tau(\lambda) > 0$, it follows that $\hat{D}_\tau(\lambda) > 0$ eventually with probability one. Therefore, by the continuous mapping theorem,

$$\widehat{mr}_{KM}(\lambda;\tau) = \frac{\hat{N}_\tau(\lambda)}{\hat{D}_\tau(\lambda)} \to \frac{N_\tau(\lambda)}{D_\tau(\lambda)} = mr(\lambda;\tau) \quad \text{almost surely.}$$

If the analysis horizon τ covers the full identifiable support, or if the remaining tail contribution beyond τ is negligible or otherwise identifiable, then the truncated functional mr (λ; τ) agrees with the unrestricted functional $mr(\lambda)$.

This completes the proof. □

Remark 2.3. *The use of a finite analysis horizon τ is a standard precaution in censored residual-life analysis. It prevents the estimator from relying on unobserved tail behavior when the largest observation is censored and makes the target functional explicitly tied to the identifiable portion of the survival curve. This precaution is particularly important for residual-life functionals, which are sensitive to the upper tail of the lifetime distribution.*

Remark 2.4. *The censored-data result establishes strong consistency of the Kaplan-Meier-based estimator on a fixed finite analysis horizon. A full asymptotic normality theory for $\widehat{mr}_{KM}(\lambda;\tau)$ is technically more involved than the complete-data case, because the estimator is a smooth functional of the Kaplan-Meier survival process rather than a ratio of ordinary empirical averages. Under standard independent censoring and support conditions, such a result can, in principle, be obtained by combining the weak convergence theory of the product-limit estimator with a functional delta-method argument applied to the numerator and denominator functionals. Classical weak convergence results for the Kaplan-Meier estimator, including Gill's large-sample theory for the product-limit estimator, provide the appropriate technical foundation for this extension. Since the main focus in this paper is the semi-parametric latent-inspection construction and its direct complete-data inferential theory, the detailed derivation of asymptotic normality and variance estimation for the censored-data estimator is left for future work. The bootstrap confidence intervals reported in the empirical studies should therefore be interpreted as practical finite-sample uncertainty measures rather than as intervals justified by the new asymptotic theory developed in this paper. Their use relies on the well-established bootstrap methodology for Kaplan-Meier-based estimators under independent censoring, which has been extensively investigated in the survival-analysis literature. Since the proposed censored-data estimator is a smooth functional of the Kaplan-Meier estimator over a fixed identifiable horizon, the bootstrap provides a natural computational tool for approximating its sampling distribution in finite samples. A rigorous theoretical investigation of bootstrap validity specifically for the proposed latent random-inspection estimator remains an interesting topic for future research.*

The proposed estimator is computationally inexpensive because it involves only empirical (or Kaplan-Meier) survival estimation followed by finite-sum calculations. No iterative optimization or simulation-based procedures are required, making the computational cost negligible for the sample sizes considered in this study.

2.4. Computational implementation

The proposed estimator is computationally straightforward because it does not require iterative optimization, numerical maximization, or any gradient-based algorithm. Once the survival function has been estimated, using the empirical estimator for complete data or the Kaplan-Meier estimator for right-censored observations, the proposed random-inspection mean residual life estimator is obtained by direct evaluation of the analytical expressions developed in Section 2.

For complete observations, the estimator is computed through explicit finite-sum expressions, whereas, for right-censored data, the Kaplan-Meier estimator is substituted into the corresponding functional over the identifiable finite analysis horizon. Consequently, the computational implementation relies only on standard survival-analysis routines that are available in most statistical software packages.

Because no iterative optimization is involved, the proposed methodology is numerically stable and free from optimization-related convergence issues. The theoretical convergence of the estimator is established by the consistency and asymptotic normality results proved in Theorems 2.1 and 2.2, which provide the statistical foundation for subsequent inference.

Algorithm 1. Computational implementation of the proposed estimator.

Input: Observed lifetime data (and censoring indicators, if applicable), together with a specified latent inspection-age distribution G and its associated parameter(s).

Step 1: Estimate the survival function using the empirical estimator for complete data or the Kaplan-Meier estimator for right-censored data.

Step 2: Substitute the estimated survival function into the proposed random-inspection mean residual life estimator defined in Section 2.

Step 3: Evaluate the corresponding finite summation or numerical integral according to the analytical expression of the estimator.

Step 4: Estimate the asymptotic variance using the theoretical results established in Section 2.

Step 5: Construct pointwise confidence intervals (or bootstrap confidence intervals when appropriate).

Step 6: When the inspection mechanism is uncertain, if required, repeat the above procedure under several scientifically plausible inspection-age distributions as part of the recommended sensitivity analysis.

Output: Estimated random-inspection mean residual life together with its associated uncertainty measures.

The above computational workflow provides a transparent and reproducible implementation of the proposed methodology and can be readily incorporated into standard statistical software for complete and right-censored lifetime data.

3. Simulation study and numerical investigation

In this section, a simulation study is presented to evaluate the finite-sample performance of the proposed semi-parametric estimator under latent random inspection. The main objective is to examine how accurately the estimator recovers the mean residual life functional when the lifetime distribution, inspection intensity, sample size, and censoring proportion vary. The simulation study also provides numerical support for the consistency results established in Section 2.

The design reflects survival and reliability settings in which the lifetime process is observed either completely or under right censoring, whereas the inspection age is not directly observed and is represented through a stochastic monitoring model. In this framework, the inspection intensity parameter λ controls the distribution of the latent inspection age. Larger values of λ correspond to more frequent inspections and place greater weight on earlier system ages, while smaller values represent less frequent monitoring and enable later ages to contribute more strongly to the residual life functional.

The performance of the estimator is assessed using empirical bias and mean squared error (MSE). These two measures are reported across lifetime distributions, aging regimes, inspection intensities, sample sizes, and censoring proportions.

3.1. Simulation design

Let T denote the lifetime variable and let the latent inspection age X follow an exponential distribution with rate $\lambda > 0$. For each configuration, the target quantity is

$$mr(\lambda) = E(T - X \mid T > X),$$

and the corresponding semi-parametric estimator is computed using the complete-data or censored-data formula developed in Section 2. For each simulation configuration, the true value of *mr* (λ) is computed from the population form

$$mr(\lambda) = \frac{\int_0^\infty \int_0^\infty \lambda e^{-\lambda x} \bar{F}(x+y)\, dydx}{\int_0^\infty \lambda e^{-\lambda x} \bar{F}(x)\, dx},$$

where $\bar{F}$ is the survival function of the specified Weibull or gamma lifetime model. For the right-censored experiments, the analysis is conducted over the same finite horizon τ used in the theoretical development of Section 2. Accordingly, the reference value for each censored configuration is the corresponding truncated population functional *mr* (λ, *τ*), rather than the unrestricted functional *mr* (λ). The Kaplan-Meier estimator and its population target are therefore evaluated over the common identifiable interval [0, *τ*]. This common-target construction ensures that the reported bias and mean squared error measure estimation performance under censoring without being confounded by truncation-induced differences between *mr* (λ, *τ*) and *mr* (λ).

For complete-data configurations, the unrestricted population functional *mr* (λ) is evaluated numerically with high accuracy. For right-censored configurations, the corresponding truncated population functional *mr* (λ, *τ*) is evaluated over the same fixed horizon [0, *τ*] used by the Kaplan-Meier estimator. The empirical bias and mean squared error are then calculated over the Monte Carlo replications as

$$Bias = \frac{1}{B}\sum_{b=1}^{B}\left\{\widehat{mr}^{(b)}(\lambda) - mr(\lambda)\right\},$$

$$MSE = \frac{1}{B}\sum_{b=1}^{B}\left\{\widehat{mr}^{(b)}(\lambda) - mr(\lambda)\right\}^2,$$

where B=1000 denotes the number of Monte Carlo replications, and $\widehat{mr}^{(b)}(\lambda)$ is the estimate obtained in the *b*-th replication.

To further demonstrate the practical relevance of modeling the inspection age as a latent random variable, a simple fixed-age benchmark is included. Under the exponential inspection model, the mean inspection age is $E(X) = 1/\lambda$. A natural naive alternative is therefore to replace the random inspection age X by its mean and evaluate the classical mean residual life at this deterministic age, namely $1/\lambda$. This benchmark reflects a common practical simplification in which the monitoring process is summarized by a single representative inspection time.

It should be stressed, however, that $mr(\lambda) = E(T - X \mid T > X)$ and $m(1/\lambda)$ are fundamentally different residual-life functionals. The proposed random-inspection functional incorporates the full distribution of the latent inspection age and weights residual-life behavior according to survival beyond that random age. By contrast, the fixed-age benchmark uses only one deterministic inspection point and accordingly ignores the variability of the monitoring process. The comparison is therefore not intended to regard $m(1/\lambda)$ as an estimator of the same target parameter, but rather to assess the potential impact of replacing a stochastic inspection mechanism by its mean. This provides a clear benchmark for evaluating the added value of the proposed latent random-inspection formulation.

Two lifetime models are considered. The first is the Weibull distribution, a standard model in reliability analysis due to its ability to represent distinct aging behaviors through its shape parameter. When the shape parameter is less than one, the Weibull distribution exhibits a decreasing failure rate, which is often associated with early-life failures or burn-in phenomena. When the shape parameter exceeds one, it exhibits an increasing failure rate, reflecting progressive degradation and aging. The second model is the gamma distribution, which is included as a flexible alternative capable of representing heterogeneous lifetime behavior and more dispersed failure patterns.

For each model, independent samples of sizes n = 30, 50, and 100 are generated. The inspection intensity is varied over selected values of λ, and right censoring is introduced to assess the estimator under incomplete observation. To generate right-censored data, censoring times C_i are sampled independently from a uniform distribution on (0, c_u). The upper endpoint c_u is chosen to achieve a desired censoring proportion p, determined by

$$\int_0^{c_u} \overline{F}(x)dx = c_u p \, .$$

This calibration produces approximately 100p% right-censored observations. The censoring proportions considered are p = 0%, 10%, and 25%. For each simulation setting, the following procedure is used:

(1) Generate an independent sample $T_1,\ldots,T_n$ from the selected lifetime distribution.

(2) Specify the latent inspection mechanism through the exponential rate λ.

(3) For censored settings, generate independent censoring times $C_1,\ldots,C_n$, and observe

$$Y_i = \min\left(T_i, C_i\right), \ \delta_i = I\left(T_i \le C_i\right),$$

where δ_i is the failure indicator.

(4) Compute the proposed estimator using the empirical survival function for complete data or the Kaplan-Meier survival estimator for right-censored data.

(5) Repeat the procedure over 1000 independent Monte Carlo replications.

(6) Compute the empirical bias and MSE relative to the corresponding population target: *mr* (λ) for complete-data configurations and *mr* (λ, τ) for right-censored configurations.

This design enables a systematic assessment of how the estimator behaves as information increases through larger sample sizes, how it responds to incomplete data through censoring, and how the latent inspection intensity affects the resulting residual life estimates.

3.2. Results for the Weibull model

Table 2 shows that MSE generally decreases with sample size, while heavier censoring increases bias and variability.

Table 2. Unified simulation results for the Weibull lifetime model. Each numerical entry reports (Bias, MSE) for the indicated censoring proportion, parameter configuration, inspection intensity λ, and sample size.

Censoring	(θ, k)	λ	$n = 30$	$n = 50$	$n = 100$
0%	1.2, 0.75	0.5	(-0.0089, 0.0389)	(-0.0065, 0.0243)	(-0.0059, 0.0109)
0%	0.8, 1.5	0.5	(-0.0245, 0.0854)	(-0.0230, 0.0472)	(-0.0034, 0.0236)
0%	1.5, 2	0.5	(-0.0023, 0.0024)	(-0.0004, 0.0013)	(-0.0013, 0.0006)
10%	1.2, 0.75	0.5	(0.0085, 0.0438)	(0.0160, 0.0263)	(0.0218, 0.0148)
10%	0.8, 1.5	0.5	(0.0467, 0.1223)	(0.0454, 0.0952)	(0.0558, 0.0471)
10%	1.5, 2	0.5	(0.0015, 0.0024)	(0.0019, 0.0015)	(0.0042, 0.0008)
25%	1.2, 0.75	0.5	(0.0996, 0.0993)	(0.1085, 0.0697)	(0.0990, 0.0484)
25%	0.8, 1.5	0.5	(0.3051, 0.5667)	(0.3997, 0.5571)	(0.4751, 0.6116)
25%	1.5, 2	0.5	(0.0165, 0.0038)	(0.0196, 0.0031)	(0.0138, 0.0016)
0%	1.2, 0.75	1.0	(-0.0114, 0.0398)	(-0.0026, 0.0232)	(-0.0087, 0.0116)
0%	0.8, 1.5	1.0	(-0.0162, 0.0634)	(-0.0146, 0.0401)	(-0.0077, 0.0184)
0%	1.5, 2	1.0	(-0.0035, 0.0023)	(-0.0013, 0.0014)	(-0.0006, 0.0007)
10%	1.2, 0.75	1.0	(0.0186, 0.0455)	(0.0176, 0.0255)	(0.0196, 0.0137)
10%	0.8, 1.5	1.0	(0.0538, 0.1159)	(0.0393, 0.0612)	(0.0516, 0.0354)
10%	1.5, 2	1.0	(0.0015, 0.0024)	(0.0037, 0.0017)	(0.0028, 0.0008)
25%	1.2, 0.75	1.0	(0.0968, 0.0906)	(0.0926, 0.0610)	(0.0858, 0.0376)
25%	0.8, 1.5	1.0	(0.2991, 0.4452)	(0.3000, 0.3576)	(0.3479, 0.3169)
25%	1.5, 2	1.0	(0.0202, 0.0047)	(0.0171, 0.0028)	(0.0159, 0.0016)

3.3. Results for the gamma model

Table 3 confirms the same pattern for gamma lifetimes, with the greatest deterioration occurring under highly dispersed distributions and 25% censoring. To provide numerical support for the fixed-age benchmark discussion, Table 4 compares the population random-inspection mean residual life with the classical mean residual life evaluated at the mean inspection age, 1/λ.

Table 3. Unified simulation results for the gamma lifetime model. Each numerical entry reports (Bias, MSE) for the indicated censoring proportion, parameter configuration, inspection intensity λ, and sample size.

Censoring	(a, b)	λ	$n = 30$	$n = 50$	$n = 100$
0%	1.2, 0.3	0.5	(-0.0552, 0.4613)	(-0.0088, 0.2961)	(0.0003, 0.1363)
0%	0.8, 0.5	0.5	(-0.0470, 0.1588)	(-0.0357, 0.0979)	(-0.0093, 0.0549)
0%	2, 1	0.5	(-0.0146, 0.0566)	(-0.0056, 0.0352)	(0.0028, 0.0184)
10%	1.2, 0.3	0.5	(0.0654, 0.5612)	(0.0884, 0.3697)	(0.0508, 0.1972)
10%	0.8, 0.5	0.5	(0.0930, 0.2653)	(0.0703, 0.1554)	(0.0555, 0.0804)
10%	2, 1	0.5	(0.0395, 0.0698)	(0.0197, 0.0435)	(0.0215, 0.0209)
25%	1.2, 0.3	0.5	(0.4132, 1.3019)	(0.4180, 0.9608)	(0.3529, 0.5225)
25%	0.8, 0.5	0.5	(0.3342, 0.7973)	(0.7309, 0.6369)	(0.4267, 0.5208)
25%	2, 1	0.5	(0.1276, 0.1448)	(0.0968, 0.0866)	(0.0871, 0.0514)
0%	1.2, 0.3	1.0	(-0.0105, 0.4741)	(0.0161, 0.2639)	(-0.0163, 0.1439)
0%	0.8, 0.5	1.0	(-0.0176, 0.1513)	(-0.0085, 0.0892)	(0.0057, 0.0468)
0%	2, 1	1.0	(-0.0059, 0.0548)	(-0.0005, 0.0330)	(-0.0001, 0.0178)
10%	1.2, 0.3	1.0	(0.0757, 0.5785)	(0.0567, 0.3179)	(0.0439, 0.1683)
10%	0.8, 0.5	1.0	(0.0455, 0.2252)	(0.0627, 0.1324)	(0.0584, 0.0663)
10%	2, 1	1.0	(0.0281, 0.0703)	(0.0181, 0.0434)	(0.0181, 0.0207)
25%	1.2, 0.3	1.0	(0.3823, 1.0942)	(0.3203, 0.7026)	(0.2924, 0.4193)
25%	0.8, 0.5	1.0	(0.3145, 0.6815)	(0.3503, 0.5106)	(0.3481, 0.3961)
25%	2, 1	1.0	(0.1005, 0.1171)	(0.0864, 0.0766)	(0.0714, 0.0370)

Table 4 shows that replacing the random inspection age by its mean may produce substantial relative discrepancies, especially under pronounced aging or heterogeneous lifetime behavior.

Table 4. Comparison between the random-inspection MRL functional and the fixed-age benchmark.

Lifetime model	Parameters	λ	$mr(\lambda)$	$m(1/\lambda)$	Relative difference
Weibull	(θ, k) = (1.2, 0.75)	0.5	1.829	2.130	-14.13%
Weibull	(θ, k) = (0.8, 1.5)	0.5	0.547	0.315	73.46%
Weibull	(θ, k) = (1.5, 2)	1.0	0.973	0.717	35.67%
Gamma	(a, b) = (1.2, 0.3)	0.5	3.782	3.685	2.63%
Gamma	(a, b) = (0.8, 0.5)	1.0	1.736	1.795	-3.27%
Gamma	(a, b) = (2, 1)	1.0	1.667	1.500	11.11%

3.4. Discussion and interpretation

The simulation results provide numerical evidence that the proposed estimator behaves coherently across the considered lifetime models, inspection intensities, sample sizes, and censoring proportions. In most configurations, the MSE decreases as the sample size increases from n=30 to n=100. This pattern is consistent with the consistency result established in Section 2. Although the decrease is not perfectly monotone in every individual setting, the overall behavior indicates improved estimation stability as more lifetime information becomes available. The results also show the expected effect of censoring. Under complete data, the estimator typically exhibits small bias and relatively low MSE. With 10% censoring, the loss of accuracy is generally moderate, particularly for larger sample sizes. Under 25% censoring, bias and MSE increase more noticeably in several parameter settings. This behavior is statistically natural because residual-life functionals depend on the upper part of the lifetime distribution, and right censoring reduces the amount of available information in that region.

The lifetime model and its parameter configuration also influence finite-sample performance. For Weibull lifetimes, the estimator tends to be more stable in increasing-failure-rate settings, where the degradation pattern is more regular. In decreasing-failure-rate settings, and in more dispersed gamma configurations, the estimator may show larger variability. This reflects the greater uncertainty associated with early-life failures, heterogeneity, and more variable tail behavior. The inspection intensity λ plays a meaningful role because it determines how strongly different parts of the lifetime distribution contribute to the target functional. Larger values of λ place more weight on earlier inspection ages, whereas smaller values enable later ages to contribute more substantially. The simulation results show that the estimator responds to changes in λ in a stable and interpretable manner, supporting its use as a sensitivity tool when the inspection mechanism is latent or model-based rather than directly observed.

The simulation study supports the practical reliability of the proposed estimator in moderate sample sizes and under moderate censoring. Furthermore, the results highlight the expected limitations of residual-life estimation: Heavier censoring and highly variable lifetime distributions naturally reduce precision. These findings should therefore be interpreted as finite-sample evidence supporting the theoretical consistency of the estimator, rather than as a claim of uniformly superior performance across all censoring proportions and lifetime configurations.

The fixed-age benchmark comparison further confirms the importance of accounting for inspection-time variability. Across several parameter settings, the random-inspection mean residual life $mr(\lambda)$ differs noticeably from the classical mean residual life evaluated at the mean inspection age, $m(1/\lambda)$. This discrepancy is natural because $mr(\lambda)$ depends on the entire inspection-age distribution and on the survival-weighted contribution of different ages, whereas $m(1/\lambda)$ reduces the monitoring process to a single deterministic time point. Consequently, replacing X by $E(X)$ may obscure the effect of latent inspection uncertainty, particularly under heterogeneous lifetime behavior, pronounced aging patterns, or censoring. These findings support the proposed semi-parametric framework as a more faithful inferential tool when inspection ages are unavailable but can be represented through a stochastic monitoring model.

In this simulation study, the primary focus is on point-estimation accuracy through empirical bias and mean squared error, which are directly supported by the theoretical consistency results established in Section 2. For the right-censored estimator, the current theoretical development does not derive an asymptotic distribution or asymptotic variance estimator. Consequently, inferential performance measures such as coverage probabilities and the accuracy of asymptotic variance estimation are intentionally not reported in the present study. A comprehensive evaluation of interval estimation and bootstrap calibration for the censored-data estimator will be the subject of future methodological work.

An additional feature evident from Tables 2 and 3 is the noticeable increase in empirical bias and MSE under heavier censoring for certain parameter configurations. This behavior is expected because right censoring effectively reduces the amount of observable lifetime information, particularly in the upper tail of the survival distribution, where estimation of the mean residual life is most sensitive. As the censoring proportion increases, fewer uncensored observations remain available to estimate the Kaplan-Meier survival function, leading to greater sampling variability and consequently larger estimation errors. The effect is more pronounced for parameter settings associated with heavier-tailed lifetime distributions or greater intrinsic variability, since these settings rely more heavily on observations in regions where censoring is most severe. Nevertheless, the overall simulation results continue to demonstrate the expected improvement in estimation accuracy with increasing sample size, indicating that the proposed estimator remains numerically stable despite the inevitable efficiency loss caused by substantial censoring.

Overall, the simulation results provide consistent empirical support for the proposed semi-parametric estimator. Across the considered lifetime models, inspection intensities, sample sizes, and censoring levels, the estimator exhibits stable finite-sample behavior that is fully consistent with the theoretical properties established in Section 2. These findings provide a sound numerical foundation for the real-data sensitivity analyses presented in the next section.

4. Real-data sensitivity analysis under latent inspection

In this section, the empirical performance and interpretability of the proposed semi-parametric framework are illustrated through three lifetime datasets arising from biomedical survival analysis, classical engineering reliability, and modern degradation-based prognostics. The analyses are intentionally framed as model-based sensitivity studies, rather than validation exercises based on observed inspection ages. In all three datasets, the available information consists of survival times, failure times, or extracted run-to-failure lifetimes, whereas the corresponding inspection ages are not directly recorded. Therefore, the latent inspection process is represented through the exponential monitoring model introduced in Section 2, and the inspection-intensity parameter λ is varied to examine how the estimated mean residual life changes across plausible stochastic monitoring scenarios.

This sensitivity-based interpretation is essential to the proposed methodology. In many retrospective survival, reliability, and predictive-maintenance studies, the lifetime or event-time information is available, but the effective age at which a unit is inspected, assessed, or found to be at risk may be missing, irregularly recorded, indirectly specified, or governed by an external monitoring protocol not included in the dataset. Classical fixed-age mean residual life analysis does not explicitly account for this additional layer of observation uncertainty. The proposed framework addresses this gap by combining a nonparametric estimate of the lifetime distribution with a parametric model for the latent inspection mechanism. In this way, it enables residual-life inference without imposing a fully parametric model on the lifetime distribution.

The three datasets are selected to represent distinct empirical settings and data structures. The first is a right-censored biomedical survival dataset involving kidney catheter infection times, which demonstrates the use of the Kaplan-Meier version of the estimator. The second is a classical mechanical reliability dataset consisting of aircraft air-conditioning failure times, which illustrates the complete-data implementation in a traditional engineering context. The third is extracted from the NASA C-MAPSS turbofan engine degradation repository, a widely used benchmark in prognostics and remaining useful life analysis, and illustrates the applicability of the method to run-to-failure lifetimes obtained from modern degradation trajectories.

The purpose of considering these datasets is not to assume that the same inspection mechanism generates all observations, nor to claim that the latent inspection intensity is directly identifiable from the data alone. Rather, the goal is to demonstrate how the proposed estimator can be used when the lifetime distribution is empirically available but the observation mechanism is unavailable, incomplete, or only indirectly specified. The resulting curves should therefore be interpreted as residual-life sensitivity profiles over λ, showing how conclusions depend jointly on the empirical survival structure and the assumed latent monitoring intensity. This perspective provides a practical and statistically coherent way to assess residual lifetime behavior in settings where inspection uncertainty cannot be ignored but is not directly observed.

4.1. Biomedical survival data: kidney catheter infection times

The first dataset considered is the kidney catheter infection-time dataset available in the survival package in R. The dataset records infection times together with censoring indicators, making it a natural example for demonstrating the right-censored version of the proposed estimator. Since the data contain no directly observed inspection ages, the lifetime survival function is estimated nonparametrically using the Kaplan-Meier estimator, while the inspection-intensity parameter λ is treated as a sensitivity parameter for the latent inspection mechanism. Consequently, the resulting curve should be interpreted as a family of mean residual life estimates corresponding to plausible stochastic monitoring intensities, rather than as an estimate based on observed inspection times.

Table 5 reports descriptive statistics for the infection times. The sample mean is 101.63, while the median is 39.5, with a standard deviation of 130.91 and skewness of 1.42. The substantial separation between the mean and median, together with the positive skewness, indicates a markedly right-skewed distribution. This pattern suggests pronounced heterogeneity in infection-time behavior: Many subjects experience infection relatively early, whereas a smaller subset remains infection-free for considerably longer periods. Such empirical features make the dataset particularly suitable for evaluating how the proposed residual-life functional responds to censoring and to changes in the assumed latent monitoring intensity.

Table 5. Descriptive statistics for the infection times in the kidney catheter dataset.

Min	Max	Mean	Mode	Median	Standard deviation
2	526	101.63	30	39.5	130.91

The left panel in Figure 2 presents the Kaplan-Meier estimate of the survival function, with censored observations marked by plus signs. The curve declines sharply at early times and then develops a long right tail. This behavior is consistent with heterogeneous biomedical survival data, where early infections occur for a noticeable part of the sample, while a subset of subjects exhibits substantially longer infection-free durations. The right panel in Figure 2 shows the estimated mean residual life under latent random inspection as a function of the inspection intensity parameter λ. The range $\lambda \in (0.01, 0.2)$ corresponds to mean latent inspection ages $1/\lambda$ ranging approximately from 100 to 5 time units. Thus, smaller values of λ represent later or less frequent latent inspections, whereas larger values correspond to earlier and more frequent latent inspections.

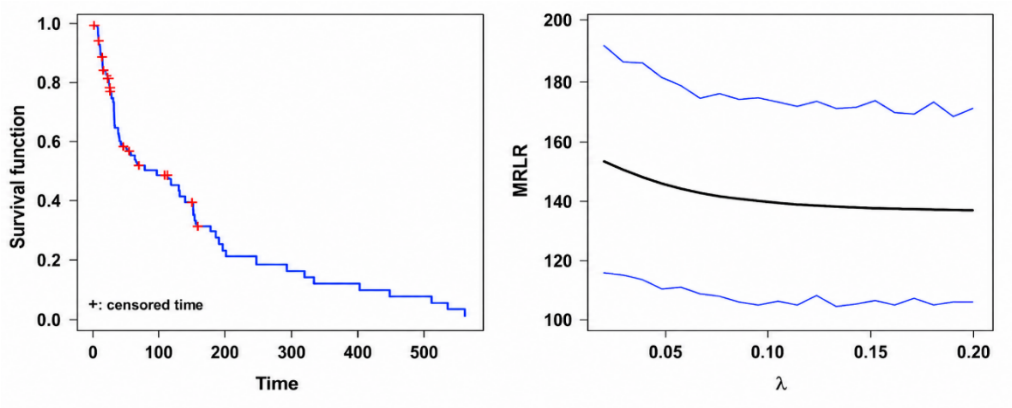

Figure 2. Kidney-catheter analysis. Left: Kaplan-Meier survival estimate with censored observations. Right: Estimated $m_r(\lambda)$ with 95% pointwise bootstrap confidence bands. Main finding: The estimated residual infection-free time decreases as the latent inspection intensity increases.

The estimated curve decreases as λ increases. This indicates that assigning greater weight to earlier inspection ages leads to smaller estimated residual lifetimes for this dataset. This pattern is consistent with the sharp early decline in the Kaplan-Meier curve: When the latent inspection mechanism emphasizes earlier ages, the estimator is more strongly influenced by the region where infection risk is relatively high. The convex shape of the curve further suggests that the effect of increasing λ is stronger at lower inspection intensities and becomes more gradual as the assumed inspection process shifts toward earlier ages. To quantify sampling uncertainty, pointwise bootstrap confidence bands are constructed using B = 1000 resamples. For each bootstrap sample, the Kaplan-Meier survival function is recomputed, and the proposed estimator is evaluated over the same grid of λ values. Approximate 95% confidence bands are then obtained from the bootstrap distribution. The bands provide a visual assessment of uncertainty and indicate that the decreasing pattern is stable across the considered range of latent inspection intensities. From a clinical perspective, this behavior indicates that the expected remaining infection-free time depends not only on the observed survival experience but also on the assumed monitoring pattern. For example, if patients are monitored more frequently during the early post-catheter period, the estimated residual infection time reflects the higher probability of observing early infections. Therefore, the proposed estimator provides a useful tool for evaluating expected remaining infection-free duration under different plausible follow-up strategies.

This example illustrates the practical value of the proposed framework for censored survival data. The analysis does not require observed inspection ages. Rather, it evaluates how residual-life conclusions vary under plausible stochastic monitoring intensities. This is particularly relevant in biomedical applications where event times are recorded but the effective inspection age is not directly available in the dataset.

The kidney catheter analysis illustrates the value of treating the inspection age as a latent random quantity rather than replacing it by a fixed evaluation time. As λ varies, the proposed estimator reweights the contribution of early and late survival ages, thereby revealing the sensitivity of residual-life conclusions to the assumed monitoring regime. From an applied biomedical perspective, the estimated mean residual life can be interpreted as the expected remaining infection-free period under a given stochastic follow-up mechanism. This information may assist researchers and clinicians in understanding the likely duration of catheter survival and in designing appropriate monitoring strategies when exact inspection times are unavailable. Importantly, this analysis is not intended to recover the unobserved inspection process; instead, it provides a principled sensitivity assessment of residual survival under plausible monitoring intensities. In this sense, the proposed framework offers a clinically interpretable extension of classical mean residual life analysis for incomplete-observation settings in which inspection ages are not directly recorded.

The principal value of this application lies in its treatment of right-censored biomedical data. The Kaplan-Meier-based estimator preserves incomplete follow-up information while enabling the residual infection-free period to be examined under alternative latent monitoring intensities. The decreasing $m_r(\lambda)$ profile indicates that the estimated remaining infection-free duration is sensitive to the weighting of early follow-up times, where infection events are more concentrated. Clinically, this finding emphasizes that residual-life conclusions may depend materially on censoring and on the assumed follow-up mechanism. The proposed analysis therefore provides a useful sensitivity assessment for biomedical studies in which event times and censoring indicators are recorded but exact inspection or assessment ages are unavailable.

4.2. Engineering reliability data: aircraft air-conditioning failures

The second application examines a benchmark dataset from engineering reliability, consisting of failure times, measured in operating hours, for aircraft air-conditioning units installed in Boeing 720 aircraft. This dataset is a classical reliability example originally discussed in connection with decreasing failure rate behavior [24] and is consistent with standard applied life-data analysis settings [25]. In contrast to the kidney catheter data, which involve right-censored biomedical survival times, the aircraft dataset consists of complete mechanical failure observations. Therefore, the empirical survival function is used to estimate the lifetime distribution nonparametrically. Because the original data do not report the inspection schedule or monitoring history associated with the observed failures, the exponential inspection model is used to represent a plausible latent observation mechanism. Varying the inspection-intensity parameter λ then enables the analysis to quantify how the estimated residual-life behavior changes across assumed monitoring regimes.

Table 6 presents descriptive statistics for the aircraft failure times. The observations range from 10 to 310 operating hours, with a mean of 83.52, a median of 61, a standard deviation of 70.80, and skewness of 0.95. These summaries indicate a moderately right-skewed distribution with considerable dispersion. Such variability is characteristic of mechanical reliability data, where failure times may be influenced by differences in operating conditions, maintenance history, component quality, thermal stress, vibration, and environmental exposure. This dataset therefore provides a useful complete-data reliability example for assessing the proposed residual-life functional in the presence of heterogeneous mechanical failure behavior and latent inspection uncertainty.

Table 6. Descriptive statistics for the aircraft air-conditioning failure data.

Min	Max	Mean	Median	Standard deviation
10	310	83.52	61	70.80

The left panel in Figure 3 shows the empirical reliability function for the aircraft failure times. The curve decreases steadily over the observed time range, with a relatively marked decline during the earlier and middle portions of the sample. This pattern is consistent with heterogeneous mechanical degradation, although no specific aging class is imposed in the analysis. The right panel in Figure 3 displays the estimated mean residual life under latent random inspection as a function of λ, together with bootstrap 95% confidence bands. As in the previous application, λ is not treated as an observed inspection rate. Instead, it is a sensitivity parameter representing plausible latent monitoring intensities under the exponential inspection model.

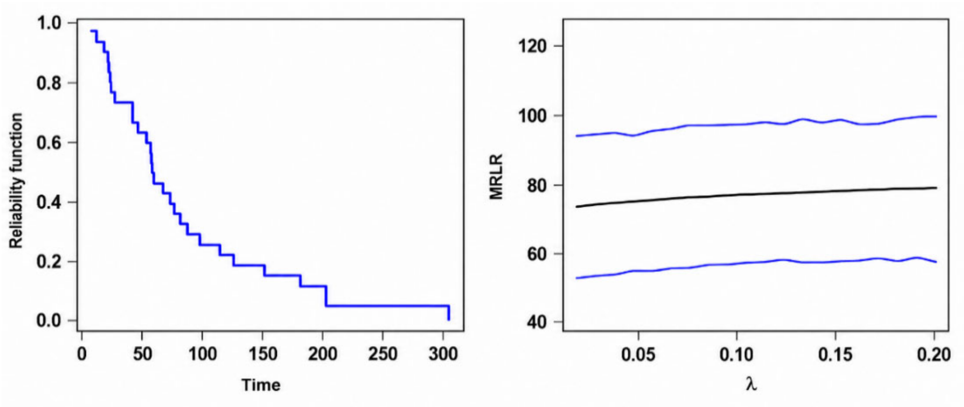

Figure 3. Aircraft air-conditioning failure analysis. Left: Empirical reliability function of the observed failure times. Right: Estimated $m_r(\lambda)$ with 95% pointwise bootstrap confidence bands. Main finding: The estimated residual life increases mildly and then stabilizes as λ increases.

For the aircraft data, the estimated curve exhibits a mild increasing pattern followed by gradual stabilization as λ increases. This behavior differs from the kidney catheter example, showing that the proposed estimator responds to the empirical structure of the lifetime data rather than imposing a predetermined qualitative trend. In this dataset, shifting the latent inspection distribution toward earlier ages is associated with slightly larger residual-life estimates over the considered range of λ, although the effect becomes weaker as the inspection intensity increases. From an engineering perspective, this result suggests that components remaining operational beyond the early stages of service tend to have a relatively stable or slightly increasing expected remaining lifetime under the assumed monitoring mechanism. This finding should be interpreted as a sensitivity result rather than as a direct recommendation regarding preventive replacement frequency.

The bootstrap bands remain relatively stable across the range of λ, suggesting that the qualitative pattern is not driven by a narrow set of values for the inspection intensity. From a reliability perspective, this analysis provides a model-based way to examine how residual-life estimates depend on assumptions about the monitoring process when only historical failure times are available. This is relevant for engineering systems where inspection policies may vary across operating environments or may not be fully documented in historical records.

The aircraft air-conditioning failure data provide a clear engineering demonstration of the importance of modeling inspection timing explicitly, rather than treating it as fixed or irrelevant. By varying λ, the proposed estimator exposes how strongly residual-life conclusions depend on the assumed monitoring intensity and on the relative contribution of early and late failure ages. The mildly increasing or stable sensitivity pattern suggests that units surviving the early-failure period may exhibit more favorable residual-life behavior, possibly due to component heterogeneity, burn-in effects, or the early removal of weaker units. This distinction is crucial in reliability practice: It separates the underlying failure-time behavior from the external inspection mechanism and shows that maintenance conclusions may change when different monitoring regimes are assumed. Therefore, the proposed framework offers reliability engineers a rigorous and interpretable sensitivity tool for assessing residual life when failure times are observed but inspection ages remain latent, unrecorded, or operationally uncertain.

The aircraft application highlights heterogeneity in mechanical failure behavior. The mildly increasing and subsequently stable $m_r(\lambda)$ profile suggests that units surviving the early operating period may form a more reliable subpopulation, consistent with burn-in effects, early removal of weaker components, or unobserved component heterogeneity. From a maintenance perspective, this pattern indicates that a single fixed-age residual-life summary may obscure differences between early-failing and longer-surviving units. The proposed framework therefore provides a structured way to assess how heterogeneous failure histories interact with alternative latent inspection regimes when historical inspection schedules are unavailable.

4.3. Modern degradation data: NASA C-MAPSS turbofan engine lifetimes

The empirical assessment is further strengthened by considering the NASA C-MAPSS turbofan engine simulation data, a widely used benchmark for degradation-based reliability analysis and remaining useful life studies [26]. The C-MAPSS repository is a widely recognized benchmark in prognostics and health management, particularly in studies of remaining useful life prediction and predictive maintenance. It consists of multivariate time-series trajectories generated from simulated turbofan engines monitored over successive operating cycles. In the training data, each engine starts from normal operating conditions, undergoes progressive degradation, and is observed until failure. Thus, the final recorded cycle of each training trajectory provides a natural run-to-failure lifetime for the corresponding engine.

This application does not aim to develop a full sensor-based remaining-useful-life prediction model. Instead, the multivariate trajectories are used to extract engine lifetimes, which are then analyzed within the proposed mean residual life framework. This use of the C-MAPSS data is fully aligned with the stated objective: To study residual-life inference when lifetime information is available but the inspection or monitoring mechanism is latent, unrecorded, or only indirectly specified. The sensor measurements are therefore not modeled directly; rather, the extracted run-to-failure times serve as degradation-derived lifetime observations.

In this application, the FD001 training subset is used. This subset contains 100 run-to-failure training trajectories and 100 test trajectories generated under a single operating condition and a single fault mode, namely high-pressure compressor degradation. For each engine unit, the lifetime is extracted as $T_i = \max_j t_{ij}$, where t_{ij} denotes the operating cycle of engine i at observation j. This yields a sample of 100 engine lifetimes measured in operating cycles. Table 7 reports the descriptive statistics of the extracted NASA C-MAPSS FD001 engine lifetimes. The lifetimes range from 128 to 362 operating cycles, with a mean of 206.31, a median of 199, and a standard deviation of 46.34. These summaries show appreciable between-engine variability in run-to-failure duration, providing context for the residual-life sensitivity analysis.

Table 7. Descriptive statistics for the NASA C-MAPSS FD001 engine lifetimes.

Min	Max	Mean	Median	Standard deviation
128	362	206.31	199	46.34

The relevance of the C-MAPSS application lies in its connection to contemporary predictive-maintenance settings. Unlike classical small failure-time datasets, C-MAPSS provides degradation-driven run-to-failure information generated from monitored engineering systems. Although this analysis does not use the sensor trajectories directly, the extracted lifetimes provide a clean and meaningful basis for evaluating the proposed residual-life functional. In practical predictive-maintenance environments, inspection frequency, data-acquisition intensity, and monitoring policies may differ across systems or may not be fully documented in retrospective records. The proposed framework offers a principled way to examine how mean residual life estimates respond to such latent monitoring assumptions through inspection-intensity parameter λ. Thus, this example demonstrates the applicability of the method to modern degradation-derived lifetime data while remaining consistent with the central focus in this paper on latent random inspection.

The left panel in Figure 4 presents the empirical reliability function of the extracted engine lifetimes. The curve summarizes the run-to-failure behavior of the 100 engine units and reflects variability in degradation progression across engines. Such variability is expected in C-MAPSS-type degradation data, where trajectories may differ in initial wear, degradation onset, and stochastic operating behavior. The right panel in Figure 4 shows the estimated mean residual life under latent random inspection as a function of λ, together with bootstrap 95% confidence bands. Here, λ is again interpreted as a model-based monitoring intensity rather than an observed inspection rate. Smaller values of λ correspond to later or less frequent latent inspections, while larger values place greater weight on earlier operating cycles.

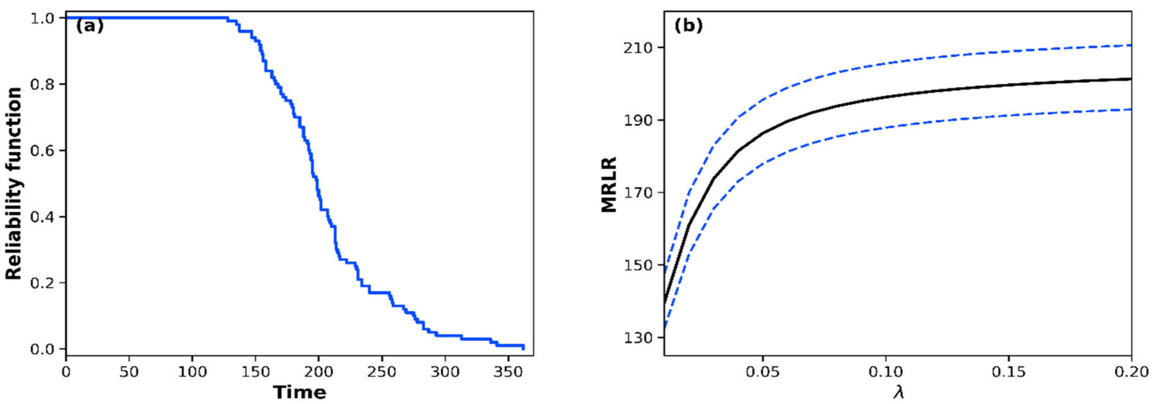

Figure 4. NASA C-MAPSS FD001 engine-lifetime analysis. Left: Empirical reliability function of the extracted run-to-failure lifetimes. Right: Estimated $m_r(\lambda)$ with 95% pointwise bootstrap confidence bands. Main finding: The residual-life profile quantifies sensitivity to alternative latent monitoring intensities.

This application complements the previous two examples in an important way. The kidney catheter data illustrate right-censored biomedical survival data, and the aircraft air-conditioning data illustrate classical mechanical failure times. The C-MAPSS application, in contrast, reflects modern degradation-based reliability data from a prognostics and predictive-maintenance context. The proposed estimator can therefore be implemented not only with traditional lifetime samples, but also with extracted run-to-failure lifetimes from contemporary degradation datasets. From an applied reliability perspective, the C-MAPSS analysis shows how residual-life inference can be connected to monitoring intensity in predictive maintenance. When inspection or monitoring times are not explicitly modeled in the original data, the proposed framework enables analysts to examine how remaining-life summaries change under different plausible latent inspection intensities. This makes the method relevant for modern reliability applications where failure-time information is available but the observation mechanism must be represented through a model-based monitoring structure.

The NASA C-MAPSS analysis provides the strongest predictive-maintenance illustration of the proposed framework. In this application, the extracted lifetimes summarize degradation-driven run-to-failure behavior, whereas the corresponding inspection ages are not directly observed and must be represented through a stochastic monitoring mechanism. Varying λ therefore reveals how estimated remaining useful life changes under different inspection intensities, rather than forcing the analysis to rely on a single fixed inspection time. This is particularly important for degradation systems, where early and late operating cycles may contain substantially different reliability information and where maintenance decisions are highly sensitive to the assumed observation regime. The resulting sensitivity pattern shows that ignoring inspection-time uncertainty can obscure important features of residual-life behavior. Thus, the proposed latent random-inspection framework provides a rigorous, interpretable, and practically relevant tool for linking degradation-based lifetime data with residual-life inference under uncertain monitoring, making it especially suitable for condition-based and predictive-maintenance applications.

The C-MAPSS application illustrates the direct relevance of the proposed methodology to predictive and condition-based maintenance. The extracted run-to-failure lifetimes summarize degradation progression across engines, while variation in $m_r(\lambda)$ quantifies how remaining useful life estimates respond to alternative monitoring intensities. This is particularly important in predictive-maintenance systems, where data-acquisition frequency and inspection policies may differ across assets or may be incompletely documented in retrospective databases. The resulting sensitivity profile can therefore support maintenance scheduling, inspection-frequency assessment, and risk-informed decisions by showing how remaining-life conclusions change under plausible latent monitoring policies.

4.4. Comparative interpretation

The three real-data applications provide complementary evidence for the practical relevance of the proposed latent random-inspection framework. They cover three distinct empirical settings: Right-censored biomedical survival data, complete engineering failure-time data, and degradation-derived run-to-failure lifetimes from a modern prognostics benchmark. This diversity is important because it shows that the method is not tied to a single data structure or application domain. Instead, the same inferential principle can be applied whenever the lifetime process is observed but the corresponding inspection or assessment ages are latent, unavailable, or only indirectly described through a monitoring mechanism.

A key message from these examples is that the sensitivity profile of $m_r(\lambda)$ is determined jointly by the empirical lifetime distribution and the assumed inspection mechanism. The inspection-intensity parameter λ does not alter the observed lifetime data or the nonparametric survival estimate; rather, it changes how early and late ages are weighted through the latent monitoring model. Thus, variation in $m_r(\lambda)$ across λ-values should be interpreted as the effect of changing the assumed observation regime, not as a change in the underlying failure-time distribution. This distinction is central to the proposed methodology and is particularly important in retrospective studies, where the event times may be available while the actual inspection process is incompletely documented.

The contrasting patterns observed across the three datasets are therefore informative rather than problematic. For the kidney catheter data, the decreasing sensitivity pattern indicates that residual-life conclusions are strongly influenced by early survival behavior, which is consistent with early catheter removals or failures among a subset of patients. For the aircraft air-conditioning data, the mildly increasing or stable profile suggests that units surviving the early-failure period may retain more favorable residual-life characteristics, possibly reflecting component heterogeneity, burn-in effects, or the early removal of weaker units. For the NASA C-MAPSS lifetimes, the sensitivity profile links degradation-driven run-to-failure behavior with assumptions about monitoring intensity, making the analysis directly relevant to condition-based and predictive-maintenance settings. These differences demonstrate that the proposed estimator does not impose a predetermined monotonic trend; instead, it adapts to the empirical survival structure of each dataset and to the selected latent inspection model.

From an applied perspective, the analysis clarifies how the proposed framework should be used. When actual inspection ages are available, their empirical distribution or a fitted inspection-time model may be incorporated directly. When inspection ages are unavailable, as in these applications, the method provides a structured sensitivity analysis over a scientifically plausible range of λ values. The resulting curves are not intended to identify a unique unobserved inspection history. Rather, they quantify how residual-life conclusions vary under alternative monitoring regimes. This interpretation is especially useful in survival, reliability, and predictive-maintenance studies where observation schedules are irregular, incompletely recorded, or governed by external operational policies.

The three applications demonstrate distinct roles of the proposed framework. The kidney-catheter example emphasizes inference under clinical censoring and uncertain follow-up, the aircraft example highlights heterogeneous mechanical failure and possible burn-in behavior, and the C-MAPSS example illustrates remaining-useful-life assessment for predictive maintenance. These contrasting applications show that the method adapts to different lifetime structures and practical objectives without imposing a common qualitative pattern across datasets.

Overall, the comparative analysis reinforces the central practical contribution of the proposed framework. Residual-life inference is influenced not only by the underlying lifetime distribution but also by how the inspection or observation process is modeled. By explicitly incorporating this latent monitoring mechanism, the proposed methodology transforms an otherwise unmodeled source of uncertainty into a transparent and statistically interpretable component of the inferential process. Consequently, it provides a more realistic and scientifically grounded basis for residual-life assessment in incomplete-observation and latent-monitoring settings while preserving the flexibility and robustness of nonparametric lifetime estimation.

5. Discussion

In this paper, a semi-parametric framework was developed for mean residual life analysis when the inspection age is latent, random, or only indirectly represented through a stochastic monitoring mechanism. By targeting the conditional residual-life functional $(T - X \mid T > X)$, the proposed approach extends classical fixed-age analysis and incorporates observation-time uncertainty directly into lifetime-distribution inference. The lifetime distribution is estimated nonparametrically, whereas the latent inspection-age distribution is specified parametrically. This separation yields a flexible and interpretable framework for reliability, survival, and predictive-maintenance settings characterized by incomplete observation, irregular monitoring, or unavailable inspection histories.

The estimator was first developed for complete lifetime data and then extended to right-censored observations through the Kaplan-Meier estimator. Under exponential inspection, explicit finite-sum representations were obtained from ordered observations and step-function survival estimates, eliminating the need for numerical double integration. Strong consistency was established under mild regularity conditions, and asymptotic normality was derived for the complete-data estimator. A plug-in estimator of the asymptotic variance was also proposed, providing analytically justified standard errors and large-sample confidence intervals. Collectively, these results supply a rigorous inferential foundation while retaining a computationally transparent implementation.

The simulation study showed coherent finite-sample behavior across Weibull and gamma lifetime distributions, distinct aging regimes, inspection intensities, sample sizes, and censoring proportions. The comparison with the fixed-age benchmark further demonstrated that replacing a random inspection age by its mean may conceal important effects of monitoring uncertainty. The three real-data analyses, kidney catheter infection times, aircraft air-conditioning failures, and NASA C-MAPSS turbofan lifetimes, illustrated the practical value of the framework in biomedical, engineering, and predictive-maintenance applications. In each case, the sensitivity profile of $m_r(\lambda)$ clarified how residual-life conclusions vary across plausible monitoring intensities when inspection ages are unavailable or incompletely recorded.

The proposed methodology nevertheless has several limitations. First, the latent inspection-age distribution G must be specified parametrically and may be unavailable or misspecified in practice. Second, the baseline theory assumes independence between the inspection age X and the lifetime T, an assumption that may fail when monitoring frequency responds to deterioration or clinical condition. Third, although the lifetime distribution F is estimated nonparametrically, residual-life estimation remains sensitive to limited upper-tail information, particularly under heavy censoring. Accordingly, conclusions should be interpreted cautiously when the inspection mechanism is poorly understood, dependence is plausible, or the identifiable follow-up horizon is short.

From an applied perspective, the framework is most useful as a structured sensitivity-analysis tool. In reliability engineering, it can support inspection planning, preventive maintenance, and replacement decisions by quantifying expected remaining life under alternative monitoring regimes. In survival analysis, it provides an interpretable residual-life summary that explicitly reflects uncertainty in the observation process. In predictive maintenance, it enables practitioners to assess how assumptions about data-acquisition or inspection intensity influence remaining-life conclusions. The method therefore complements, rather than replaces, classical fixed-age and fully parametric residual-life procedures.

From a practical perspective, the proposed methodology provides a flexible decision-support framework for lifetime analysis when inspection ages are unavailable, irregularly recorded, or only indirectly specified. In reliability engineering, it enables practitioners to evaluate expected remaining lifetime under alternative inspection policies without imposing a fully parametric lifetime model. In biomedical survival studies, it offers a principled approach for assessing residual survival under uncertain follow-up schedules while preserving the flexibility of nonparametric survival estimation. In predictive maintenance applications, the methodology supports condition-based maintenance by quantifying how different monitoring intensities influence the estimated remaining useful life of engineering systems. More broadly, the proposed framework provides a practical sensitivity-analysis tool that explicitly incorporates uncertainty in the inspection process, thereby yielding more realistic residual-life assessments for modern lifetime-data applications.

Overall, the proposed framework treats inspection-time uncertainty as an explicit component of the inferential model rather than as an ignored nuisance. By separating the lifetime distribution from the latent monitoring mechanism, it provides a statistically coherent and practically interpretable extension of classical mean residual life analysis. The following subsections summarize practical guidance, robustness considerations, the scope of the simulation study, and the principal directions for future research.

5.1. Practical guidance for selecting the inspection-age distribution

In practical applications, the latent inspection-age distribution should be selected using substantive knowledge of the monitoring process rather than mathematical convenience alone. When inspections occur approximately at random and independently over time, the exponential distribution provides a natural baseline because of its memoryless property and analytical tractability. If operational records suggest that inspection intensity changes with age, more flexible families such as Weibull, gamma, or other scientifically justified parametric models may be preferable. When little information is available, the exponential model should be treated as a reference specification and accompanied by sensitivity analyses under several plausible alternatives. Stable conclusions across these alternatives increase confidence in the resulting inference, whereas substantial discrepancies signal the need for additional information about the monitoring mechanism.

5.2. Computational efficiency

The proposed methodology is computationally efficient because it combines standard survival-function estimation with explicit analytical evaluation of the proposed estimator. Once the empirical or Kaplan-Meier estimator has been constructed, no iterative optimization, numerical maximization, or simulation-based estimation is required. Consequently, the computational cost is dominated by well-established survival-analysis procedures that are readily available in standard statistical software. This computational simplicity makes the proposed framework suitable for routine applications involving moderate and relatively large datasets. Although a dedicated computational-time benchmark is not included in this work, such an empirical evaluation across a wide range of sample sizes would provide useful complementary information and represents an interesting direction for future research.

5.3. Robustness to inspection-model misspecification

Because the lifetime distribution is estimated nonparametrically through the empirical or Kaplan-Meier estimator, the principal parametric source of misspecification is the assumed inspection-age distribution G. Misspecification may arise from an incorrect value of the inspection-intensity parameter λ or from selection of an inappropriate distributional family. Robustness should therefore be assessed by evaluating the estimator over scientifically plausible parameter values and competing inspection models. If the resulting estimates remain stable across these specifications, the corresponding residual-life conclusions may be regarded as robust. Conversely, substantial variation indicates that the inspection mechanism materially affects the analysis and should be incorporated explicitly into uncertainty assessment and practical decision-making. Formal robustness diagnostics, data-driven model selection, and model-averaging procedures for competing inspection-age distributions are natural topics for further investigation.

5.4. Scope of the simulation study

The simulation study was designed primarily to evaluate point-estimation performance through bias and mean squared error across a range of lifetime, inspection, sample-size, and censoring configurations. Additional inferential criteria, including empirical confidence-interval coverage probabilities and average interval lengths, would provide valuable complementary information. Their systematic assessment would require repeated interval construction across a substantially larger design and is therefore beyond the scope of this study. A comprehensive evaluation of interval-estimation performance remains an important direction for future methodological work.

5.5. Limitations

A fundamental assumption underlying the proposed framework is that the lifetime variable T and the latent inspection age X are statistically independent. This assumption is appropriate for many retrospective reliability and survival studies in which inspection times are determined by external administrative schedules, routine maintenance policies, or random monitoring mechanisms unrelated to the actual degradation state of the unit. However, it may be restrictive in condition-based maintenance systems, adaptive inspection strategies, or personalized clinical follow-up, where inspection frequency is influenced by observed degradation, health status, or previous failure history. In such settings, the inspection mechanism becomes informative, and the independence assumption may no longer hold. Extending the proposed methodology to accommodate dependent inspection processes represents an important direction for future research and would require joint modeling of the lifetime and inspection mechanisms or the development of semi-parametric estimators under informative inspection schemes.

5.6. Future research directions

First, an important question is the robustness of the proposed methodology under misspecified inspection-age distributions. While it is assumed in this paper that the inspection mechanism belongs to a specified parametric family, future work may entail investigating how sensitive the estimator and the associated inference procedures are to departures from this assumption and developing robust alternatives that remain reliable under distributional misspecification. Second, it would be valuable to estimate the inspection-intensity parameter λ directly from the observed data rather than treating it as a sensitivity parameter. This raises the question of how latent monitoring mechanisms can be inferred from incomplete lifetime information using likelihood-based, Bayesian, or semi-parametric estimation strategies. Third, the proposed framework may be extended to more complex observation schemes, including dependent inspection processes, covariate-adjusted mean residual life models, recurrent-event settings, and more general censoring mechanisms. Such developments would substantially broaden the applicability of the methodology to modern biomedical and engineering reliability studies.

Finally, an important computational direction is the development of fully data-driven inspection-age models using machine-learning and adaptive estimation techniques. Combining flexible monitoring-mechanism estimation with the proposed semi-parametric inference framework may provide more accurate residual-life prediction in large-scale predictive-maintenance and survival-analysis applications.

6. Conclusions

A semi-parametric framework has been developed for mean residual life inference under latent random inspection. Combining nonparametric lifetime estimation with a parametric inspection-age model makes it possible to examine observation-time uncertainty without imposing a fully parametric lifetime distribution. Under exponential inspection, explicit finite-sum formulas support direct computation for complete data and for right-censored data over an identifiable finite horizon.

The theoretical results establish strong consistency and provide asymptotic normality and a plug-in variance estimator for the complete-data estimator. The simulation study supports the finite-sample behavior of the method across the investigated lifetime models and censoring settings. The biomedical, aircraft-reliability, and NASA C-MAPSS applications illustrate its use as a sensitivity-analysis tool when inspection ages are unavailable. These applications quantify the consequences of assumed monitoring intensities; they do not identify the actual latent inspection mechanism.

Practical use should account for the assumed independence between lifetime and inspection age, the specification of the inspection-age distribution, and the availability of upper-tail information under censoring. Further work on informative inspection, robustness to model misspecification, and interval estimation for the censored-data estimator would extend the scope of the framework.

Use of Generative-AI tools declaration

AI tools were used to assist with English-language proofreading and editorial revision. The author retains sole responsibility for the manuscript content.

Acknowledgments

The author sincerely thanks the four anonymous reviewers for their valuable comments and constructive suggestions, which have substantially improved the quality and clarity of this manuscript. This research is supported by the Ongoing Research Funding program (ORF-2026-1454), King Saud University, Riyadh, Saudi Arabia.

Data and code availability

The datasets analyzed in the real-data applications are publicly available or derived from publicly accessible benchmark datasets identified in the manuscript. The simulation code and numerical routines used in this study are available from the corresponding author upon reasonable request.

Conflict of interest

The author states that there is no conflict of interest.

References

1. J. Navarro, P. J. Hernández, Mean residual life functions of finite mixtures, order statistics and coherent systems, *Metrika*, **67** (2008), 277–298. https://doi.org/10.1007/s00184-007-0133-8
2. L. Sun, Q. Zhao, A class of mean residual life regression models with censored survival data, *J. Stat. Plann. Inference*, **140** (2010), 3425–3441. https://doi.org/10.1016/j.jspi.2010.05.012
3. A. Patra, C. Kundu, Stochastic comparisons and ageing properties of RLRT (ITRT) based on variance residual life, *Commun. Stat. Theory Methods*, **51** (2022), 4276–4295. https://doi.org/10.1080/03610926.2020.1812655
4. B. Wang, X. Song, Q. Zhao, Proportional mean residual life model with varying coefficients for right-censored data, *Stat. Med.*, **44** (2025), e70008. https://doi.org/10.1002/sim.70008
5. Y. Lin, X. Yang, F. Bai, Estimation and inference for the additive mean residual life model with change point in the continuous covariate, *Stat. Med.*, **44** (2025), e70133. https://doi.org/10.1002/sim.70133
6. A. Shafiq, A. B. Çolak, S. A. Lone, T. N. Sindhu, T. Muhammad, Reliability modeling and analysis of mixture of exponential distributions using artificial neural network, *Math. Methods Appl. Sci.*, **47** (2024), 3308–3328. https://doi.org/10.1002/mma.8178
7. A. Shafiq, A. B. Çolak, T. N. Sindhu, Reliability investigation of exponentiated Weibull distribution using IPL through numerical and artificial neural network modeling, *Qual. Reliab. Eng. Int.*, **38** (2022), 3616–3631. https://doi.org/10.1002/qre.3155
8. T. N. Sindhu, A. B. Çolak, S. A. Lone, A. Shafiq, Reliability study of generalized exponential distribution based on inverse power law using artificial neural network with Bayesian regularization, *Qual. Reliab. Eng. Int.*, **39** (2023), 2398–2421. https://doi.org/10.1002/qre.3352
9. A. Shafiq, T. N. Sindhu, S. A. Lone, T. A. Abushal, M. K. H. Hassan, An updated software reliability model using the Shanker model and failure data, *Qual. Reliab. Eng. Int.*, **40** (2024), 2078–2095. https://doi.org/10.1002/qre.3512
10. M. Tao, J. Xu, S. Zhang, J. Tan, J. Xu, Precision sensitivity optimization for parallel robotic system based on multi-source preventive maintenance, *Int. J. Reliab. Qual. Saf. Eng.*, **31** (2024), 2450019. https://doi.org/10.1142/S0218539324500190
11. L. Wang, Z. Zhu, X. Zhao, Dynamic predictive maintenance decisions based on system remaining useful life prediction and three inspection strategies, *Int. J. Reliab. Qual. Saf. Eng.*, **32** (2025), 2550011. https://doi.org/10.1142/S0218539325500111

12. T. Nguyen, Y. Liao, Y. Lin, P. Chang, System reliability analysis for a stochastic liner container shipping service with stochastic terminals, *Int. J. Reliab. Qual. Saf. Eng.*, **30** (2023), 2250019. https://doi.org/10.1142/S021853932250019X
13. A. Leccadito, A. Staino, P. Toscano, A novel robust method for estimating the covariance matrix of financial returns with applications to risk management, *Financ. Innov.*, **10** (2024), 116. https://doi.org/10.1186/s40854-024-00642-2
14. M. Asadi, A. Berred, The mean residual life at random age and its connection to variability measures, *Probab. Eng. Inf. Sci.*, **39** (2025), 326–343. https://doi.org/10.1017/S026996482400010X
15. N. Cai, Y. Zheng, Preservation of generalized ageing classes on the residual life at random times, *J. Stat. Plann. Inference*, **142** (2012), 148–154. https://doi.org/10.1016/j.jspi.2011.07.004
16. M. Asadi, M. Finkelstein, On variability of the mean remaining lifetime at random age, *TEST*, **33** (2024), 717–730. https://doi.org/10.1007/s11749-023-00917-3
17. A. Patra, C. Kundu, On generalized orderings and ageing classes for residual life and inactivity time at random time, *Metrika*, **82** (2019), 691–704. https://doi.org/10.1007/s00184-018-0701-0
18. Z. Su, L. Fang, A novel method to calculate mean survival time for time-to-event data, *Commun. Stat. Simul. Comput.*, **41** (2012), 611–620. https://doi.org/10.1080/03610918.2011.598996
19. B. Park, J. Sohn, S. Lee, Nonparametric estimation of mean residual life function under random censorship, *J. Korean Stat. Soc.*, **22** (1993), 147–157.
20. D. M. Bradley, R. C. Gupta, Limiting behaviour of the mean residual life, *Ann. Inst. Stat. Math.*, **55** (2003), 217–226. https://doi.org/10.1007/BF02530495
21. A. Alvarez-Iglesias, J. Newell, C. Scarrott, J. Hinde, Summarising censored survival data using the mean residual life function, *Stat. Med.*, **34** (2015), 1965–1976. https://doi.org/10.1002/sim.6431
22. E. E. A. Aly, Nonparametric tests for comparing two mean residual life functions, *Lifetime Data Anal.*, **3** (1997), 353–366. https://doi.org/10.1023/A:1009654112947
23. M. Ghebremichael, Nonparametric estimation of mean residual functions, *Lifetime Data Anal.*, **15** (2009), 107–119. https://doi.org/10.1007/s10985-008-9089-x
24. F. Proschan, Theoretical explanation of observed decreasing failure rate, *Technometrics*, **5** (1963), 375–383. https://doi.org/10.1080/00401706.1963.10490105
25. W. Nelson, *Applied life data analysis*, New York: John Wiley & Sons, 1982. https://doi.org/10.1002/0471725234
26. A. Saxena, K. Goebel, Turbofan engine degradation simulation data set, *NASA Ames Prognostics Data Repository*, 2008. Available from: https://ti.arc.nasa.gov/tech/dash/groups/pcoe/prognostic-data-repository/.

Appendix

In this appendix, a detailed derivation of the asymptotic variance of the complete-data semi-parametric estimator is provided. The derivation makes explicit the ratio-of-means representation, the joint central limit theorem for the numerator and denominator, the delta-method calculation, and the resulting plug-in variance estimator.

Let $T_1,\ldots,T_n$ be independent and identically distributed copies of the non-negative lifetime variable T. Let G denote the specified distribution function of the latent inspection age, which is assumed to be independent of T. For a continuous inspection distribution, define

$$\phi_G(t) = G(t), \qquad \psi_G(t) = \int_0^t G(u)\,du.$$

If G is not continuous, $G(t)$ in the denominator may be replaced by the appropriate left-limit $G(t-)$, as stated in the main text. Under the integrability assumptions in Conditions (C1)–(C4), define the population quantities

$$a = E\{\psi_G(T)\}, \qquad b = E\{\phi_G(T)\} > 0.$$

The target random-inspection mean residual life can then be written as

$$m_r = \frac{a}{b}.$$

A.1. Ratio-of-empirical-means representation

For each observation, set

$$U_i = \psi_G(T_i), \qquad V_i = \phi_G(T_i), \qquad i = 1, \ldots, n.$$

Let

$$\bar{U}_n = \frac{1}{n}\sum_{i=1}^{n} U_i, \qquad \bar{V}_n = \frac{1}{n}\sum_{i=1}^{n} V_i.$$

The complete-data estimator introduced in Section 2 can therefore be expressed as

$$\hat{m}_r = \frac{\bar{U}_n}{\bar{V}_n}.$$

This representation is useful because the estimator is a smooth function of two ordinary sample means. In particular, if

$$g(x, y) = \frac{x}{y}, \qquad y > 0,$$

then

$$\hat{m}_r = g(\bar{U}_n, \bar{V}_n), \qquad m_r = g(a, b).$$

A.2. Joint central limit theorem

Define the random vector

$$Z_i = \begin{pmatrix} U_i \\ V_i \end{pmatrix}, \qquad E(Z_i) = \begin{pmatrix} a \\ b \end{pmatrix}.$$

Under Condition (C4), the required second moments are finite. Hence, by the multivariate central limit theorem,

$$\sqrt{n}\left[\begin{pmatrix} \bar{U}_n \\ \bar{V}_n \end{pmatrix} - \begin{pmatrix} a \\ b \end{pmatrix}\right] \xrightarrow{d} N_2(0, \Sigma),$$

where

$$\Sigma = \begin{pmatrix} \operatorname{Var}(U_1) & \operatorname{Cov}(U_1, V_1) \\ \operatorname{Cov}(U_1, V_1) & \operatorname{Var}(V_1) \end{pmatrix}.$$

A.3. Delta-method calculation

The gradient of $g(x, y) = x / y$ is

$$\nabla g(x, y) = \begin{pmatrix} 1 / y \\ -x / y^2 \end{pmatrix}.$$

Evaluating the gradient at (*a, b*) gives

$$\nabla g(a, b) = \begin{pmatrix} 1 / b \\ -a / b^2 \end{pmatrix} = \frac{1}{b}\begin{pmatrix} 1 \\ -m_r \end{pmatrix},$$

because $m_r = a / b$. The delta method therefore yields

$$\sqrt{n}\left(\hat{m}_r - m_r\right) \xrightarrow{d} N\left(0, \sigma_r^2\right),$$

with

$$\sigma_r^2 = \nabla g(a, b)^T \Sigma \nabla g(a, b).$$

Substituting the gradient and expanding the quadratic form gives

$$\sigma_r^2 = \frac{1}{b^2}\operatorname{Var}(U_1) + \frac{a^2}{b^4}\operatorname{Var}(V_1) - \frac{2a}{b^3}\operatorname{Cov}(U_1, V_1).$$

Using $a = m_r b$, the same expression becomes

$$\sigma_r^2 = \frac{1}{b^2}\left[\operatorname{Var}\left(U_1\right) + m_r^2 \operatorname{Var}\left(V_1\right) - 2m_r \operatorname{Cov}\left(U_1, V_1\right)\right].$$

The quantity in square brackets is the variance of $U_1 - m_r V_1$. Consequently, the asymptotic variance admits the compact representation

$$\boxed{\sigma_r^2 = \frac{\operatorname{Var}\{U_1 - m_r V_1\}}{b^2}}$$

or, equivalently,

$$\boxed{\sigma_r^2 = \frac{E\left[\{\psi_G\left(T\right) - m_r \phi_G\left(T\right)\}^2\right]}{\left[E\{\phi_G\left(T\right)\}\right]^2}}.$$

The final equality follows from

$$E\left\{U_1 - m_r V_1\right\} = a - m_r b = 0.$$

Thus, the first-order influence contribution of an observation T_i is

$$IF\left(T_i\right) = \frac{\psi_G\left(T_i\right) - m_r \phi_G\left(T_i\right)}{b}.$$

A.4. Plug-in estimator of the asymptotic variance

The unknown population quantities are replaced by their empirical counterparts. Define

$$\hat{b} = \overline{V}_n = \frac{1}{n}\sum_{i=1}^{n} \phi_G\left(T_i\right)$$

and

$$\hat{\xi}_i = U_i - \hat{m}_r V_i = \psi_G\left(T_i\right) - \hat{m}_r \phi_G\left(T_i\right).$$

Because

$$\frac{1}{n}\sum_{i=1}^{n} \hat{\xi}_i = \overline{U}_n - \hat{m}_r \overline{V}_n = 0,$$

a natural plug-in estimator of σ_r^2 is

$$\hat{\sigma}_r^2 = \frac{1}{\hat{b}^2}\left(\frac{1}{n}\sum_{i=1}^{n}\hat{\xi}_i^2\right).$$

An alternative finite-sample version replaces 1/*n* by 1/(*n*-1):

$$\hat{\sigma}_{r,unb}^2 = \frac{1}{\hat{b}^2}\left(\frac{1}{n-1}\sum_{i=1}^{n}\hat{\xi}_i^2\right).$$

Both versions are asymptotically equivalent. Since the limiting distribution is

$$\sqrt{n}\left(\hat{m}_r - m_r\right) \xrightarrow{d} N\left(0, \sigma_r^2\right),$$

the estimated standard error of $\hat{m}_r$ is

$$\widehat{\mathrm{se}}\left(\hat{m}_r\right) = \sqrt{\frac{\hat{\sigma}_r^2}{n}}.$$

Accordingly, an approximate two-sided (1-α) confidence interval for m_r is

$$\hat{m}_r \pm z_{1-\alpha/2}\sqrt{\frac{\hat{\sigma}_r^2}{n}},$$

where $z_{1-\alpha/2}$ is the $\left(1-\alpha/2\right)$ quantile of the standard normal distribution.

A.5. Exponential inspection model

When the latent inspection age follows an exponential distribution with rate $\lambda > 0$,

$$G\left(t\right) = 1 - e^{-\lambda t}, \quad t \geq 0.$$

Hence,

$$\phi_G\left(t\right) = 1 - e^{-\lambda t}$$

and

$$\psi_G\left(t\right) = \int_0^t \left(1 - e^{-\lambda u}\right) du = t - \frac{1 - e^{-\lambda t}}{\lambda}.$$

Therefore,

$$U_i = T_i - \frac{1 - e^{-\lambda T_i}}{\lambda}, \qquad V_i = 1 - e^{-\lambda T_i},$$

and the variance estimator is obtained directly from

$$\hat{\xi}_i = T_i - \frac{1-e^{-\lambda T_i}}{\lambda} - \hat{m}_r(\lambda)\left(1-e^{-\lambda T_i}\right).$$

Thus,

$$\hat{\sigma}_r^2(\lambda) = \frac{\frac{1}{n}\sum_{i=1}^{n}\hat{\xi}_i^2}{\left\{\frac{1}{n}\sum_{i=1}^{n}\left(1-e^{-\lambda T_i}\right)\right\}^2}.$$

This explicit formula shows that, under exponential inspection, the asymptotic variance and its plug-in estimator can be computed directly from the observed complete lifetime sample without numerical differentiation or iterative optimization.