mathematics

MDPI

Article

Component-Wise Shewhart-CUSUM Monitoring for Progressively Type-II Censored Weibull Lifetime Processes

Sinan Çalık and Ayse Bugatekin *

Department of Statistics, Faculty of Science, Fırat University, 23119 Elazig, Turkey; scalik@firat.edu.tr
* Correspondence: aturan@firat.edu.tr

Abstract

Progressive Type-II censoring presents important challenges for monitoring lifetime processes, particularly when distributional changes affect multiple parameters. This study proposes a component-wise combined Shewhart–CUSUM monitoring framework for progressively Type-II censored Weibull lifetime processes, with minimum density power divergence estimation (MDPDE) incorporated as an optional robust estimation component. The Weibull scale and shape estimates are transformed into finite-sample signed normal scores using Monte Carlo reference distributions, enabling directional monitoring of individual parameter changes. A two-sided Shewhart mechanism and four directional CUSUM statistics are jointly calibrated to achieve a nominal in-control average run length of approximately 370. Simulation results under early, uniform, and late progressive censoring demonstrate effective detection of scale, shape, and simultaneous parameter changes. The CUSUM component provides sensitivity to moderate persistent shifts, whereas the Shewhart component contributes to the immediate detection of large transient departures. An external comparison with an existing EWMA-MAX chart shows shorter out-of-control average run lengths for the proposed procedure across the 36 investigated configurations. Robustness analyses indicate that positive MDPDE tuning can improve shape-parameter estimation but does not consistently enhance the monitoring performance relative to maximum likelihood estimation. Its benefits depend on the censoring scheme, contamination direction, and affected parameter. The framework is further illustrated using Kevlar 49/epoxy stress-rupture data, where an immediate Shewhart signal is followed by directional CUSUM signals confirming persistent distributional changes. Overall, the proposed approach provides a censoring-aware and interpretable monitoring framework that combines complementary detection mechanisms with an optional, condition-dependent robust estimation component.

Keywords: MDPDE; density power divergence; Shewhart–CUSUM chart; progressive Type-II censoring; lifetime data

MSC: 62P30; 62N05

Check for updates

Academic Editor: Stelios Psarakis

Received: 8 October 2026
Revised: 1 November 2026
Accepted: 6 November 2026
Published: 9 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Statistical process control (SPC) has become one of the most effective and widely adopted methodologies for monitoring process stability, improving product quality, and supporting continuous quality improvement in modern manufacturing, reliability engineering, healthcare, and service systems. As industrial processes become increasingly automated and product quality continues to improve, early detection of process deterioration has become more critical than ever. The primary objective of SPC is to distinguish

common-cause variability from assignable causes before process degradation leads to unacceptable product quality, system failures, or economic losses. Among numerous SPC techniques, control charts remain the most popular monitoring tools due to their conceptual simplicity, computational efficiency, and practical applicability across a broad range of industrial and engineering problems [1]. Traditionally, control charts are classified into two major categories according to whether historical process information is incorporated into the monitoring statistic. Memoryless control charts, represented by the classical Shewhart chart, make decisions solely based on the current observation and therefore exhibit excellent performance in detecting large and abrupt process shifts. In contrast, memory-type control charts accumulate information from previous observations and are considerably more sensitive to small and persistent changes that may gradually deteriorate process performance. Among these, the cumulative sum (CUSUM) chart proposed by Page [2] and the exponentially weighted moving average (EWMA) chart introduced by Roberts [3] have become the two most influential memory-type monitoring schemes. Their statistical properties and practical advantages have been extensively investigated over the past decades, particularly following the seminal work of Lucas and Saccucci [4], which established the theoretical foundation for modern EWMA chart design. Consequently, memory-type control charts have become indispensable tools in applications where early detection of small parameter shifts is essential for maintaining process reliability and operational safety.

Although memory-type charts generally outperform Shewhart charts in detecting small and sustained process shifts, no single monitoring scheme performs optimally under all possible change scenarios. Shewhart charts react almost immediately to large and sudden shifts but are relatively insensitive to gradual changes, whereas CUSUM and EWMA charts exhibit the opposite behavior due to their cumulative memory structures. This complementary detection capability has motivated extensive research on hybrid and combined monitoring schemes that integrate the strengths of both memoryless and memory-type procedures. Early attempts by Westgard et al. [5] demonstrated that combining Shewhart and CUSUM decision rules could substantially improve overall monitoring performance across different shift magnitudes. Since then, numerous combined monitoring structures have been proposed, including adaptive, weighted, and hybrid variants that seek to achieve a better balance between rapid detection of large shifts and high sensitivity to small process changes. Recent contributions by Haq and Munir [6], Hu et al. [7], Hu et al. [8], Tyagi et al. [9], and other researchers have further demonstrated that combined Shewhart–CUSUM and Shewhart–EWMA schemes frequently outperform their individual counterparts in terms of average run length (ARL), especially when the magnitude of the process shift is unknown in advance. These studies clearly indicate that combining memoryless and memory-type statistics has become one of the most active research directions in contemporary SPC. Nevertheless, nearly all existing combined monitoring schemes have been developed under the assumption of complete observations obtained from conventional production processes, while considerably less attention has been devoted to monitoring lifetime-related quality characteristics collected under censoring mechanisms.

In many reliability engineering, biomedical, aerospace, energy, transportation, and industrial applications, the primary quality characteristic is not a directly measurable product dimension but rather a lifetime or time-to-failure variable. Lifetime data are commonly encountered in accelerated life testing, reliability qualification experiments, fatigue analysis, maintenance scheduling, and survival studies, where complete observations are often unavailable because experiments terminate before all test units fail. Consequently, censoring mechanisms naturally arise during data collection, making conventional SPC procedures developed for complete observations inappropriate or statistically inefficient. Among various censoring schemes, progressive Type-II censoring has attracted substantial attention

because it provides greater experimental flexibility, reduces testing costs and duration, and better reflects practical life-testing environments compared with conventional Type-II censoring. Over the past few years, considerable progress has been achieved in developing monitoring procedures specifically designed for progressively Type-II censored lifetime data. For example, Yu et al. [10] proposed several EWMA-based monitoring schemes for Weibull processes under progressive Type-II censoring using likelihood-ratio and weighted likelihood statistics. Zhao et al. [11] subsequently introduced a self-starting monitoring framework based on goodness-of-fit statistics for progressively censored observations, while Enami and Alamri [12] developed a Phase-II monitoring procedure for progressively censored proportional hazard processes. Other important contributions include the conditional control charts of Mukherjee et al. [13], failure-censored monitoring procedures proposed by Mohammadipour et al. [14], weighted EWMA approaches for censored Weibull lifetimes developed by Wang et al. [15], as well as more recent studies on adaptive progressive censoring and reliability inference for Weibull-type lifetime models [16–19]. Collectively, these studies demonstrate that monitoring progressively censored lifetime processes has become an increasingly important research topic in statistical quality control. However, despite these significant developments, existing methods remain largely based on likelihood-ratio statistics, maximum likelihood estimation, or goodness-of-fit procedures, and almost exclusively employ individual EWMA or CUSUM-type structures. The integration of combined monitoring strategies with statistically robust monitoring statistics for progressively Type-II censored lifetime processes has received very limited attention, thereby motivating the development of more effective and practically reliable monitoring frameworks.

In recent years, another important research direction in statistical process monitoring has been the development of robust control charts capable of maintaining satisfactory detection performance under model misspecification, non-normality, contaminated observations, and outliers. Classical Shewhart, CUSUM, and EWMA charts are generally developed under ideal distributional assumptions and often rely on conventional estimators such as the sample mean and variance. Consequently, even a small proportion of contaminated observations may substantially distort parameter estimation, inflate false alarm rates, or delay the detection of genuine process shifts. To address these limitations, numerous robust monitoring procedures have been proposed by incorporating resistant estimators, adaptive weighting strategies, or robust scale measures into conventional memory-type charts. Recent studies by [20–25] demonstrated that robust EWMA-type monitoring procedures generally provide superior performance in the presence of outliers and non-normal process distributions. Similarly, adaptive memory-type charts, including those developed by Capizzi and Masarotto [26] and subsequent extensions, further improve monitoring efficiency by dynamically adjusting chart parameters according to the current process condition. Collectively, these studies clearly indicate that robustness has become one of the major research themes in contemporary SPC. Nevertheless, the overwhelming majority of these robust monitoring procedures have been developed for complete observations, while the robustness issue has rarely been investigated for progressively Type-II censored lifetime data, where both censoring and contamination may simultaneously influence monitoring performance.

Alongside the rapid development of robust control charts, information-theoretic and divergence-based statistical methods have emerged as powerful alternatives to conventional likelihood-based inference. Unlike classical likelihood approaches, divergence measures quantify the discrepancy between the assumed probability model and the observed data, providing an attractive balance between statistical efficiency and robustness. Among various divergence measures, the Density Power Divergence (DPD) proposed by Basu et al. [27] has become one of the most influential frameworks for robust statistical inference because it introduces a tuning parameter that controls the trade-off between efficiency under the

assumed model and robustness against contaminated observations. Subsequent studies have further developed divergence-based robust inference and demonstrated its usefulness in sequential change detection and statistical process monitoring [28–30]. More recently, Park et al. [31] successfully incorporated generalized Kullback–Leibler divergence into statistical process monitoring and developed robust control charts for high-quality processes, illustrating the considerable potential of information-theoretic monitoring strategies. In parallel, entropy-based monitoring frameworks and information-driven control chart designs have also attracted increasing attention [32]. Despite these significant methodological advances, existing divergence-based monitoring approaches have almost exclusively focused on complete observations and conventional quality characteristics. To the best of our knowledge, no existing study has integrated divergence-based robust statistics with combined Shewhart–CUSUM monitoring for progressively Type-II censored lifetime processes.

The above review reveals that the existing literature has evolved along four largely independent research directions. The first direction focuses on combined monitoring schemes that integrate the complementary advantages of Shewhart and memory-type control charts to improve detection performance across different shift magnitudes. The second direction concerns monitoring procedures specifically designed for progressively Type-II censored lifetime data, where likelihood-ratio statistics, maximum likelihood estimators, or goodness-of-fit statistics are predominantly employed. The third direction aims to improve robustness against outliers and model misspecification through resistant estimators and adaptive monitoring strategies. Finally, the fourth direction introduces divergence-based statistical methodologies that provide robust alternatives to conventional likelihood inference. Although each of these research directions has made substantial progress individually, they have developed almost independently, and their integration remains largely unexplored. In particular, current monitoring schemes for progressively Type-II censored lifetime processes still suffer from several important limitations. First, existing combined monitoring procedures generally assume complete observations and therefore cannot appropriately accommodate censoring mechanisms encountered in reliability experiments. Second, most monitoring statistics for progressively censored data rely heavily on likelihood-based estimation and may become sensitive to contamination or model deviations. Third, although density power divergence has demonstrated excellent robustness in statistical estimation and hypothesis testing, its potential for monitoring progressively Type-II censored lifetime processes has received very limited attention. Furthermore, no existing study has incorporated density power divergence-based robust parameter estimation into a combined Shewhart–CUSUM monitoring framework for progressively Type-II censored Weibull processes.

Motivated by the above observations, this paper proposes a robust divergence-based Shewhart–CUSUM monitoring scheme for progressively Type-II censored Weibull lifetime processes. Instead of relying on conventional maximum likelihood estimation, the proposed methodology employs the minimum density power divergence estimator (MDPDE) to obtain robust parameter estimates under progressive Type-II censoring. These estimates are subsequently incorporated into a combined Shewhart–CUSUM monitoring framework designed to detect both small and large process shifts. The MDPDE tuning parameter allows the degree of robustification to be varied, including the likelihood-based case as a benchmark, so that the effect of robust estimation can be evaluated under different progressive censoring and contamination settings. The performance of the proposed monitoring scheme is comprehensively evaluated through extensive Monte Carlo simulation studies under various shift magnitudes, censoring schemes, and contamination scenarios. The proposed method is compared with conventional likelihood-based monitoring procedures using several performance criteria, including the Average Run Length (ARL), Standard Deviation of Run Length (SDRL), and Median Run Length (MRL). Finally, the practical use-

fulness of the proposed approach is illustrated using a real progressively Type-II censored lifetime dataset.

2. Proposed Robust Divergence-Based Monitoring Scheme

This section presents the proposed monitoring methodology for progressively Type-II censored lifetime processes. First, the progressive Type-II censoring mechanism and the underlying Weibull lifetime model are briefly introduced to establish the necessary notation. Next, the minimum density power divergence (MDPD) framework is presented as the basis for robust parameter estimation under progressive Type-II censoring. Subsequently, the proposed MDPD-based parameter estimation procedure is incorporated into a combined Shewhart–CUSUM monitoring scheme to enhance robustness against contaminated observations while maintaining high sensitivity to both small and large process shifts. Finally, the performance evaluation criteria and parameter design procedure adopted throughout the simulation study are described.

2.1. Progressive Type-II Censored Lifetime Model

Lifetime data frequently arise in reliability engineering, biomedical studies, industrial life testing, and maintenance planning, where the primary interest is to characterize and monitor the lifetime behavior of products or systems. In many practical situations, observing the exact failure times of all experimental units is economically infeasible or excessively time-consuming because life tests are terminated before every unit fails. Consequently, censoring naturally arises during data collection, making statistical procedures developed exclusively for complete samples inappropriate. Among the various censoring mechanisms, progressive Type-II censoring has received considerable attention because of its flexibility in experimental design and its ability to reduce testing time and cost while retaining sufficient information for reliable statistical inference [33].

Suppose that n identical units are initially placed on a life test and that only $m < n$ failures are intended to be observed. Let

$$\boldsymbol{R} = (R_1, R_2,\ldots,R_m)$$

denote a prespecified progressive Type-II censoring scheme, where R_i represents the number of surviving units randomly removed from the experiment immediately after the i-th observed failure. The censoring scheme satisfies

$$n = m + \sum_{i=1}^{m} R_i. \tag{1}$$

Let

$$X_{1:m:n}^{(\boldsymbol{R})},\ X_{2:m:n}^{(\boldsymbol{R})},\ \ldots,\ X_{m:m:n}^{(\boldsymbol{R})}$$

denote the progressively Type-II censored order statistics corresponding to the observed failure times. After observing $X_{i:m:n}^{(\boldsymbol{R})}$, R_i surviving units are randomly withdrawn from the experiment. This procedure continues until the mth failure is observed, at which point the life test is terminated. The resulting mm observed failure times, together with the prespecified censoring vector $\boldsymbol{R}$, constitute the available information for subsequent parameter estimation and process monitoring.

Conventional Type-II censoring is obtained as a special case when no units are removed before the final observed failure, that is,

$$R_1 = R_2 = \cdots = R_{m-1} = 0, \quad R_m = n - m$$

Furthermore, when

$$R_1 = R_2 = \cdots = R_m = 0,$$

the censoring scheme reduces to a complete sample. Progressive Type-II censoring therefore provides a flexible framework that includes both complete sampling and conventional Type-II censoring as special cases. In practical reliability studies, such a scheme may represent situations in which surviving test units are removed because of maintenance schedules, inspection requirements, resource limitations, or other operational constraints.

Let $f(x;\boldsymbol{\theta})$, $F(x;\boldsymbol{\theta})$, and $S(x;\boldsymbol{\theta}) = 1 - F(x;\boldsymbol{\theta})$ denote the probability density, cumulative distribution, and survival functions of the underlying lifetime distribution, respectively. For a progressively Type-II censored sample, the likelihood function can be expressed as

$$L(\boldsymbol{\theta}) = C_{\boldsymbol{R}} \prod_{i=1}^{m} f(x_{i:m:n};\boldsymbol{\theta}) \left[1 - F(x_{i:m:n};\boldsymbol{\theta})\right]^{R_i}, \tag{2}$$

where $C_{\boldsymbol{R}}$ is a combinatorial constant determined by the censoring scheme and is independent of the unknown model parameters [34].

Throughout this study, the underlying lifetime distribution is assumed to follow a Weibull model because of its flexibility in representing a wide range of failure-rate behaviors encountered in reliability engineering and life-testing applications. Depending on the value of its shape parameter, the Weibull distribution can represent decreasing, constant, and increasing hazard-rate patterns.

Let X denote a Weibull random variable with scale parameter $\alpha > 0$ and shape parameter $\beta > 0$. Its probability density function is

$$f(x;\alpha,\beta) = \frac{\beta}{\alpha}\left(\frac{x}{\alpha}\right)^{\beta-1} exp\left[-\left(\frac{x}{\alpha}\right)^{\beta}\right], \quad x > 0,$$

and its cumulative distribution function is

$$F(x;\alpha,\beta) = 1 - \exp\left[-\left(\frac{x}{\alpha}\right)^{\beta}\right].$$

Accordingly, the survival function is

$$S(x;\alpha,\beta) = \exp\left[-\left(\frac{x}{\alpha}\right)^{\beta}\right],$$

while the hazard rate function is given by

$$h(x;\alpha,\beta) = \frac{f(x;\alpha,\beta)}{S(x;\alpha,\beta)} = \frac{\beta}{\alpha}\left(\frac{x}{\alpha}\right)^{\beta-1}.$$

Substituting the Weibull density and survival functions into the progressive Type-II likelihood yields

$$L(\alpha,\beta) = C_{\boldsymbol{R}} \prod_{i=1}^{m} \left[\frac{\beta}{\alpha}\left(\frac{x_{i:m:n}}{\alpha}\right)^{\beta-1} exp\left\{-(R_i+1)\left(\frac{x_{i:m:n}}{\alpha}\right)^{\beta}\right\}\right]. \tag{3}$$

Equivalently,

$$L(\alpha,\beta) = C_{\boldsymbol{R}}\left(\frac{\beta}{\alpha}\right)^{m} \prod_{i=1}^{m} \left(\frac{x_{i:m:n}}{\alpha}\right)^{\beta-1} exp\left[-\sum_{i=1}^{m}(R_i+1)\left(\frac{x_{i:m:n}}{\alpha}\right)^{\beta}\right]. \tag{4}$$

Equation (4) explicitly shows how the progressive censoring scheme enters the likelihood through the factors $R_i + 1$. Thus, even when two samples originate from the same underlying Weibull distribution, different progressive censoring schemes generally lead to different likelihood contributions and different amounts of statistical information.

The unknown Weibull parameters may conventionally be estimated by maximizing the likelihood in Equation (4). Since closed-form solutions for the scale and shape parameters are generally unavailable under progressive Type-II censoring, the corresponding

maximum likelihood estimates are typically obtained numerically. Although maximum likelihood estimation has desirable asymptotic properties under a correctly specified model, it may become sensitive to atypical failure times, contaminated observations, or model misspecification. Such sensitivity is particularly relevant in sequential monitoring because unstable parameter estimates may lead to inflated false-alarm rates or delayed detection of genuine process changes.

To address this limitation, the present study employs the minimum density power divergence estimator (MDPDE) to obtain robust parameter estimates for progressively Type-II censored Weibull lifetime data. Rather than treating the observed failure times as independent complete observations, the proposed estimation framework incorporates the progressive censoring mechanism through their joint probabilistic structure. The resulting robust estimates form the basis of the divergence-based combined Shewhart–CUSUM monitoring scheme developed in the

2.2. Density Power Divergence

Robust statistical inference has received considerable attention because classical likelihood-based procedures may be sensitive to model misspecification, contaminated observations, and outliers. Among the available robust inference approaches, the density power divergence (DPD) introduced by Basu et al. [27] provides a particularly useful framework by establishing a continuous trade-off between statistical efficiency under the assumed model and robustness against departures from that model. This trade-off is controlled through a non-negative tuning parameter, allowing the classical likelihood-based approach to be recovered as a limiting case while providing increased resistance to atypical observations for positive tuning values.

Let $g(x)$ denote the true underlying probability density and let $f(x;\boldsymbol{\theta})$ denote a parametric model indexed by the parameter vector $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ For a tuning parameter $\tau > 0$, the density power divergence between g and $f_{\boldsymbol{\theta}}$ is defined as

$$d_\tau(g, f_{\boldsymbol{\theta}}) = \int \left\{ f(x;\boldsymbol{\theta})^{1+\tau} - \left(1 + \frac{1}{\tau}\right) g(x) f(x;\boldsymbol{\theta})^{\tau} + \frac{1}{\tau} g(x)^{1+\tau} \right\} dx. \tag{5}$$

The divergence satisfies

$$d_\tau(g, f_{\boldsymbol{\theta}}) \geq 0$$

with equality if and only if $g = f_{\boldsymbol{\theta}}$ almost everywhere. Hence, the DPD provides a non-negative measure of discrepancy between the true distribution and the assumed parametric model. An important limiting property is obtained as $\tau \to 0$, for which

$$d_\tau(g, f_{\boldsymbol{\theta}}) \to D_{KL}(g \| f_{\boldsymbol{\theta}})), \tag{6}$$

where D_{KL} denotes the Kullback–Leibler divergence. Consequently, likelihood-based inference is recovered as the limiting case of the DPD framework.

For estimation purposes, the final term in Equation (5),

$$\frac{1}{\tau} \int g(x)^{1+\tau} dx,$$

does not depend on $\boldsymbol{\theta}$ and can therefore be omitted from the minimization criterion. For independent observations $X_1, X_2, \ldots, X_N$ from g, replacing the expectation with its empirical counterpart leads to the conventional DPD objective

$$H_{\tau,N}(\boldsymbol{\theta}) = \int f(x;\boldsymbol{\theta})^{1+\tau} dx - \left(1 + \frac{1}{\tau}\right) \frac{1}{N} \sum_{j=1}^{N} f\left(X_j;\boldsymbol{\theta}\right)^{\tau}. \tag{7}$$

The corresponding minimum density power divergence estimator (MDPDE) is defined by

$$\hat{\boldsymbol{\theta}}_{\tau} = arg\min_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} H_{\tau,N}(\boldsymbol{\theta}) \tag{8}$$

At $\tau = 0$, interpreted through the limiting criterion, the MDPDE coincides with the maximum likelihood estimator. For $\tau > 0$, observations having low probability under the fitted model receive reduced effective influence through the factor $f\left(X_j;\boldsymbol{\theta}\right)^{\tau}$. Thus, relatively small positive values of τ generally retain greater likelihood efficiency, whereas increasing τ provides stronger resistance to contamination and atypical observations. The choice of τ therefore determines the efficiency–robustness trade-off inherent in the DPD methodology.

Under appropriate regularity conditions, the MDPDE possesses desirable large-sample properties, including consistency and asymptotic normality, while positive tuning values can provide substantially improved robustness relative to maximum likelihood estimation under contaminated sampling models [27,28]. These properties have motivated the use of density-power-divergence methods not only in robust parameter estimation and hypothesis testing but also in sequential change detection and statistical process monitoring [30,31].

The conventional empirical criterion in Equation (7), however, is formulated for independent observations and cannot be applied directly to the individual observed failure times of a progressively Type-II censored sample. Under progressive censoring, the ordered failure times and the withdrawal mechanism jointly determine the available statistical information. Therefore, treating the observed failure times as independent complete observations would disregard the dependence and survival contributions induced by the censoring scheme. For this reason, the present study extends the DPD principle to the joint density of the progressively Type-II censored sample established in Section 2.1. The resulting censoring-aware MDPDE and its incorporation into the component-wise Shewhart–CUSUM monitoring framework are developed in Section 3.

3. Proposed MDPDE-Based Shewhart–CUSUM Monitoring Scheme

3.1. Censoring-Aware Robust Parameter Estimation Based on the MDPDE

The proposed monitoring framework begins with robust estimation of the Weibull parameters from the progressively Type-II censored sample observed at each monitoring period. As discussed in Section 2.2, conventional likelihood-based estimation may be sensitive to atypical failure times and contaminated observations, and this sensitivity may subsequently affect the performance of a sequential monitoring procedure. The minimum density power divergence estimator (MDPDE) is therefore employed to provide a robust input to the proposed combined Shewhart–CUSUM monitoring scheme.

Let

$$\boldsymbol{X}^{(t)} = \left(X^{(t)}_{1:m:n},\ X^{(t)}_{2:m:n}, \ldots, X^{(t)}_{m:m:n}\right)$$

denote the progressively Type-II censored sample observed at monitoring period tt under the prespecified censoring scheme

$$\boldsymbol{R} = (R_1, R_2, \ldots, R_m), \quad n = m + \sum_{i=1}^{m} R_i.$$

The Weibull parameter vector is defined as

$$\boldsymbol{\theta} = (\alpha,\beta)^T \in \Theta,\ \Theta = \{(\alpha,\beta){:}\alpha > 0, \beta > 0\},$$

where α and β denote the scale and shape parameters, respectively. As established in Section 2.1, the joint density of the progressively Type-II censored sample is

$$p^{(\mathbf{R})}_{\boldsymbol{\theta}}(\mathbf{x}) = C_R \prod_{i=1}^{m} f(x_i;\boldsymbol{\theta})[S(x_i;\boldsymbol{\theta})]^{R_i}, \mathbf{x} \in \chi_m \tag{9}$$

where

$$\chi_m = \{\mathbf{x} = (x_1, x_2, \ldots, x_m) : 0 \le x_1 \le x_2 \le \cdots \le x_m < \infty\}$$

and

$$C_{\boldsymbol{R}} \prod_{i=1}^{m} \left[n - i + 1 - \sum_{j=1}^{i-1} R_j \right],$$

with the empty sum for $i = 1$ defined as zero. The constant $C_{\boldsymbol{R}}$ is retained explicitly because, although it is independent of the Weibull parameters, it enters the density-power terms and therefore cannot be omitted from the DPD-based estimation criterion.

An important feature of Equation (9) is that the observed failure times cannot be treated as mm independent complete observations. Their probabilistic contribution is determined jointly by the ordered failure times and the progressive withdrawal mechanism. Therefore, directly applying the conventional observation-wise MDPDE criterion in Equation (7) would disregard the dependence and survival information induced by the censoring scheme. To preserve this structure, the DPD principle is applied to the full joint density $p_{\boldsymbol{\theta}}^{(\mathbf{R})}$.

For a fixed tuning parameter $\tau > 0$, the censoring-aware empirical DPD objective function for the sample observed at monitoring period t is defined as

$$K_{\tau,t}(\boldsymbol{\theta}) = \int_{\chi_m} \left[p_{\boldsymbol{\theta}}^{(\mathbf{R})}(\mathbf{u}) \right]^{1+\tau} d\mathbf{u} - \left(1 + \frac{1}{\tau} \right) \cdot \left[p_{\boldsymbol{\theta}}^{(\mathbf{R})}\left(\mathbf{X}^{(t)} \right) \right]^{\tau}, \tag{10}$$

The term in the density power divergence that does not depend on $\boldsymbol{\theta}$ has been omitted from Equation (10), since it has no effect on the minimization. The proposed censoring-aware MDPDE is consequently defined by

$$\hat{\boldsymbol{\theta}}_{\tau,t} = arg\min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} K_{\tau,t}(\boldsymbol{\theta}) \tag{11}$$

where

$$\hat{\boldsymbol{\theta}}_{\tau,t} = \left(\hat{\alpha}_{\tau,t}, \hat{\beta}_{\tau,t} \right)^T$$

Thus, unlike an observation-wise robustification in which arbitrary weights could be assigned separately to the observed failure times, Equation (10) incorporates progressive censoring directly through the joint density. In particular, both the failure contributions $f(x_i;\boldsymbol{\theta})$ and the survival contributions $[S(x_i;\boldsymbol{\theta})]^{R_i}$ enter the robust estimation criterion through $p_{\boldsymbol{\theta}}^{(\mathbf{R})}$. Consequently, different censoring schemes generally induce different MDPDE objective functions even when the underlying Weibull parameters are identical. The proposed density power divergence criterion is constructed from the joint density of the progressively Type-II censored sample and therefore incorporates both the observed failures and the progressive withdrawal mechanism. However, this construction should not be interpreted as automatically guaranteeing bounded sensitivity to contamination of an individual lifetime observation. For positive tuning parameters, the divergence modifies the contribution of the complete censored sample configuration rather than assigning separate robustness weights to individual failure times. Consequently, the effect of an atypical lifetime may depend on its position among the observed failures, the censoring scheme, and the resulting change in the joint likelihood. The observation-level sensitivity of the estimator must therefore be assessed separately.

The tuning parameter τ controls the balance between model efficiency and robustness. In the limiting case $\tau \to 0$, the DPD criterion approaches its likelihood-based counterpart, so that

$$\hat{\boldsymbol{\theta}}_{\tau,t} \to \hat{\boldsymbol{\theta}}_{MLE,t}. \tag{12}$$

For positive values of τ, censored sample configurations receiving relatively small probability under the fitted model have reduced influence on the estimation criterion. The proposed construction therefore introduces the robustness mechanism of density power

divergence while explicitly preserving the stochastic structure induced by progressive Type-II censoring. The use of DPD methods for censored survival data has previously been established for conventional right-censoring settings; however, the present formulation differs by constructing the estimation criterion directly from the joint distribution associated with a prespecified progressive Type-II censoring scheme.

When N_0 independent progressively Type-II censored samples,

$$\boldsymbol{X}^{(1)}, \boldsymbol{X}^{(2)}, \ldots, \boldsymbol{X}^{(N_0)}$$

are available during Phase I under the same censoring scheme $\boldsymbol{R}$, the corresponding pooled objective function is defined as

$$K_{\tau,0}(\boldsymbol{\theta}) = \int \left[p_{\boldsymbol{\theta}}^{(\mathbf{R})}(\mathbf{u})\right]^{1+\tau} d\mathbf{u} - \left(1 + \frac{1}{\tau}\right) \frac{1}{N_0} \sum_{j=1}^{N_0} \left[p_{\boldsymbol{\theta}}^{(\mathbf{R})}\left(\mathbf{X}^{(t)}\right)\right]^{\tau}. \tag{13}$$

The in-control parameter estimate is then obtained as

$$\hat{\boldsymbol{\theta}}_{\tau,0} = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{arg\min}\, K_{\tau,0}(\boldsymbol{\theta})$$

The Phase-I estimate $\hat{\boldsymbol{\theta}}_{\tau,0}$ provides the robust reference model for subsequent monitoring, whereas $\hat{\boldsymbol{\theta}}_{\tau,t}$ represents the estimated process state at monitoring period t. In the following subsection, the period-specific estimates are transformed to the log-parameter scale and converted into component-wise finite-sample normal scores. This construction retains the robustness provided by the MDPDE while allowing changes in the Weibull scale and shape parameters to be monitored separately on a common standardized scale.

Computational Form of the Censoring-Aware Objective

To facilitate numerical evaluation of the first term in Equation (10), consider the transformation

$$y_i = \left(\frac{u_i}{\alpha}\right)^{\beta}, \quad i = 1, 2, \ldots, m. \tag{14}$$

Then,

$$\int_{\chi_m} \left[p_{\alpha,\beta}^{(\mathbf{R})}(\mathbf{u})\right]^{1+\tau} d\mathbf{u} = C_{\mathbf{R}}^{1+\tau} \left(\frac{\beta}{\alpha}\right)^{m\tau} \boldsymbol{A}_{\tau,\beta,\mathbf{R}}, \tag{15}$$

where

$$\boldsymbol{A}_{\tau,\beta,\mathbf{R}} = \int_{0<y_1<y_2<\cdots<y_m} \prod_{i=1}^{m} y_i^{a_{\tau,\beta}} exp\Big[-(1+\tau)\sum_{i=1}^{m}(R_i+1)y_i\Big] d\mathbf{y}, \tag{16}$$

with

$$a_{\tau,\beta} = \tau\left(1 - \frac{1}{\beta}\right).$$

The integral in Equation (16) is finite provided $a_{\tau,\beta} > -1$. This condition is automatically satisfied for $\beta \geq 1$. For $0 < \beta < 1$, the admissible tuning parameter must satisfy

$$\tau < \frac{\beta}{1-\beta} \tag{17}$$

This admissibility condition is imposed throughout the subsequent numerical analysis.

Although Equation (16) does not generally admit a simple product-form solution, its ordered integration region can be transformed by introducing the spacings

$$v_1 = y_1 \text{ and } v_i = y_i - y_{i-1}, \quad i \geq 2$$

so that

$$y_i = \sum_{j=1}^{i} v_j, v_i > 0.$$

This transformation replaces the ordered region $0 < y_1 < y_2 < \cdots < y_m < \infty$ by the positive orthant and facilitates stable numerical evaluation of $\boldsymbol{A}_{\tau,\beta,\mathbf{R}}$. An additional computational advantage is that $\boldsymbol{A}_{\tau,\beta,\mathbf{R}}$ depends on τ,β and the censoring scheme $\boldsymbol{R}$, but not on the Weibull scale parameter α. Therefore, for each fixed pair (τ,$\boldsymbol{R}$), the integral can be evaluated over a prespecified grid of β values and subsequently interpolated during numerical optimization, substantially reducing the computational burden of the Monte Carlo study. In the special case $\beta = 1$, corresponding to the exponential distribution, $a_{\tau,\beta} = 0$, and Equation (16) reduces to a product of elementary exponential integrals, providing a useful analytical benchmark for validating the numerical implementation. This computational simplification is particularly relevant because direct repeated evaluation of the ordered integral within every optimization step would otherwise be computationally demanding. The integral in Equation (16) was evaluated numerically over the ordered integration domain using a cumulative trapezoidal integration procedure on a nonuniform grid. To reduce computational cost, the logarithm of the integral was evaluated over a grid of admissible shape-parameter values and interpolated using a natural cubic spline. When a candidate shape parameter fell outside the interpolation interval, the integral was evaluated directly rather than restricting the estimate to the interpolation boundary. Parameter estimation was performed on the logarithmic parameter scale using the Nelder–Mead optimization algorithm, with the admissibility condition in Equation (17) enforced during numerical evaluation.

The period-specific censoring-aware MDPDE in Equation (11) is obtained by numerical minimization of $K_{\tau,t}(\boldsymbol{\theta})$. To ensure positivity of the Weibull parameters throughout the optimization, the reparameterization

$$\eta_1 = \log(\alpha), \quad \eta_2 = \log(\beta),$$

is employed, so that

$$\alpha = exp(\eta_1), \quad \beta = exp(\eta_2).$$

The optimization is therefore performed over the unconstrained parameter vector

$$\boldsymbol{\eta} = (\eta_1, \eta_2)^T.$$

For each candidate value of β, the integral term $\boldsymbol{A}_{\tau,\beta,\mathbf{R}}$ in Equation (16) is evaluated using the transformed representation described above, or obtained from the corresponding precomputed interpolation grid. The resulting parameter estimates are then obtained as

$$\hat{\alpha}_{\tau,t} = exp(\hat{\eta}_{1,\tau,t}), \quad \hat{\beta}_{\tau,t} = exp(\hat{\eta}_{2,\tau,t}) \tag{18}$$

The same numerical procedure is applied to the Phase-I objective function $K_{\tau,t}(\boldsymbol{\theta})$. in Equation (13). Suitable likelihood-based estimates may be used as initial values to improve numerical stability and computational efficiency. The admissibility condition in Equation (17) is enforced whenever $0 < \beta < 1$. The log-parameter representation introduced above is also used in the monitoring stage developed in Section 3.2, where the period-specific robust estimates are compared on a common covariance-standardized scale.

3.2. Component-Wise Robust Normal-Score Monitoring Statistics

After obtaining the censoring-aware MDPDE at each monitoring period, the next step is to construct monitoring statistics that preserve the information associated with changes in the individual Weibull parameters. Because the Weibull scale and shape parameters are strictly positive and may exhibit different finite-sample variability, the period-specific parameter estimates are first represented on the logarithmic scale. Let

$$\hat{\boldsymbol{\eta}}_{\tau,t} = \begin{pmatrix} \hat{\eta}_{\alpha,\tau,t} \\ \hat{\eta}_{\beta,\tau,t} \end{pmatrix} = \begin{pmatrix} log\hat{\alpha}_{\tau,t} \\ log\hat{\beta}_{\tau,t} \end{pmatrix} \tag{19}$$

The logarithmic representation preserves parameter positivity and expresses multiplicative changes in the original Weibull parameters as additive changes on the transformed scale. In particular, increases and decreases in either parameter retain their direction on the log scale, which is important for sequential monitoring of both upward and downward process changes.

The finite-sample distributions of $\hat{\eta}_{\alpha,\tau,t}$ and $\hat{\eta}_{\beta,\tau,t}$ depend on the in-control Weibull parameters, the robustness tuning parameter τ, the number of observed failures mm, and the progressive Type-II censoring scheme $\boldsymbol{R}$. For the relatively small progressively censored samples considered in Phase-II monitoring, these distributions need not be symmetric or centered exactly at the corresponding log-transformed model parameters. Therefore, rather than imposing an asymptotic parametric reference distribution, the component-wise in-control distributions are determined by Monte Carlo simulation under the fitted Phase-I reference model.

For each fixed combination of τ and $\boldsymbol{R}$, a sufficiently large number B_0 of independent progressively Type-II censored samples are generated from the Phase-I reference model. The period-specific parameter estimates are obtained according to Section 3.1 and transformed to the logarithmic scale. Let

$$F_{\alpha,0,\tau,\boldsymbol{R}}(x) = P_0\big(\hat{\eta}_{\alpha,\tau,t} \le x\big),$$

and

$$F_{\beta,0,\tau,\boldsymbol{R}}(x) = P_0\Big(\hat{\eta}_{\beta,\tau,t} \le x\Big), \tag{20}$$

denote the corresponding finite-sample in-control marginal distribution functions. In practice, these distributions are estimated empirically from the Monte Carlo reference samples.

At monitoring period t, the two parameter components are transformed through their respective in-control probability integral transforms:

$$U_{\alpha,t,\tau} = F_{\alpha,0,\tau,\boldsymbol{R}}\big(\hat{\eta}_{\alpha,\tau,t}\big)$$

and

$$U_{\beta,t,\tau} = F_{\beta,0,\tau,\boldsymbol{R}}\Big(\hat{\eta}_{\beta,\tau,t}\Big). \tag{21}$$

The corresponding component-wise normal-score monitoring statistics are defined as

$$Z_{\alpha,t,\tau} = \Phi^{-1}(U_{\alpha,t,\tau}),$$

and

$$Z_{\beta,t,\tau} = \Phi^{-1}(U_{\beta,t,\tau}), \tag{22}$$

where $\Phi^{-1}(.)$ denotes the standard normal quantile function. Under the in-control model and an adequately estimated reference distribution, each transformed component is approximately standard normal in its marginal distribution.

An important feature of Equation (22) is that the direction of parameter departure is retained. Large positive values of $Z_{\alpha,t,\tau}$ indicate unusually large values of the estimated log-scale parameter relative to its in-control distribution, whereas large negative values indicate unusually small values. The same interpretation applies to $Z_{\beta,t,\tau}$ for the Weibull shape parameter. Consequently, upward and downward changes are not collapsed into a single non-negative distance measure.

This component-wise construction is particularly useful when the scale and shape estimators exhibit different finite-sample responses to process changes. Rather than combining their departures into a single quadratic statistic, the proposed transformation preserves the marginal evidence associated with each parameter and places both components on a common normal-score scale. Statistical dependence between $Z_{\alpha,t,\tau}$ and $Z_{\beta,t,\tau}$ is not assumed

to vanish; instead, their joint in-control behavior is retained in the Monte Carlo calibration of the combined monitoring procedure developed in Section 3.4.

When empirical distribution functions are used, probability values are restricted away from zero and one to avoid infinite normal scores. For a reference sample of size B_0, a midpoint empirical probability may be defined as

$$\hat{U}_{j,t,\tau} = \frac{N_{j,t} + 0.5}{B_0 + 1} \tag{23}$$

where $N_{j,t}$ denotes the number of simulated in-control reference estimates for component j not exceeding the current value $\hat{\eta}_{j,\tau,t}$. This construction ensures that

$$0 < \hat{U}_{j,t,\tau} < 1$$

and therefore produces finite values of the corresponding normal-score statistic.

For $\tau = 0$, the period-specific parameter estimates are obtained by maximum likelihood, whereas for $\tau > 0$ they are obtained using the censoring-aware MDPDE defined in Section 3.1. Thus, $\tau = 0$ provides the likelihood-based non-robust benchmark, while positive values of τ introduce robustness through the parameter-estimation stage. The same component-wise normal-score transformation is subsequently applied for all values of τ, allowing the effect of robust estimation to be evaluated under a common monitoring structure.

The pair

$$\boldsymbol{Z}_{t,\tau} = \begin{pmatrix} Z_{\alpha,t,\tau} \\ Z_{\beta,t,\tau} \end{pmatrix} \tag{24}$$

therefore provides the input to the combined Shewhart–CUSUM monitoring scheme. Because both positive and negative values contain information about departures from the in-control state, the subsequent monitoring procedure is formulated to detect changes in both directions for each Weibull parameter.

Although the proposed monitoring framework tracks the Weibull scale and shape parameters through separate directional statistics, a signal from a particular component should not be interpreted as definitive identification of the underlying parameter change. Because the parameter estimators and their monitoring statistics may be statistically dependent, a change in one parameter can also affect the signaling behavior of the other component. Accordingly, the component-wise structure provides parameter-specific diagnostic information rather than guaranteed identification of the true source of a distributional change.

3.3. Combined Shewhart–CUSUM Monitoring Scheme

The component-wise normal-score statistics $Z_{\alpha,t,\tau}$ and $Z_{\beta,t,\tau}$ defined in Section 3.2 quantify the standardized departures of the current robust Weibull scale and shape estimates, respectively, from their finite-sample in-control reference distributions. Unlike a non-negative omnibus distance measure, these statistics retain the direction of parameter departure: positive values indicate unusually large parameter estimates relative to the in-control condition, whereas negative values indicate unusually small estimates.

To provide sensitivity to both abrupt and persistent changes in either Weibull parameter, the two component-wise statistics are incorporated into a combined Shewhart–CUSUM monitoring structure. The Shewhart component is designed to respond immediately when either parameter exhibits an unusually large departure in absolute magnitude. The CUSUM component, in contrast, accumulates moderate but persistent departures separately for the positive and negative directions of each parameter. Consequently, the proposed monitoring scheme can respond to increases or decreases in the Weibull scale parameter, the shape parameter, or both, without requiring the direction of a future process change to be specified in advance.

The combined structure therefore preserves the complementary roles of Shewhart and CUSUM monitoring while retaining the parameter-specific directional information contained in the robust estimates. The two parameter components are not assumed to be statistically independent; their joint finite-sample behavior under progressive Type-II censoring is accounted for through Monte Carlo calibration of the overall stopping rule.

3.3.1. Shewhart Component

To detect abrupt changes in either Weibull parameter, define the component-wise Shewhart monitoring statistic as

$$M_{t,\tau} = \max\{|Z_{\alpha,t,\tau}|,|Z_{\beta,t,\tau}|\}. \tag{25}$$

Let $H_S > 0$ denote the Shewhart control limit. At monitoring period tt, the Shewhart component signals an out-of-control condition whenever

$$M_{t,\tau} > H_S. \tag{26}$$

Equivalently, a Shewhart signal is generated if

$$|Z_{\alpha,t,\tau}| > H_S$$

or

$$|Z_{\beta,t,\tau}| > H_S.$$

Because the component-wise normal scores retain the direction of departure, taking their absolute values allows the Shewhart component to detect unusually large deviations in either direction. Thus, a single common upper limit on $M_{t,\tau}$ is sufficient to represent a two-sided monitoring rule for both Weibull parameters. The dependence between $Z_{\alpha,t,\tau}$ and $Z_{\beta,t,\tau}$ is not ignored when determining H_S; the control limit is calibrated using their joint in-control distribution under the specified progressive censoring scheme.

3.3.2. CUSUM Component

To accumulate evidence of moderate but persistent departures in either direction, separate positive and negative CUSUM statistics are constructed for the scale and shape components. For the scale parameter, define

$$C^+_{\alpha,t,\tau} = \max\left\{0, C^+_{\alpha,t-1,\tau} + Z_{\alpha,t,\tau} - k\right\}$$

and

$$C^-_{\alpha,t,\tau} = \max\left\{0, C^-_{\alpha,t-1,\tau} - Z_{\alpha,t,\tau} - k\right\}. \tag{27}$$

Similarly, for the shape parameter,

$$C^+_{\beta,t,\tau} = \max\left\{0, C^+_{\beta,t-1,\tau} + Z_{\beta,t,\tau} - k\right\}$$

and

$$C^-_{\beta,t,\tau} = \max\left\{0, C^-_{\beta,t-1,\tau} - Z_{\beta,t,\tau} - k\right\}. \tag{28}$$

All four CUSUM statistics are initialized at zero:

$$C^+_{\alpha,0,\tau} = C^-_{\alpha,0,\tau} = C^+_{\beta,0,\tau} = C^-_{\beta,0,\tau} = 0. \tag{29}$$

Here, $k > 0$ denotes the common reference or allowance value. The positive CUSUMs accumulate persistent positive departures in the corresponding component-wise normal scores, whereas the negative CUSUMs accumulate persistent negative departures. Thus, increases and decreases in either Weibull parameter are monitored separately while retaining a common standardized monitoring scale.

Let $H_C > 0$ denote the common CUSUM control limit. For notational convenience, define

$$C_{t,\tau}^{max} = \max\left\{C_{\alpha,t,\tau}^{+}, C_{\alpha,t,\tau}^{-}, C_{\beta,t,\tau}^{+}, C_{\beta,t,\tau}^{-}\right\}. \tag{30}$$

The CUSUM component signals an out-of-control condition whenever

$$C_{t,\tau}^{max} > H_C. \tag{31}$$

Equivalently, a signal is generated when any one of the four directional CUSUM statistics exceeds H_C. A single moderate departure need not trigger an alarm; however, repeated departures in the same parameter and direction accumulate over successive monitoring periods and may eventually produce a CUSUM signal. This construction provides sensitivity to sustained changes while preserving information about both the affected Weibull parameter and the direction of departure.

3.3.3. Combined Decision Rule

The proposed robust combined Shewhart–CUSUM chart operates the component-wise Shewhart and four directional CUSUM mechanisms simultaneously. At monitoring period t, an out-of-control signal is generated if either

$$M_{t,\tau} > H_S$$

or

$$C_{t,\tau}^{max} > H_C.$$

Equivalently, the stopping time of the proposed monitoring procedure is defined as

$$\mathcal{T} = \inf\left\{t \geq 1 : M_{t,\tau} > H_S \text{ or } C_{t,\tau}^{max} > H_C\right\}. \tag{32}$$

If neither condition is satisfied, the process is regarded as being in control at monitoring period t, and monitoring continues with the next progressively Type-II censored sample.

The two components play complementary roles in the combined scheme. An unusually large instantaneous departure in either $Z_{\alpha,t,\tau}$ or $Z_{\beta,t,\tau}$, irrespective of direction, may trigger the Shewhart component through $M_{t,\tau}$. In contrast, smaller but persistent positive or negative departures are accumulated separately by the corresponding directional CUSUM statistic. Consequently, the proposed procedure is designed to detect abrupt as well as sustained increases or decreases in the Weibull scale parameter, the shape parameter, or both.

It should be emphasized that the Shewhart and CUSUM components are not calibrated independently. Because any of the component-wise Shewhart or directional CUSUM conditions can generate an alarm, the overall in-control performance depends jointly on the robustness tuning parameter τ, the CUSUM reference value k, the control limits H_S and H_C, the dependence between the two component-wise normal scores, and the progressive Type-II censoring scheme $\boldsymbol{R}$. Therefore, H_S and H_C are calibrated jointly so that the complete stopping rule achieves the prespecified in-control average run length. The calibration and design-selection procedure is developed in Section 3.4.

For clarity, the proposed monitoring procedure can be summarized as follows. At monitoring period tt, a new progressively Type-II censored sample $\boldsymbol{X}^{(t)}$ is obtained under the prespecified censoring scheme $\boldsymbol{R}$. The period-specific Weibull parameter estimate $\hat{\boldsymbol{\theta}}_{t,\tau}$ is first obtained according to Section 3.1 and transformed to the log-parameter scale. The component-wise normal scores $Z_{\alpha,t,\tau}$ and $Z_{\beta,t,\tau}$ are then calculated according to Section 3.2. The Shewhart statistic $M_{t,\tau}$ is evaluated according to Equation (25), while the four directional CUSUM statistics are updated according to Equations (27) and (28). If either the Shewhart condition in Equation (26) or the CUSUM condition in Equation (31) is satisfied, the process is declared out-of-control; otherwise, monitoring proceeds to period $t + 1$.

For clarity, the complete Phase-I calibration and Phase-II monitoring procedure of the proposed component-wise Shewhart–CUSUM scheme is summarized in Figure 1.

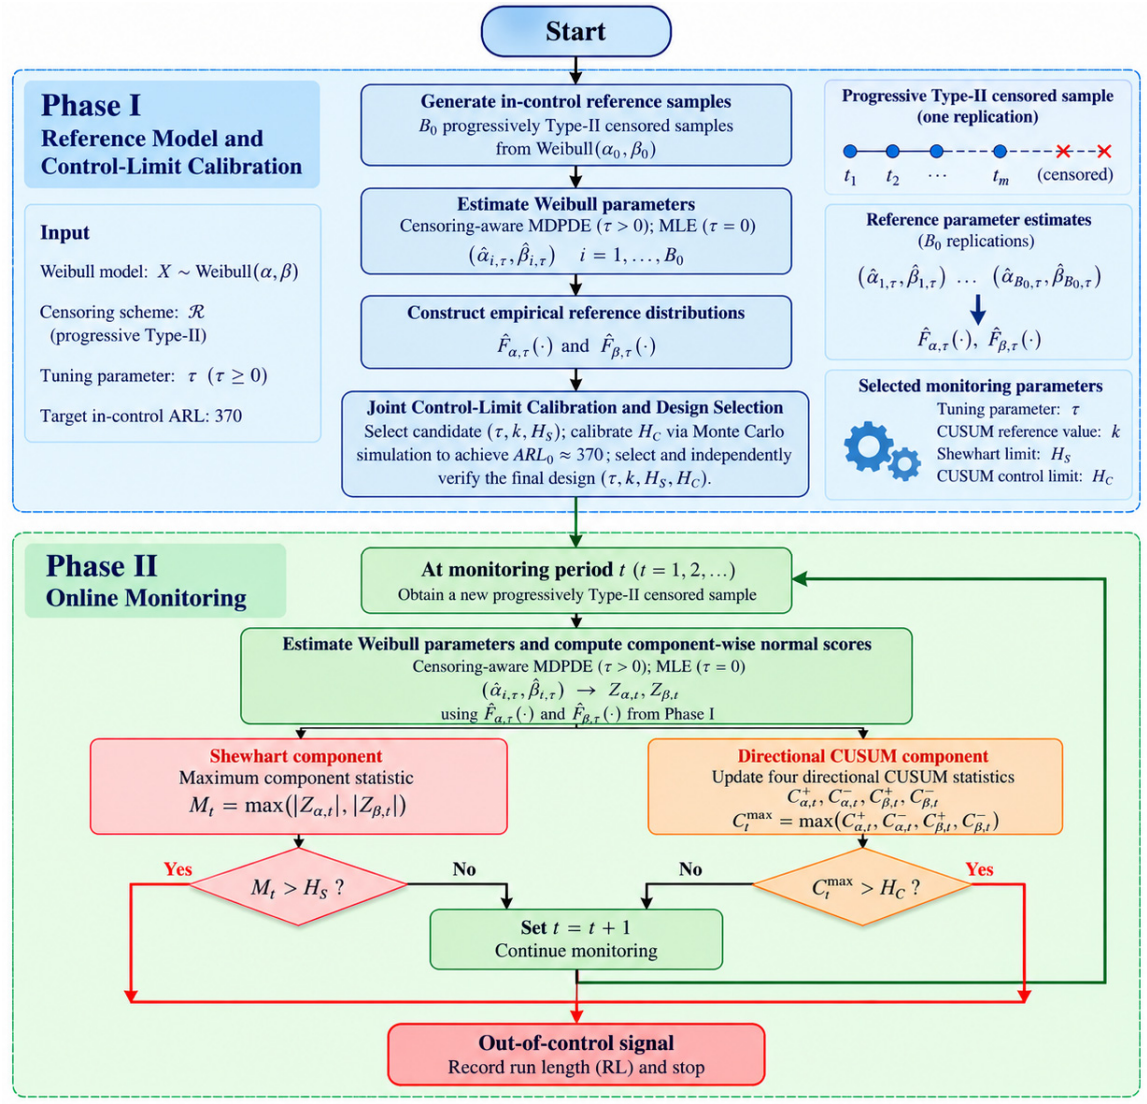

Figure 1. Flowchart of the proposed component-wise Shewhart–CUSUM monitoring procedure.

3.4. Control-Limit and Parameter Design

The performance of the proposed combined Shewhart–CUSUM monitoring scheme is governed by four design quantities: the robustness tuning parameter τ, the CUSUM reference value k, the Shewhart control limit H_S, and the CUSUM control limit H_C. In addition, the finite-sample behavior of the chart depends on the in-control Weibull parameter vector, the number of observed failures, and the progressive Type-II censoring scheme. Consequently, the control limits are calibrated for the specific monitoring configuration under consideration.

Let

$$\psi = (\tau, k, H_S, H_C)$$

denote the design vector of the proposed chart. For a fixed censoring scheme $\boldsymbol{R}$, sampling configuration (n, m), and in-control Weibull parameter vector $\boldsymbol{\theta}_0$, the in-control average run length is defined as

$$ARL_0(\psi) = E_{\boldsymbol{\theta}_0}\left(\mathcal{T} \middle| \psi, \boldsymbol{R}, n, m\right), \tag{33}$$

where Υ is the stopping time defined in Equation (32). Throughout this study, the target in-control performance is set to

$$ARL_0^* = 370, \tag{34}$$

which is adopted as a conventional benchmark widely used for comparing control-chart performance.

Because either the Shewhart or CUSUM component can trigger an alarm, their control limits must not be calibrated independently. For each prespecified combination of τ, k and censoring configuration, the pair (H_S,H_C) is therefore calibrated jointly so that

$$ARL_0(\tau, k, H_S,\ H_C|\boldsymbol{R}, n, m, \theta_0) \approx ARL_0^*. \tag{35}$$

Candidate values of H_S for which the Shewhart component alone yields an in-control ARL not exceeding ARL_0^* are excluded before joint calibration. This feasibility condition is necessary because the run length of the combined chart cannot exceed that of its Shewhart component when either component is allowed to trigger a signal.

The control limits are jointly calibrated through Monte Carlo simulation under the in-control Weibull reference model. For each prespecified combination of the tuning parameter, CUSUM reference value, and progressive Type-II censoring scheme, independent censored samples are generated from the Phase-I reference distribution. The corresponding Weibull parameter estimates are obtained according to Section 3.1 and transformed into component-wise normal scores using the fixed finite-sample reference distributions defined in Section 3.2. At each monitoring period, the Shewhart statistic and four directional CUSUM statistics are evaluated according to Equations (25), (27) and (28). The run length is recorded as the first period in which either the Shewhart condition in Equation (26) or the CUSUM condition in Equation (31) is satisfied. Repeating this procedure over (B) independent monitoring sequences yields the Monte Carlo estimate of the in-control average run length given in Equation (36).

During control-chart calibration, the normal-score transformed statistic $Z_{\alpha,t,\tau}$ is calculated according to Equation (24), the CUSUM statistic is updated according to Equation (27), and the run length is recorded as the first monitoring period at which either $Z_{t,\tau} > H_S$ or $C_{t,\tau} > H_C$. Repeating this procedure over B independent replications yields

$$\hat{ARL}_0 = \frac{1}{B}\sum_{b=1}^{B} RL_b, \tag{36}$$

where RL_b denotes the run length obtained from the bth replication. A candidate pair (H_S,H_C) is regarded as adequately calibrated when the estimated in-control ARL lies within a prespecified numerical tolerance around ARL_0^*. The specific tolerance and numbers of Monte Carlo replications used for calibration and final verification are given in Section 4.

The tuning parameter τ determines the efficiency–robustness trade-off of the MDPDE. The limiting case $\tau = 0$ corresponds to likelihood-based estimation and is retained as a non-robust benchmark, whereas positive values of τ define robust versions of the proposed chart. Importantly, the same component-wise normal-score monitoring structure is used for all values of τ. Thus, the effect of τ enters the monitoring procedure through the robustness of the estimated Weibull parameters rather than through an additional divergence measure. The common reference value k determines the magnitude of a standardized directional departure that is tolerated before evidence begins to accumulate in the corresponding CUSUM statistic.

After calibration, the out-of-control performance of a design is evaluated under specified changes in the Weibull parameters. Let

$$\alpha_1 = \rho_\alpha \alpha_0, \quad \beta_1 = \rho_\beta \beta_0, \tag{37}$$

where $\rho_\alpha > 0$ and $\rho_\beta > 0$ denote multiplicative changes in the scale and shape parameters, respectively. The in-control state corresponds to

$$(\rho_\alpha,\rho_\beta) = (1,1).$$

Scale changes are represented by $\rho_\alpha \neq 1$ and $\rho_\beta = 1$. shape changes by $\rho_\alpha = 1$ and $\rho_\beta \neq 1$, and simultaneous distributional changes by $\rho_\alpha \neq 1$ and $\rho_\beta \neq 1$. Because the component-wise normal scores retain the sign of parameter departure and the directional CUSUM statistics separately accumulate positive and negative evidence, increases and decreases in either Weibull parameter can be monitored within the same combined framework.

For an out-of-control parameter vector $\boldsymbol{\theta}_1$, the corresponding average run length is defined as

$$ARL_1(\psi;\boldsymbol{\theta}_1) = E_{\boldsymbol{\theta}_1}(\mathcal{T}). \tag{38}$$

Smaller values of ARL_1, subject to approximately equivalent ARL_0, indicate faster detection. However, a design is not regarded as acceptable solely because it performs well on average across several alternatives. As a basic detection requirement, each candidate design must satisfy

$$ARL_{1,d}^{(s)} < ARL_{0,d} \quad \textit{for every prespecified core shift scenario } s. \tag{39}$$

This requirement prevents configurations that become less responsive than under the in-control condition for one or more important shift directions from being selected on the basis of favorable performance elsewhere. Among designs satisfying this requirement, overall detection performance is subsequently compared across the prespecified collection of scale, shape, and simultaneous shift scenarios.

Robustness is evaluated separately under contaminated sampling scenarios. The nominal control limits and reference quantities are retained without recalibration when contamination is introduced. Two aspects of robustness are considered: in-control contamination, where the underlying Weibull parameters remain unchanged but atypical lifetime observations are introduced, and out-of-control contamination, where contamination is present simultaneously with a genuine parameter change. The former assesses resistance to contamination-induced false alarms or excessive insensitivity, whereas the latter evaluates whether genuine process changes remain detectable in the presence of atypical observations. This distinction follows the robustness evaluation proposed for the revised monitoring structure.

In addition to the average run length, the standard deviation and median of the run-length distribution are used to characterize monitoring performance:

$$SDRL = \sqrt{\frac{1}{B-1}\sum_{b=1}^{B}(RL_b - A\hat{R}L)^2}, \tag{40}$$

and

$$MRL = median(RL_1,RL_2,\ldots,RL_B). \tag{41}$$

Reporting *ARL*, *SDRL*, and *MRL* together provides a more complete description of the run-length distribution, particularly when it is strongly right-skewed.

Accordingly, the design procedure consists of three stages. First, candidate values of τ, k and the control-limit allocation are specified and jointly calibrated under the nominal in-control model to achieve $ARL_0 \approx 370$. Second, calibrated designs are screened according to the minimum detection requirement in Equation (39) and subsequently compared across the prespecified nominal shift scenarios. Finally, contamination scenarios are applied only after nominal design selection to evaluate the robustness of the selected monitoring configurations.

The Monte Carlo calibration described above is conditional on the specified in-control Weibull parameters. When these parameters are estimated from Phase-I data, the resulting control limits and nominal in-control average run length are conditional on the fitted reference model. Consequently, the reported calibration does not incorporate the additional uncertainty arising from Phase-I parameter estimation. Unconditional false-alarm performance would require accounting for the sampling variability of the Phase-I estimates, for example through an appropriate bootstrap procedure that repeats both parameter estimation and control-limit calibration.

Allocation of the Shewhart and CUSUM Control Limits

Since the combined monitoring scheme contains two control limits, the in-control ARL_0 requirement alone does not uniquely determine the pair (H_S,H_C). To obtain an identifiable and reproducible design, a grid-based allocation procedure is adopted. For each prespecified combination of the robustness parameter τ and the CUSUM reference value k, a finite set of candidate Shewhart limits,

$$\boldsymbol{H}_S = \left\{H_S^{(1)}, H_S^{(2)}, \ldots, H_S^{(L)}\right\}, \tag{42}$$

is considered.

Before calibrating the corresponding CUSUM limit, each candidate $H_S^{(l)}$ is evaluated using the Shewhart component alone under the in-control model. Candidate values for which

$$ARL_{0,Shewhart}\left(H_S^{(l)}\right) \leq ARL_0^*$$

are excluded from further consideration, since the combined-chart ARL cannot exceed the Shewhart-only ARL when either component is allowed to trigger a signal. Therefore, only Shewhart limits satisfying

$$ARL_{0,Shewhart}\left(H_S^{(l)}\right) > ARL_0^*$$

are retained for joint calibration.

For each retained candidate $H_S^{(l)}$ $l = 1, 2, \ldots, L$, the corresponding CUSUM limit $H_C^{(l)}$ is determined by Monte Carlo calibration such that the combined stopping rule satisfies

$$\hat{ARL}_0\left(\tau, k, H_S^{(l)}, H_C^{(l)}\right) \approx ARL_0^*. \tag{43}$$

The calibration of $H_C^{(l)}$ is performed numerically while $H_S^{(l)}$, τ, and k are held fixed. Consequently, every admissible pair $(H_S^{(l)}, H_C^{(l)})$ provides approximately the same nominal in-control protection, and differences in out-of-control performance are not attributable to unequal false-alarm rates.

The calibrated designs are subsequently evaluated over a prespecified set of core out-of-control scenarios,

$$\boldsymbol{S} = \left\{\boldsymbol{\theta}_1^{(1)}, \boldsymbol{\theta}_1^{(2)}, \ldots, \boldsymbol{\theta}_1^{(K)}\right\}, \tag{44}$$

including scale, shape, and simultaneous parameter changes in both directions. Before aggregate performance is considered, each candidate design must satisfy the minimum detection requirement defined in Equation (39) for every core shift scenario. Designs violating this requirement for any scenario are excluded from further consideration. This screening condition prevents a design with poor or paradoxical detection behavior in one shift direction from being selected solely because of favorable performance under other alternatives.

$$G_d = \frac{1}{K}\sum\nolimits_{s=1}^{K} \frac{ARL_{1,d}^{(s)}}{\min\limits_{d' \in D_A} ARL_{1,d'}^{(s)}}, \tag{45}$$

where D_A denotes the set of candidate designs satisfying the ARL_0 requirement. Since each ratio in Equation (45) is at least one,

$$G_d \geq 1$$

smaller values of G_d indicate better overall detection performance, with $G_d = 1$ represents the ideal lower bound and is attained only when the same design achieves the minimum ARL_1 under every considered shift scenario.

The final design is selected according to

$$d^* = arg\min_{d \in D_A} G_d, \tag{46}$$

subject to the prescribed ARL_0 tolerance. The shift scenarios used in $\boldsymbol{S}$, the candidate values of τ and k, the Shewhart-limit grid, and the Monte Carlo tolerance are specified a priori in the simulation design. Contamination scenarios are not used to select the nominal control limits; instead, they are reserved for evaluating the robustness of the selected chart.

4. Simulation Study

A comprehensive Monte Carlo simulation study is designed to evaluate the finite-sample performance of the proposed robust component-wise combined Shewhart–CUSUM monitoring scheme. The simulation experiments are intended to assess three main aspects of the proposed methodology: its ability to maintain the prespecified in-control performance, its sensitivity to different types, directions, and magnitudes of changes in the Weibull lifetime distribution, and its robustness against contaminated observations. Particular attention is also given to the effect of the progressive Type-II censoring pattern on monitoring performance.

The simulation study is organized in several stages. First, the in-control configuration and candidate design parameters are specified, and the component-wise finite-sample reference distributions required for the normal-score transformations are established under the nominal in-control model. The Shewhart and CUSUM control limits are then jointly calibrated to achieve a target in-control average run length of approximately 370. The calibrated designs are subsequently evaluated under scale, shape, and simultaneous changes in the Weibull parameters in both directions. Designs failing to satisfy the minimum detection requirement specified in Equation (39) are excluded, and the remaining admissible configurations are compared using the aggregate detection criterion defined in Section Allocation of the Shewhart and CUSUM Control Limits. Robustness is subsequently investigated under contaminated lifetime distributions without recalibrating the nominal reference distributions or control limits. Finally, the selected combined monitoring scheme is compared with the non-robust benchmark and the corresponding single-component monitoring procedures under equivalent in-control conditions. The average run length (ARL), standard deviation of the run length ($SDRL$), and median run length (MRL) are used as the principal performance measures throughout the study.

4.1. Simulation Design and In-Control Configuration

The simulation study is conducted under a Weibull lifetime model with in-control scale and shape parameters $\alpha_0 = 1$ and $\beta_0 = 2$, respectively. This configuration provides a standardized lifetime scale and represents an increasing failure-rate process, while allowing the effects of parameter changes and progressive censoring to be examined without unnecessary scale dependence. At each monitoring period, $n = 20$ units are initially placed on test, and the experiment is terminated after $m = 10$ failures have been observed, corresponding to an overall censoring proportion of 50%.

To investigate the effect of the timing of removals independently of the overall censoring proportion, three progressive Type-II censoring schemes are considered:

$$\boldsymbol{R}_E = \{10,0,0,0,0,0,0,0,0,0\}$$

$$\boldsymbol{R}_U = \{1,1,1,1,1,1,1,1,1,1\}$$

and

$$\boldsymbol{R}_L = \{0,0,0,0,0,0,0,0,0,10\}$$

representing early, approximately uniform, and late removal patterns, respectively. Each scheme satisfies

$$\sum_{i=1}^{m} R_i = n - m = 10$$

and therefore has the same overall censoring proportion. Consequently, differences in monitoring performance across the three configurations can be attributed primarily to the timing of withdrawals rather than to differences in the amount of censoring.

The robustness tuning parameter τ and the common CUSUM reference value k are examined over prespecified candidate sets. The case $\tau = 0$ corresponds to maximum-likelihood estimation and is retained as the likelihood-based non-robust benchmark, whereas positive values of τ correspond to robust MDPDE-based designs. For each combination of τ and censoring scheme, the finite-sample marginal reference distributions of the log-transformed scale and shape estimators, $F_{\alpha,0,\tau,\boldsymbol{R}}$ and $F_{\beta,0,\tau,\boldsymbol{R}}$, are estimated by Monte Carlo simulation under the nominal in-control model. These reference distributions are subsequently used to obtain the component-wise normal scores $Z_{\alpha,\tau,\boldsymbol{R}}$ and $Z_{\beta,\tau,\boldsymbol{R}}$ defined in Section 3.2.

For every candidate configuration, the Shewhart limit H_S and the common directional CUSUM limit H_C are jointly calibrated according to the procedure described in Section 3.4 and Section Allocation of the Shewhart and CUSUM Control Limits to achieve a target in-control average run length of approximately 370. The complete simulation configuration is summarized in Table 1.

Table 1. Simulation design and in-control configuration.

Design Component	Setting
In-control scale parameter	$\alpha_0 = 1$
In-control shape parameter	$\beta_0 = 2$
Initial number of units	$n = 20$
Number of observed failures	$m = 10$
Censoring proportion	50%
Early censoring	$\boldsymbol{R}_E = \{10,0,0,0,0,0,0,0,0,0\}$
Uniform censoring	$\boldsymbol{R}_U = \{1,1,1,1,1,1,1,1,1,1\}$
Late censoring	$\boldsymbol{R}_L = \{0,0,0,0,0,0,0,0,0,10\}$
DPD tuning parameter	$\tau \in \{0,0.05,0.10,0.20,0.30\}$
CUSUM reference value	$k \in \{0.25,0.50,0.75,1.00\}$
Target in-control ARL	$ARL_0^* = 370$
Reference-distribution replications	$B_0 = 50{,}000$
Calibration replications	$B_{cal} = 20{,}000$
Final verification replications	$B_{ver} = 50{,}000$

Note: The three censoring schemes have the same censoring proportion but differ in the timing of withdrawals, allowing the effect of the censoring pattern to be evaluated independently of the censoring level. The case $\tau = 0$ corresponds to the likelihood-based non-robust benchmark.

For each candidate monitoring configuration, the finite-sample reference distributions were constructed using $B_0 = 50{,}000$ independent Monte Carlo replications. Control-limit calibration was performed using $B_{cal} = 20{,}000$ replications, followed by an independent final verification based on $B_{ver} = 50{,}000$ additional in-control sequences. These simulation

stages were conducted independently so that the final in-control verification was not based on the same Monte Carlo realizations used for control-limit calibration.

4.2. Control-Limit Calibration and Design Selection

For each combination of the MDPDE tuning parameter τ, common CUSUM reference value k, and progressive Type-II censoring scheme, the control limits of the component-wise combined Shewhart–CUSUM chart are calibrated under the nominal in-control Weibull model. Before control-limit calibration, the finite-sample marginal reference distributions $F_{\alpha,0,\tau,\boldsymbol{R}}$ and $F_{\beta,0,\tau,\boldsymbol{R}}$ are estimated using $B_0 = 50{,}000$ independent progressively Type-II censored samples generated from the corresponding in-control reference model. These distributions are subsequently held fixed and used to transform the period-specific log-parameter estimates into the component-wise normal scores $Z_{\alpha,\tau,\boldsymbol{R}}$ and $Z_{\beta,\tau,\boldsymbol{R}}$.

The dependence between the two component-wise normal scores was also examined using 10,000 independent in-control reference samples. The Pearson correlations between Z_α and Z_β, and -0.452 under early, uniform, and late progressive censoring, respectively. The corresponding Kendall coefficients were 0.035, -0.147, and -0.298. Thus, the association was negligible under early censoring, moderately negative under uniform censoring, and more pronounced under late censoring. These findings show that marginal normal-score standardization does not imply independence between the two components and support the use of joint Monte Carlo calibration of the stopping rule rather than an independence-based control-limit approximation.

Following the allocation procedure described in Section Allocation of the Shewhart and CUSUM Control Limits, a prespecified grid of candidate Shewhart limits,

$$\boldsymbol{H}_S = \left\{H_S^{(1)}, H_S^{(2)}, \ldots, H_S^{(L)}\right\}$$

is considered for each combination of τ, k, and censoring scheme. Each candidate $H_S^{(l)}$ is first evaluated using the component-wise Shewhart mechanism alone, with

$$M_{t,\tau} = \max\{|Z_{\alpha,t,\tau}|, |Z_{\beta,t,\tau}|\}.$$

Candidates for which the Shewhart-only in-control ARL does not exceed the target value are excluded according to the feasibility condition described in Section Allocation of the Shewhart and CUSUM Control Limits.

For each retained $H_S^{(l)}$, the corresponding common directional CUSUM limit $H_C^{(l)}$ is determined numerically using $B_{cal} = 20{,}000$ Monte Carlo replications. The same $H_C^{(l)}$ is applied to $C^+_{\alpha,t,\tau}, C^-_{\alpha,t,\tau}, C^+_{\beta,t,\tau}, C^-_{\beta,t,\tau}$. A candidate pair $\left(H_S^{(l)}, H_C^{(l)}\right)$ is regarded as adequately calibrated when

$$\left|\hat{ARL}_0 - 370\right| \leq 5,$$

corresponding to the admissible interval

$$365 \leq \hat{ARL}_0 \leq 375.$$

This calibration is performed separately for every combination of

$$\tau \in \{0, 0.05, 0.10, 0.20, 0.30\}, \quad k \in \{0.25, 0.50, 0.75, 1.00\},$$

and each of the early, uniform, and late censoring schemes defined in Table 1. Consequently, all admissible candidate designs provide approximately equivalent nominal in-control protection, allowing their out-of-control performance to be compared on a common basis.

To obtain a monitoring design that is not optimized for a single alternative, the admissible configurations are evaluated over a common set of prespecified out-of-control scenarios.

Using the multiplicative shift representation defined in Equation (37), the design-selection set is specified as

$$\boldsymbol{S} = \left\{ \begin{matrix} (0.80,1),(0.90,1),(1.10,1),(1.20,1),(1,0.80),(1,0.90),(1,1.10), \\ (1,1.20),(0.90,0.90),(0.90,1.10),(1.10,0.90),(1.10,1.10) \end{matrix} \right\}.$$

Each ordered pair represents $(\rho_\alpha, \rho_\beta)$. The first four scenarios correspond to isolated scale changes, the next four to isolated shape changes, and the remaining four to simultaneous changes in both Weibull parameters. Both upward and downward changes are included because the component-wise signed normal scores and directional CUSUM statistics are specifically constructed to preserve and monitor the direction of parameter departure.

Before aggregate design comparison, each calibrated configuration is required to satisfy the minimum detection condition in Equation (39) for every scenario in $\boldsymbol{S}$. Thus, any design for which

$$ARL_{1,d}^{(s)} \geq ARL_{0,d}^{(s)}$$

for at least one core shift scenario is excluded from further consideration. This preliminary screening prevents inadequate detection under one type or direction of process change from being masked by favorable performance under other alternatives.

Among the remaining admissible configurations, the standardized relative detection score G_d defined in Equation (45) is calculated across all scenarios in $\boldsymbol{S}$. The design yielding the smallest G_d, subject to the prescribed in-control ARL requirement and the minimum detection condition, is selected as the final nominal configuration for the corresponding censoring scheme. This criterion favors a configuration with strong overall detection performance rather than one optimized for a single favorable alternative.

Finally, the in-control performance of each provisionally selected configuration is independently verified using $B_{ver} = 50{,}000$ additional Monte Carlo replications. The component-wise reference distributions and candidate control limits are held fixed during this verification stage. A configuration is retained as a final nominal design only if the independently verified in-control average run length remains within the prespecified interval $365 \leq \hat{ARL}_0 \leq 375$. If the provisionally selected configuration fails this verification criterion, the next-best admissible configurations, ordered according to G_d, are independently examined until a design satisfying the verification requirement is identified. Among the independently verified configurations, the design with the smallest G_d is retained. The resulting $\hat{ARL}_0$, $SDRL_0$, and MRL_0 values are reported as the final in-control performance measures. This independent verification step ensures that the nominal design selection is not driven by Monte Carlo variability arising during the initial control-limit calibration.

As shown in Table 2, all three selected configurations achieved independently verified in-control ARL values within the prespecified interval of 365–375. The selected reference value was $k = 0.25$ for all censoring schemes, whereas the optimal robustness level and control-limit allocation varied with the timing of progressive withdrawals. Early censoring favored a robust MDPDE-based configuration with $\tau = 0.20$, while the likelihood-based benchmark ($\tau = 0$) was selected under the uniform and late censoring schemes. These results indicate that the preferred degree of robustification may depend on the information structure induced by the progressive censoring pattern.

Table 2. Selected control-chart designs and independently verified in-control performance.

Censoring Scheme	τ^*	k^*	H_S^*	H_C^*	G_d	ARL_0	$SDRL_0$	MRL_0
Early	0.20	0.25	3.75	9.3342	1.1010	367.88	351.22	261
Uniform	0.00	0.25	4.25	9.2378	1.0856	371.46	350.89	265
Late	0.00	0.25	4.25	9.3921	1.0608	373.82	353.22	265

To assess whether the selected CUSUM reference value was influenced by the lower boundary of the original search grid, the design-selection procedure was repeated with the additional reference values k ∈ {0.05, 0.10, 0.15, 0.20} for all three censoring schemes. For each candidate, the control limits were recalibrated to an in-control ARL of approximately 370, and G_d was computed using a common scenario-wise normalization across all candidates with k ≤ 0.25 (Supplementary Table S1). Under every censoring scheme, G_d increased monotonically as k decreased below 0.25. Even at k = 0.20, the best attainable G_d exceeded that of the k = 0.25 design by 0.036–0.060, a difference well beyond the Monte Carlo uncertainty. Smaller reference values therefore do not improve overall detection performance, and the selection of k = 0.25 is not an artefact of the lower grid boundary.

4.3. Out-of-Control Performance Under Parameter Shifts

The out-of-control performance of the selected monitoring configurations was evaluated under the 12 prespecified shift scenarios used in the design-selection stage. These scenarios include isolated changes in the Weibull scale parameter, isolated changes in the shape parameter, and simultaneous changes in both parameters. For each censoring scheme, the final configuration reported in Table 2 was retained without further modification, and the run-length distribution was estimated using 20,000 independent Monte Carlo replications for each out-of-control scenario. The resulting ARL_1, $SDRL_1$, and MRL_1 values are reported in Table 3.

Table 3. Out-of-control run-length performance of the selected combined Shewhart–CUSUM charts.

		Early			**Uniform**			**Late**		
ρ_α	ρ_β	*ARL*	*SDRL*	*MRL*	*ARL*	*SDRL*	*MRL*	*ARL*	*SDRL*	*MRL*
0.80	1.00	9.80	2.54	10	9.69	2.55	9	10.53	2.99	10
0.90	1.00	23.87	10.65	22	24.17	10.75	22	27.05	12.82	24
1.10	1.00	23.42	12.19	21	25.23	12.63	23	28.84	15.20	25
1.20	1.00	9.93	4.05	10	11.03	3.94	10	11.89	4.46	11
1.00	0.80	12.74	5.15	12	14.90	6.27	14	19.03	8.70	17
1.00	0.90	39.44	24.18	33	45.01	29.17	37	51.62	34.66	42
1.00	1.10	49.58	31.18	42	64.01	44.65	52	79.92	60.72	63
1.00	1.20	19.45	7.93	18	23.61	10.35	22	29.07	14.04	26
0.90	0.90	22.36	10.04	20	20.21	8.36	19	22.26	9.53	21
0.90	1.10	21.65	8.10	20	24.73	10.51	23	28.94	13.50	26
1.10	0.90	18.15	8.79	17	23.51	12.54	21	27.70	16.43	24
1.10	1.10	23.85	10.65	22	22.86	9.30	21	26.03	11.05	24

All selected configurations satisfied the minimum detection requirement in Equation (39), with $ARL_1 < ARL_0$ under every considered alternative. The magnitude of the run length generally decreased as the parameter change became more pronounced, indicating faster detection of larger departures from the in-control Weibull model. Moreover, both upward and downward changes in the scale and shape parameters were successfully detected, consistent with the directional structure of the component-wise normal scores and the four CUSUM components.

The results in Table 3 demonstrate that the proposed monitoring structure is responsive to all investigated forms of parameter change. For isolated scale shifts, the selected charts exhibit particularly rapid detection, with ARL_1 values between approximately 9.7 and 11.9 for 20% changes and between approximately 23.4 and 28.8 for 10% changes. The relatively similar run-length values for downward and upward scale shifts indicate that the component-wise signed monitoring structure provides effective detection in both directions.

Shape changes are generally more difficult to detect than scale changes of comparable proportional magnitude. For a 10% decrease in the shape parameter, ARL_1 ranges from 39.44

to 51.62 across the three censoring schemes, whereas a 10% increase produces values from 49.58 to 79.92. Nevertheless, all of these values remain substantially below the corresponding in-control ARL_s reported in Table 2. Increasing the magnitude of the shape shift to 20% substantially reduces the run lengths, confirming the expected improvement in detection as the departure from the reference model becomes larger.

The simultaneous-shift scenarios also yield consistently low run lengths. In particular, the combined changes considered in the study produce ARL_1 values below 30 for all three censoring schemes. Overall, early censoring tends to provide the fastest detection in several isolated shape-change scenarios, whereas the differences among censoring schemes are less pronounced for large scale changes. These findings demonstrate that the timing of progressive withdrawals affects monitoring sensitivity even when the overall censoring proportion is held fixed.

To assess the influence of the initial monitoring state, a supplementary conditional steady-state analysis was performed using a 50-period in-control warm-up. Monitoring sequences that signaled during the warm-up were excluded, while the accumulated directional CUSUM statistics of the surviving sequences were retained when the parameter shift was introduced. Across the three progressive censoring schemes, approximately 91–92% of the sequences survived the warm-up period. The resulting post-change average run lengths were generally shorter than their zero-state counterparts. For example, under late censoring, the average run length for a 10% increase in the Weibull shape parameter decreased from 79.92 to 73.32. Similar reductions were observed under early and uniform censoring. These findings indicate that the proposed monitoring procedure retains its detection capability when parameter changes occur after an initial in-control monitoring period, rather than only when monitoring begins from zero CUSUM states. The detailed results are provided in Supplementary Table S2.

4.4. Estimator-Level Robustness Assessment

Before evaluating the robustness of the complete monitoring procedure, an additional estimator-level simulation was conducted to examine how positive DPD tuning affects the finite-sample behavior of the Weibull parameter estimates. The same sampling configuration used in the main simulation study, with $n = 20$ and $m = 10$, was considered under the early, uniform, and late progressive Type-II censoring schemes. The likelihood-based estimator $\tau = 0$ and the MDPDE with $\tau = 0.20$ were compared under the nominal Weibull model $(\alpha,\beta) = (1,2)$, as well as under 5% lower-tail and upper-tail contamination. The contamination mechanisms were identical to those subsequently used in the chart-level robustness analysis. For each configuration, 10,000 independent replications were generated, and estimator performance was evaluated in terms of bias and mean squared error (MSE). No numerical convergence failures occurred in any of the investigated configurations.

The results reveal a clear parameter-specific efficiency–robustness pattern. For the Weibull shape parameter, the MDPDE produced smaller MSE than the likelihood-based estimator in every investigated combination of censoring scheme and contamination condition. The reduction in shape-parameter MSE ranged from approximately 16.6% to 35.6%. For example, under late censoring, the shape-parameter MSE decreased from 0.9123 to 0.5875 under the uncontaminated model and from 0.8835 to 0.5711 under upper-tail contamination. Similar improvements were observed under early and uniform censoring.

The behavior of the scale estimator was less uniform. Positive DPD tuning provided a modest MSE reduction under early censoring with lower-tail contamination, from 0.0280 to 0.0268, but the likelihood-based estimator generally retained smaller scale-parameter MSE in the remaining configurations. The difference was particularly pronounced under upper-tail contamination, where the MDPDE scale MSE exceeded that of the likelihood-

based estimator for all three censoring schemes. Thus, positive DPD tuning should not be interpreted as providing uniform improvement for every Weibull parameter. Rather, the results indicate a parameter-specific trade-off in which the principal finite-sample gain of $\tau = 0.20$ is associated with stabilization of shape estimation, whereas scale estimation may incur an efficiency cost.

This estimator-level behavior provides additional context for the chart-level contamination results considered in the following subsection. Because the combined monitoring procedure is driven jointly by the scale and shape estimates, improved robustness of one parameter component does not necessarily imply uniformly improved run-length behavior of the complete chart. The interaction between estimator behavior, contamination direction, and progressive withdrawal pattern therefore remains relevant when assessing the robustness of the overall monitoring procedure.

4.5. Detection of Large Transient Shifts

The comparison in Table 4 indicates that the combined and CUSUM-only procedures have very similar performance for the moderate persistent shifts considered in the main simulation design. To examine the specific contribution of the Shewhart component, an additional experiment was therefore conducted for large transient departures. Scale and shape multipliers of 0.50 and 2.00 were considered, with each departure lasting for either one or three monitoring periods before the process returned to the in-control model. Detection performance was measured by the conditional probability of generating a signal before the transient departure ended, given that no signal had occurred during the preceding 20 in-control periods. The experiment was based on 2000 monitoring sequences for each configuration, while the finite-sample normal-score reference distributions retained the original $B_0 = 50{,}000$ Monte Carlo size. The directional CUSUM statistics were initialized at zero only at the beginning of the 20-period in-control monitoring phase. At the onset of the transient shift, the accumulated CUSUM values of the surviving monitoring sequences were retained without resetting, so the reported detection probabilities reflect the monitoring state reached after the in-control prehistory.

Table 4. Estimator-level MSE comparison between likelihood-based estimation and MDPDE ($\tau = 0.20$).

Scheme	Condition	MLE MSE_{α}	MDPDE MSE_{α}	Δ_{α} (%)	MLE MSE_{β}	MDPDE MSE_{β}	Δ_{β} (%)
Early	Clean	0.0253	0.0264	−4.46	0.3668	0.2799	23.71
	Lower-tail	0.0280	0.0268	4.41	0.2810	0.2274	19.06
	Upper-tail	0.0650	0.0927	−42.59	0.3490	0.2910	16.62
Uniform	Clean	0.0265	0.0306	−15.61	0.6190	0.4379	29.25
	Lower-tail	0.0302	0.0319	−5.70	0.4151	0.3123	24.76
	Upper-tail	0.0358	0.0486	−35.74	0.5584	0.4014	28.12
Late	Clean	0.0306	0.0391	−27.57	0.9123	0.5875	35.61
	Lower-tail	0.0355	0.0419	−17.90	0.6674	0.4547	31.87
	Upper-tail	0.0334	0.0480	−43.49	0.8835	0.5711	35.36

Note: $\Delta = 100(1 - \frac{MSE_{MDPDE}}{MSE_{MLE}})$. Positive values indicate an MSE reduction under MDPDE, whereas negative values indicate a larger MSE relative to likelihood-based estimation. Results are based on 10,000 Monte Carlo replications for each configuration.

In Table 5, the transient-shift results reveal a complementary but nonuniform contribution of the two monitoring components. For one-period large departures, the CUSUM-only procedure exhibited limited immediate detection capability, whereas the Shewhart pathway substantially improved the detection probability of the combined procedure relative to CUSUM-only monitoring. For example, under early censoring with a scale multiplier of 2.00, the one-period detection probability increased from 0.031 for the CUSUM-only procedure

to 0.878 for the combined procedure, with 0.876 attributable to activation of the Shewhart pathway. Similar behavior was observed under uniform and late censoring, where the corresponding combined detection probabilities were 0.767 and 0.499, compared with 0.032 and 0.021 for CUSUM-only monitoring, respectively.

Table 5. Detection probabilities for large transient shifts under the single-component and combined monitoring schemes.

Censoring Scheme	Transient Shift	Duration L	Shewhart-Only P_S	CUSUM-Only P_C	Combined P_{Comb}	Shewhart Path $P_{Comb,S}$	CUSUM Path $P_{Comb,C}$
Early	Scale 0.50	1	0.677	0.013	0.228	0.217	0.015
		3	0.966	0.794	0.806	0.486	0.451
	Scale 2.00	1	0.920	0.031	0.878	0.876	0.030
		3	0.999	0.982	0.999	0.996	0.050
	Shape 0.50	1	0.584	0.022	0.435	0.428	0.019
		3	0.931	0.600	0.856	0.800	0.155
	Shape 2.00	1	0.168	0.007	0.070	0.062	0.008
		3	0.436	0.242	0.307	0.163	0.169
Uniform	Scale 0.50	1	0.679	0.018	0.234	0.220	0.020
		3	0.967	0.799	0.846	0.502	0.483
	Scale 2.00	1	0.921	0.032	0.767	0.763	0.028
		3	0.999	0.973	0.995	0.979	0.116
	Shape 0.50	1	0.510	0.017	0.178	0.167	0.017
		3	0.880	0.465	0.617	0.409	0.304
	Shape 2.00	1	0.090	0.006	0.007	0.002	0.006
		3	0.245	0.134	0.132	0.007	0.126
Late	Scale 0.50	1	0.560	0.019	0.199	0.191	0.013
		3	0.905	0.724	0.756	0.460	0.422
	Scale 2.00	1	0.763	0.021	0.499	0.492	0.018
		3	0.990	0.900	0.957	0.859	0.255
	Shape 0.50	1	0.333	0.014	0.138	0.130	0.013
		3	0.707	0.272	0.466	0.321	0.194
	Shape 2.00	1	0.034	0.005	0.004	0.000	0.004
		3	0.095	0.084	0.073	0.000	0.073

However, the combined procedure did not uniformly outperform the Shewhart-only chart. Under early censoring with a one-period scale multiplier of 0.50, the detection probability was 0.677 for Shewhart-only monitoring but only 0.228 for the combined procedure. This difference highlights an important trade-off: when the procedures are calibrated to comparable in-control performance, incorporating the CUSUM component may require more conservative signaling limits for the Shewhart pathway, thereby reducing sensitivity to certain isolated large departures. Consequently, the combined procedure should not be interpreted as uniformly superior to its individual components, but rather as a monitoring strategy that balances immediate detection and cumulative evidence.

When the transient departure persisted for three monitoring periods, the contribution of the CUSUM mechanism increased substantially. For example, under early censoring with a scale multiplier of 0.50, the CUSUM-only detection probability increased from 0.013 for a one-period departure to 0.794 for a three-period departure, while the combined detection probability increased from 0.228 to 0.806. Within the combined procedure, the Shewhart and CUSUM pathways were activated with probabilities of 0.486 and 0.451, respectively. A similar pattern was observed under uniform and late censoring. These results indicate that the Shewhart component provides an immediate signaling pathway for certain large

short-lived departures, whereas the directional CUSUM component becomes increasingly influential as the change persists.

Detection remained more difficult for upward changes in the Weibull shape parameter, particularly under uniform and late censoring. For a one-period shape multiplier of 2.00, the combined detection probabilities were 0.070, 0.007, and 0.004 under early, uniform, and late censoring, respectively. This finding is consistent with the persistent-shift experiments, which also indicated reduced sensitivity to upward shape changes. Overall, the results demonstrate that the practical value of the combined Shewhart–CUSUM framework lies in providing complementary signaling mechanisms rather than uniformly improved detection performance. Its relative effectiveness depends on the magnitude, duration, and direction of the parameter change, as well as the progressive censoring configuration.

4.6. Robustness Under Contamination

The robustness of the proposed monitoring framework was further investigated by introducing atypical observations into the progressively Type-II censored lifetime samples. To isolate the effect of robust parameter estimation, the likelihood-based monitoring scheme corresponding to $\tau = 0$ was compared with a common robust MDPDE-based counterpart using $\tau = 0.20$. For each censoring scheme, the comparison charts were calibrated to provide approximately equivalent nominal in-control protection before contamination was introduced. Their nominal reference distributions and control limits were subsequently held fixed throughout the contamination experiments; no contamination-specific recalibration was performed.

Two complementary contamination settings were considered. The first corresponds to in-control contamination, in which the underlying Weibull process remains at $(\alpha,\beta) = (\alpha_0,\beta_0)$, but a proportion of the individual lifetimes is generated from an atypical distribution. This setting evaluates the extent to which contamination alone distorts the nominal run-length behavior and induces premature signals. The second corresponds to out-of-control contamination, in which the same contamination mechanism is superimposed on a genuine change in the Weibull parameters. This setting evaluates whether sensitivity to an actual process change is preserved when atypical observations are simultaneously present.

Let ε denote the contamination proportion. The contaminated lifetime distribution is represented by the mixture

$$F_\varepsilon(x;\alpha,\beta,c) = (1-\varepsilon)F_W(x;\alpha,\beta) + \varepsilon F_W(x;c\alpha,\beta), \tag{47}$$

where $F_\varepsilon(x;\alpha,\beta)$ denotes the Weibull distribution associated with the underlying process state and $c > 0$ is a contamination-scale multiplier. The nominal Weibull model is recovered when $\varepsilon = 0$.

To investigate contamination of different severities,

$$\varepsilon \in \{0.01, 0.05, 0.10\}$$

was considered. Two contamination directions were examined. Lower-tail contamination was generated using $c = 1/3$ whereas upper-tail contamination was generated using $c = 3$.

Thus, the contamination mechanism produces atypically short or atypically long lifetimes while leaving the Weibull shape parameter unchanged. Considering both directions is particularly relevant under progressive Type-II censoring because the effect of atypical lifetimes on the observed failure sequence may depend on both their position in the lifetime distribution and the timing of progressive withdrawals.

For each censoring scheme, the likelihood-based benchmark and the robust MDPDE-based configuration were first matched to provide approximately equivalent nominal in-control ARL values. The resulting nominal ARLs were 367.28 and 367.76 for the likelihood-based and robust configurations under early censoring, 368.75 and 371.45 under uniform

censoring, and 369.75 and 368.60 under late censoring, respectively. Consequently, differences observed after contamination are not primarily attributable to unequal nominal false-alarm protection.

Under in-control contamination, the underlying process parameters remain at (α_0,β_0), and the resulting ARL, SDRL, and MRL values are used to quantify the stability of the monitoring procedures. Rather than assuming a priori that robust estimation must improve the complete monitoring procedure under every contamination mechanism, the comparison is used to determine whether and under which combinations of censoring pattern, contamination direction, and contamination level the positive MDPDE tuning parameter provides greater stability than the likelihood-based benchmark.

For the out-of-control robustness analysis, two moderate changes that were relatively challenging in the nominal simulation were considered. These correspond to a 10% increase in the scale parameter,

$$(\rho_\alpha,\rho_\beta) = (1.10, 1.00)$$

and a 10% increase in the shape parameter,

$$(\rho_\alpha,\rho_\beta) = (1.00, 1.10) \tag{48}$$

The same lower- and upper-tail contamination mechanisms and contamination proportions were superimposed on these shifted processes. The resulting contaminated ARL_1, $SDRL_1$, and MRL_1 values were compared with their uncontaminated counterparts. This analysis therefore evaluates both resistance to contamination-induced false alarms under the in-control state and preservation of detection capability when a genuine parameter change and contamination occur simultaneously.

Table 6 shows that contamination affects the in-control run-length distribution substantially, although the magnitude and direction of the effect depend on the progressive censoring pattern and the contamination mechanism. Under uniform censoring, the MDPDE-based chart provided greater stability than the likelihood-based benchmark for all three lower-tail contamination levels. For example, at $\varepsilon = 0.05$, the in-control ARL decreased from its nominal level to 46.73 under maximum-likelihood monitoring, whereas the corresponding MDPDE-based ARL remained at 96.21. A similar advantage was observed under late censoring at moderate and severe lower-tail contamination, where the MDPDE-based ARLs were 94.72 and 33.07 compared with 40.79 and 24.16 for the likelihood-based benchmark.

The behavior under upper-tail contamination was more heterogeneous. At the mild contamination level $\varepsilon = 0.01$, the MDPDE-based configurations under uniform and late censoring retained ARLs of 403.31 and 414.02, respectively, compared with 337.02 and 321.25 for the likelihood-based charts. However, this advantage did not persist uniformly as the contamination proportion increased. Under early censoring, the likelihood-based chart generally retained larger in-control ARLs than its MDPDE-based counterpart, particularly under mild contamination. These findings indicate that the protection provided by robust parameter estimation is not uniform across progressive censoring structures and contamination directions.

More generally, both monitoring approaches become increasingly sensitive to contamination as ε increases, as reflected by the substantial reduction in ARL for several severe contamination scenarios. Nevertheless, the MDPDE-based construction reduces this deterioration in a number of important settings, most notably under lower-tail contamination with uniform and late censoring. The results therefore suggest that the robustness of the complete sequential monitoring procedure is determined jointly by the estimation method, the progressive withdrawal pattern, and the location and severity of contamination.

Figure 2 provides a visual comparison of the effect of contamination on the in-control behavior of the likelihood-based and MDPDE-based monitoring schemes. Consistent with Table 6, the effect of robustification varies according to both the progressive censoring

pattern and the direction of contamination. Under uniform censoring, the MDPDE-based chart generally provides greater resistance to lower-tail contamination, with the difference becoming particularly pronounced at $\varepsilon = 0.05$. A similar advantage is observed under late censoring for moderate and severe lower-tail contamination.

Table 6. In-control run-length performance under lower- and upper-tail contamination.

			MLE			MDPDE		
Scheme	**Contamination**	$\boldsymbol{\varepsilon}$	**ARL**	**SDRL**	**MRL**	**ARL**	**SDRL**	**MRL**
Early	Lower-tail	0.01	303.07	282.32	216	137.35	126.10	99
		0.05	61.08	42.22	50	80.69	65.26	62
		0.10	30.05	15.62	27	26.10	13.06	23
	Upper-tail	0.01	96.56	85.12	71	46.90	40.39	36
		0.05	25.04	16.85	21	13.35	10.26	11
		0.10	8.56	4.74	8	7.13	5.54	6
Uniform	Lower-tail	0.01	237.90	217.26	173	276.85	262.35	195
		0.05	46.73	29.03	39	96.21	75.70	73.5
		0.10	26.61	12.28	24	31.88	16.96	28
	Upper-tail	0.01	337.02	316.46	240	403.31	387.63	283
		0.05	94.20	82.34	70	55.80	43.44	44
		0.10	21.13	12.26	19	21.64	12.53	19
Late	Lower-tail	0.01	215.78	200.19	154	197.11	176.10	143
		0.05	40.79	24.17	35	94.72	73.56	73
		0.10	24.16	11.34	22	33.07	16.92	30
	Upper-tail	0.01	321.25	301.94	228	414.02	393.65	295
		0.05	153.08	134.80	113	101.47	84.97	76
		0.10	43.33	27.39	36	36.39	21.07	31

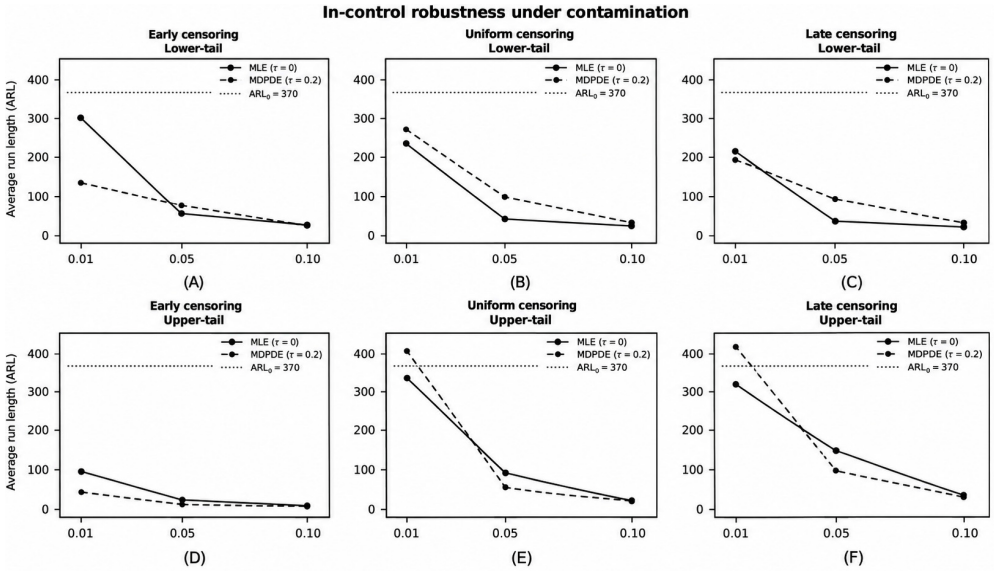

Figure 2. In-control average run length (ARL) under lower- and upper-tail contamination for the likelihood-based $\tau = 0$ and MDPDE-based $\tau = 0.20$ monitoring schemes under early, uniform, and late progressive Type-II censoring. Panels (**A**–**C**) correspond to lower-tail contamination, whereas panels (**D**–**F**) correspond to upper-tail contamination. The horizontal dotted line represents the nominal target $ARL_0^* = 370$.

In contrast, under early censoring, the MDPDE-based procedure does not consistently retain a larger in-control ARL than the likelihood-based benchmark.

For upper-tail contamination, the interaction with the censoring pattern is even more apparent. At $\varepsilon = 0.01$, the MDPDE-based chart under uniform and late censoring remains close to, or slightly above, the nominal ARL_0, reference level, whereas the corresponding likelihood-based ARLs decrease to a greater extent. However, this advantage diminishes as the contamination proportion increases, and both procedures exhibit substantial reductions in ARL under moderate and severe contamination. Under early censoring, upper-tail contamination produces a particularly rapid reduction in ARL for both procedures.

Overall, Figure 2 demonstrates that the effect of MDPDE-based robustification is not uniform across contamination settings. Rather, the degree of protection depends jointly on the contamination direction, contamination proportion, and progressive withdrawal pattern. This finding highlights the importance of evaluating robustness within the specific censoring structure under which the monitoring procedure is implemented.

Table 7 demonstrates that the effect of contamination on out-of-control detection depends strongly on the type of parameter change, contamination direction, and progressive censoring pattern. For scale shifts, both monitoring procedures generally retained relatively short run lengths under contamination. Under early censoring, the MDPDE-based chart yielded smaller ARL_1 values than the likelihood-based benchmark in most scale-shift contamination scenarios, although this advantage was not consistently observed under uniform and late censoring. These findings indicate that the benefits of robustification are conditional rather than uniform.

More pronounced differences were observed for shape shifts, particularly under lower-tail contamination. For a 10% increase in the shape parameter combined with 5% lower-tail contamination, the ARL_1 values under uniform censoring increased to 258.04 and 307.21 for the MLE-based and MDPDE-based charts, respectively. Under late censoring, the corresponding values were 255.42 and 449.41, indicating approximately 76% longer detection delay for the MDPDE-based procedure. This deterioration is consistent with a possible masking effect, whereby lower-tail contamination counteracts the influence of a relatively small shape-parameter increase on the monitoring statistics. Since the joint-density-based MDPDE responds to the complete censored sample configuration, robustification may attenuate the signal associated with a genuine parameter change rather than selectively suppressing contamination. However, the reported run-length results alone do not establish the precise masking mechanism.

A supplementary examination of the saved standardized score distributions provides evidence of an offsetting mechanism. Under late censoring with a 10% increase in the shape parameter, the mean shape score decreased from 0.294 to 0.037 for the likelihood-based procedure and from 0.290 to -0.040 for the MDPDE-based procedure after introducing 5% lower-tail contamination. Thus, contamination largely counteracted the positive score displacement associated with the genuine shape increase. The corresponding mean scale scores under contamination were -0.118 and -0.044, respectively, indicating a smaller negative scale displacement for the MDPDE-based procedure. These findings provide evidence of an offsetting mechanism, whose consequences for directional alarm behavior are examined below. To further examine the detection mechanism, the direction and parameter associated with the first alarm were recorded using the calibrated monitoring limits. Under late censoring with a 10% shape increase and 5% lower-tail contamination, the likelihood-based and MDPDE-based procedures yielded mean run lengths of 254.07 and 436.67, respectively, consistent with the corresponding results in Table 7. However, only 27.33% and 17.67% of the first alarms, respectively, were attributable to the upward shape component. For the likelihood-based procedure, downward scale alarms predominated (60.73%), whereas the MDPDE-based procedure exhibited a higher proportion of downward shape alarms (45.47%). These findings indicate that contamination substantially alters the

direction of the monitoring signals and that the longer run length of the MDPDE-based procedure cannot be interpreted solely as inferior detection of the genuine shape increase.

Table 7. Out-of-control ARL performance of the likelihood-based $\tau = 0$ and robust MDPDE-based $\tau = 0.20$ monitoring schemes under contamination.

Scheme	Shift	Contam.	ε	MLE ARL_1	MDPDE ARL_1
Early	Scale +10%	Lower	0.01	34.49	28.89
			0.05	38.65	29.11
			0.10	30.53	28.79
		Upper	0.01	16.93	17.71
			0.05	9.39	8.26
			0.10	5.58	4.36
	Shape +10%	Lower	0.01	36.37	42.53
			0.05	62.26	182.08
			0.10	53.38	51.31
		Upper	0.01	67.00	47.50
			0.05	18.86	9.64
			0.10	9.02	7.38
Uniform	Scale +10%	Lower	0.01	23.48	33.92
			0.05	33.63	50.01
			0.10	25.07	33.92
		Upper	0.01	24.72	29.52
			0.05	13.30	14.08
			0.10	9.05	11.02
	Shape +10%	Lower	0.01	90.63	76.99
			0.05	258.04	307.21
			0.10	46.43	46.12
		Upper	0.01	44.28	63.85
			0.05	56.35	48.60
			0.10	23.59	32.88
Late	Scale +10%	Lower	0.01	26.61	43.15
			0.05	37.49	36.37
			0.10	34.34	59.90
		Upper	0.01	26.67	30.81
			0.05	15.33	17.59
			0.10	13.23	19.94
	Shape +10%	Lower	0.01	87.83	57.74
			0.05	255.42	449.41
			0.10	48.36	50.88
		Upper	0.01	86.34	140.17
			0.05	38.31	79.73
			0.10	39.16	49.34

Conversely, MDPDE-based monitoring improved detection in selected contamination scenarios. Under late censoring with mild lower-tail contamination and a 10% shape increase, ARL_1 decreased from 87.83 for the likelihood-based chart to 57.74 for the MDPDE-based chart. Under early censoring with upper-tail contamination, the robust configuration also produced shorter run lengths in several shape-shift scenarios. These contrasting results suggest that the interaction between contamination and parameter changes may either reinforce or weaken the monitoring signal, depending on the censoring structure and contamination severity.

Importantly, contaminated out-of-control run lengths should be interpreted jointly with the corresponding contaminated in-control performance. A shorter ARL_1 does not

necessarily indicate superior detection when contamination also substantially reduces ARL_0, since the resulting increase in alarm frequency may occur even without a genuine process change. Conversely, a longer ARL_1 may reflect reduced sensitivity to the parameter shift, although its interpretation also depends on the associated in-control false-alarm behavior. Therefore, the robustness of the proposed monitoring framework should be assessed in terms of both false-alarm control and change-detection capability. Overall, the results support treating MDPDE as an optional, conditionally beneficial component rather than a uniformly superior alternative to maximum likelihood estimation.

Figure 3 shows that the effect of contamination on the detection of a 10% shape increase depends strongly on both the contamination direction and the progressive censoring pattern. Lower-tail contamination, particularly at $\varepsilon = 0.05$, produces a marked increase in ARL_1 under uniform and late censoring, whereas upper-tail contamination generally results in shorter detection delays. The relative performance of the MDPDE- and likelihood-based schemes also varies across the considered settings, confirming that the benefit of robustification is contamination- and censoring-dependent rather than uniform.

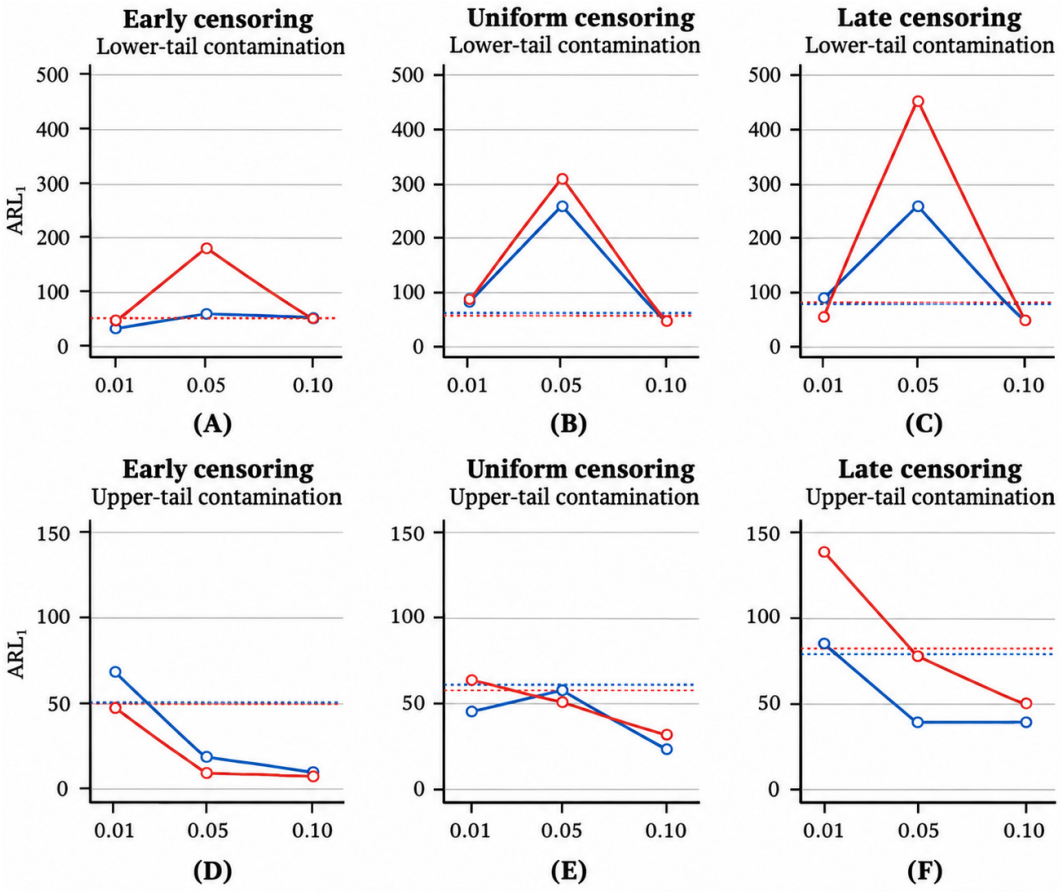

Figure 3. Out-of-control average run length ARL_1 for a 10% increase in the Weibull shape parameter under lower- and upper-tail contamination.

Overall, the contamination results demonstrate that positive MDPDE tuning does not provide uniformly improved monitoring performance. Its effects depend on the progressive censoring scheme, contamination direction, and affected Weibull parameter. In particular, improvements in parameter-estimation accuracy do not necessarily translate into better false-alarm protection or faster detection of genuine process changes. Moreover, contaminated out-of-control run lengths should be interpreted alongside the corresponding contaminated in-control performance, since shorter detection delays may also reflect increased sensitivity

to contamination rather than improved discrimination of actual process changes. These findings support treating MDPDE as an optional, condition-dependent estimation component rather than a universally superior alternative to maximum likelihood estimation.

4.7. Comparison with Single-Component Monitoring Schemes

To assess the contribution of combining the Shewhart and CUSUM mechanisms, the selected combined monitoring scheme was compared with its corresponding Shewhart-only and CUSUM-only counterparts. For a fair comparison, each single-component procedure was calibrated separately to achieve approximately the same nominal in-control average run length, $ARL_0 = 370$, as the combined chart. The same component-wise normal-score statistics, robustness tuning parameter, progressive Type-II censoring scheme, and finite-sample reference distributions were retained across the competing procedures, so that differences in out-of-control performance could be attributed primarily to the monitoring structure rather than to differences in parameter estimation or nominal false-alarm protection. The Shewhart-only procedure generated a signal whenever

$$M_{t,\tau} > H_S^{(S)},$$

where $H_S^{(S)}$ denotes the independently calibrated Shewhart-only control limit. The CUSUM-only procedure retained the four directional statistics $C^+_{\alpha,t,\tau}, C^-_{\alpha,t,\tau}, C^+_{\beta,t,\tau}, C^-_{\beta,t,\tau}$, and generated a signal whenever

$$C^{max}_{t,\tau} > H_C^{(C)}$$

where $H_C^{(C)}$ denotes the independently calibrated CUSUM-only control limit. In each case, the control limit was selected to provide $ARL_0 \approx 370$.

The three monitoring structures were subsequently evaluated under the same 12 scale, shape, and simultaneous parameter-shift scenarios considered in Section 4.3. Their ARL_1 values were compared to determine whether the combined structure preserves the rapid response of the Shewhart component to relatively large instantaneous departures while retaining the accumulation capability of the directional CUSUM component for smaller persistent changes.

The results in Table 8 demonstrate distinct detection characteristics of the Shewhart and CUSUM components under persistent parameter changes. The Shewhart-only procedure generally exhibits lower sensitivity to moderate and sustained shifts, particularly small increases in the Weibull shape parameter. For a 10% increase in the shape parameter, the Shewhart-only ARL_1 values are 476.19, 555.56, and 416.67 under early, uniform, and late censoring, respectively, compared with 49.39, 63.24, and 79.74 for the combined procedure. Similar improvements relative to Shewhart-only monitoring are observed for several moderate scale and simultaneous parameter changes. These results indicate that the principal improvement arises from the cumulative detection mechanism rather than from combining the two components itself.

In contrast, the CUSUM-only and combined procedures exhibit very similar run-length performance across most persistent-shift scenarios. For example, under early censoring with a 20% increase in the scale parameter, ARL_1 decreases only slightly from 10.36 for CUSUM-only monitoring to 9.91 for the combined procedure. Such small differences should not be interpreted as evidence of a general performance advantage without accounting for Monte Carlo uncertainty. The combined procedure therefore largely preserves CUSUM-like sensitivity to sustained changes, while the Shewhart component provides an additional immediate signaling pathway whose contribution depends on the magnitude and direction of the shift.

Table 8. Comparison of the combined Shewhart–CUSUM scheme with the corresponding single-component monitoring procedures.

Scheme	ρ_α	ρ_β	Shewhart-Only	CUSUM-Only	Combined
Early	0.80	1.00	55.87	9.73	9.81
	0.90	1.00	196.08	23.79	23.97
	1.10	1.00	79.37	23.83	23.51
	1.20	1.00	20.92	10.36	9.91
	1.00	0.80	34.13	13.11	12.91
	1.00	0.90	121.95	39.41	39.13
	1.00	1.10	476.19	49.51	49.39
	1.00	1.20	222.22	19.26	19.40
	0.90	0.90	94.34	22.29	22.62
	0.90	1.10	217.39	21.47	21.67
	1.10	0.90	41.15	18.97	18.29
	1.10	1.10	156.25	23.50	23.84
Uniform	0.80	1.00	62.11	9.75	9.76
	0.90	1.00	232.56	24.32	24.23
	1.10	1.00	131.58	25.44	25.30
	1.20	1.00	29.33	11.10	11.02
	1.00	0.80	38.31	15.15	14.92
	1.00	0.90	136.99	45.35	44.65
	1.00	1.10	555.56	63.82	63.24
	1.00	1.20	333.33	23.77	23.61
	0.90	0.90	93.46	20.29	20.11
	0.90	1.10	303.03	24.86	25.05
	1.10	0.90	56.18	23.97	23.89
	1.10	1.10	196.08	23.03	22.94
Late	0.80	1.00	53.76	10.43	10.52
	0.90	1.00	270.27	27.20	27.43
	1.10	1.00	185.19	28.86	29.03
	1.20	1.00	50.25	12.06	11.94
	1.00	0.80	55.56	19.22	19.09
	1.00	0.90	147.06	52.00	52.50
	1.00	1.10	416.67	80.28	79.74
	1.00	1.20	384.62	28.82	28.70
	0.90	0.90	90.09	22.14	22.44
	0.90	1.10	434.78	28.89	29.20
	1.10	0.90	80.00	28.11	27.47
	1.10	1.10	312.50	26.01	26.16

Taken together with the transient-shift results in Table 8, these findings highlight a trade-off between immediate detection and cumulative monitoring. The combined Shewhart–CUSUM framework offers complementary signaling mechanisms but does not uniformly outperform either individual procedure. Its practical advantage lies in accommodating different types of process changes within a single monitoring structure, rather than guaranteeing shorter run lengths or higher detection probabilities in every scenario. Accordingly, claims of comparative superiority are restricted to the Shewhart-only and CUSUM-only procedures evaluated under the specified censoring schemes and calibration settings.

In Table 9, to further evaluate the diagnostic interpretation of component-wise monitoring, the parameter and direction associated with the first alarm were examined under isolated 10% increases in the Weibull scale and shape parameters. Using 1500 monitoring sequences per configuration, the proportion of first alarms corresponding to the truly shifted parameter and direction ranged from 95.27% to 98.47% across the three censoring schemes and the MLE- and MDPDE-based procedures. Incorrect or opposite-direction

signals were also observed, with the highest incorrect attribution rate reaching 4.73% under late censoring. These findings indicate that the component-wise statistics provide useful parameter-specific diagnostic information under the investigated uncontaminated, single-parameter shifts, but do not guarantee identification of the true source of a distributional change. The reported proportions are conditional on an alarm occurring and were obtained by resampling the previously generated score distributions.

Table 9. First-alarm attribution under isolated Weibull parameter shifts.

Censoring Scheme	Method	Parameter Shift	Correct Attribution (%)	Incorrect Attribution (%)	Multiple Signals (%)
Early	MLE	Scale +10%	97.87	2.07	0.07
		Shape +10%	98.40	1.53	0.07
	MDPDE	Scale +10%	98.00	1.87	0.13
		Shape +10%	98.47	1.53	0.00
Uniform	MLE	Scale +10%	98.40	1.47	0.13
		Shape +10%	97.20	2.73	0.07
	MDPDE	Scale +10%	97.33	2.47	0.20
		Shape +10%	96.80	3.07	0.13
Late	MLE	Scale +10%	96.87	3.00	0.13
		Shape +10%	95.27	4.73	0.00
	MDPDE	Scale +10%	96.93	3.00	0.07
		Shape +10%	96.20	3.53	0.27

Figure 4 illustrates the complementary roles of the Shewhart and directional CUSUM components. Following the 10% increase in the Weibull shape parameter at $t = 31$, neither component-wise normal score exceeds the Shewhart limit, and consequently no Shewhart signal is generated. In contrast, the positive shape-direction CUSUM begins to accumulate sustained evidence after the change and crosses the CUSUM control limit at $t = 54$, producing the first combined-chart signal. The negative shape CUSUM remains close to zero, while the scale-direction CUSUMs exhibit no comparable sustained accumulation. This behavior illustrates both the directional interpretability of the proposed component-wise construction and the role of the CUSUM component in detecting moderate persistent parameter changes.

4.8. Sensitivity to Single-Lifetime Contamination

To investigate the observation-level sensitivity of the proposed joint-density-based MDPDE, a numerical contamination experiment was conducted under early, uniform, and late progressive Type-II censoring schemes. For each scheme, 1500 independent samples were generated from the Weibull distribution with baseline parameters $\alpha_0 = 1$ and $\beta_0 = 2$, using $n = 20$ units and $m = 10$ observed failures. In each replication, one of the original lifetimes was randomly selected and replaced by a contamination value z ranging from 0.01 to 20. The remaining lifetimes and the random withdrawal mechanism were held fixed to ensure a paired comparison between the uncontaminated and contaminated samples. The analysis considered $\tau \in \{0, 0.05, 0.10, 0.20, 0.30\}$ with $\tau = 0$ corresponding to maximum likelihood estimation.

The sensitivity of the scale and shape estimators was assessed through the changes in their logarithmic estimates, defined as

$$\Delta_\alpha(z, \tau) = log\hat{\alpha}_{z,\tau} - log\hat{\alpha}_{0,\tau}$$

and

$$\Delta_\beta(z,\tau) = log\hat{\beta}_{z,\tau} - log\hat{\beta}_{0,\tau}$$

respectively, where the subscripts z and 0 denote the contaminated and uncontaminated samples. The reported averages were calculated over replications in which contamination actually altered the observed censored sample, thereby excluding cases in which the contaminated lifetime was withdrawn without affecting the observations.

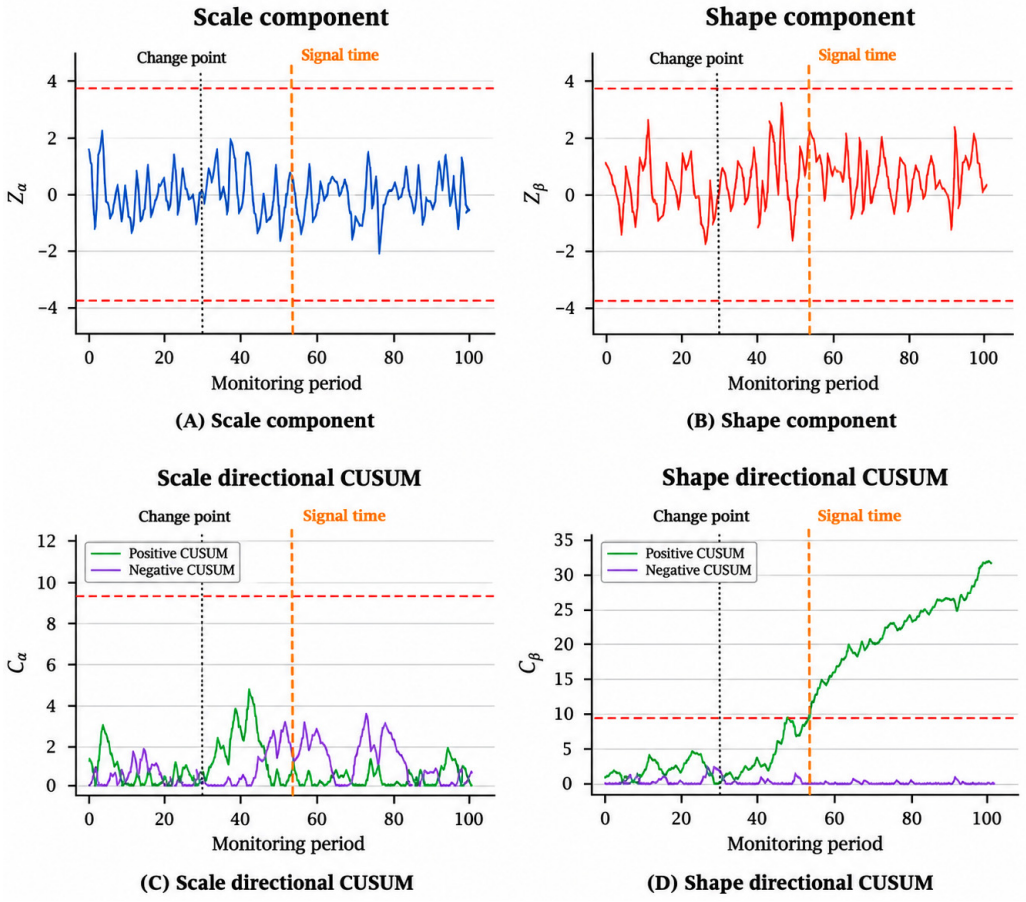

Figure 4. Illustrative monitoring trajectory of the proposed component-wise Shewhart–CUSUM scheme.

In Figure 5, the results reveal that the effect of positive τ is neither uniform across parameters nor independent of the censoring structure. For the shape parameter, $\tau = 0.20$ reduced the largest observed mean change by approximately 6%, 3%, and 8% under early, uniform, and late censoring, respectively, compared with maximum likelihood estimation. These reductions indicate relatively modest improvements in shape-estimation sensitivity rather than a general guarantee of observation-level robustness.

For the scale parameter, the response depended more strongly on the contamination direction. Under early censoring, $\tau = 0.20$ reduced the mean change associated with an extremely short lifetime $z = 0.01$ from approximately -0.143 to -0.087. In contrast, positive τ increased sensitivity to lower-tail contamination under uniform and late censoring. For an extremely long lifetime $z = 20$, the mean scale-estimation change under early censoring increased from approximately 0.91 for maximum likelihood estimation to 1.12 for $\tau = 0.20$.

Under uniform and late censoring, the upper-tail sensitivity curves levelled off as the contamination value increased. This behavior is attributable to the censoring mechanism: sufficiently long lifetimes may be withdrawn or remain unobserved beyond the final recorded failure, so further increases in their magnitude do not necessarily alter the observed censored sample. Thus, the attenuation of upper-tail sensitivity in these settings

is primarily associated with the censoring structure rather than with the density power divergence itself.

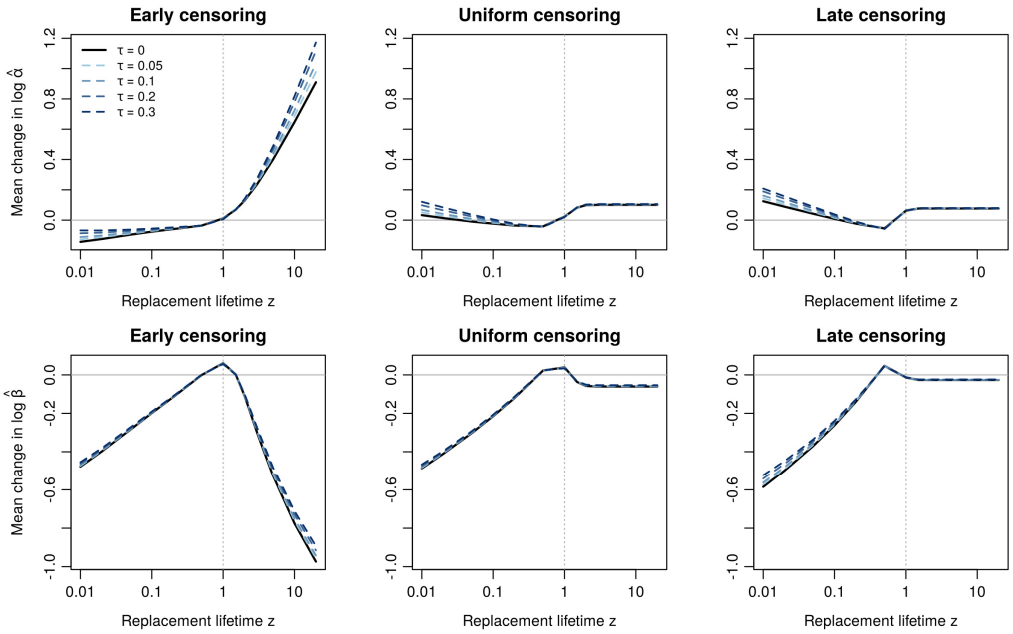

Figure 5. Sensitivity of Weibull scale and shape estimates to single-lifetime contamination under different progressive Type-II censoring schemes.

Overall, the findings demonstrate that joint-density-based MDPDE does not automatically provide robustness against contamination of individual lifetime observations. Its practical sensitivity depends on the tuning parameter, the affected Weibull parameter, the contamination direction, and the progressive censoring configuration. These results support a conditional interpretation of the robust estimation component rather than a claim of uniformly improved resistance to contamination.

4.9. Comparative Performance Analysis with an Existing EWMA-Based Control Chart

To provide an external benchmark for the proposed monitoring framework, we reproduced the joint EWMA-MAX chart of Yu et al. [10], developed for progressively Type-II censored Weibull lifetime data. The comparison was conducted under the same sampling configurations $n = 20$, $m = 10$, three censoring schemes (Early, Uniform, and Late), and 12 parameter-shift scenarios used in the present study. The EWMA-MAX chart was implemented according to Equations (4.5)–(4.7) of Yu et al. [10], with the smoothing parameter fixed at $\lambda = 0.10$. Its control limit was calibrated separately for each censoring scheme to target an in-control average run length of 370. The resulting out-of-control average run lengths and corresponding standard deviations were compared with those of the proposed combined Shewhart–CUSUM chart.

In Table 10, the external comparison reveals substantial differences in detection performance between the two monitoring approaches. In particular, the proposed combined Shewhart–CUSUM chart exhibits considerably shorter average run lengths for increases in the Weibull shape parameter. Under early censoring, for example, the ARL for a 10% shape increase was approximately 50.13 for the proposed chart, compared with 450.12 for the EWMA-MAX chart. The latter value exceeds the nominal in-control ARL, indicating an ARL-biased response, a phenomenon also discussed by Yu et al. [10]. These findings suggest that the directional monitoring structure can offer advantages in detecting certain parameter changes. Nevertheless, the comparison is restricted to the EWMA-MAX

benchmark and does not establish superiority over all EWMA-type or likelihood-based monitoring procedures.

Table 10. Independently verified in-control performance of the Yu et al. EWMA-MAX chart.

Censoring Scheme	λ	Control Limit h	ARL_0
Early	0.10	1.4835	367.81
Uniform	0.10	1.4819	374.35
Late	0.10	1.4777	365.60

Table 11 presents the out-of-control run-length performance of the proposed combined Shewhart–CUSUM chart and the EWMA-MAX chart of Yu et al. [10], with results reported as ARL_1 ($SDRL_1$). Both methods were evaluated under the same sampling configuration ($n = 20$, $m = 10$), three progressive Type-II censoring schemes, and 12 parameter-shift scenarios, with a nominal in-control ARL of 370. The proposed chart uses the selected designs reported in Table 2, whereas the EWMA-MAX chart was implemented using the joint monitoring statistic of Yu et al. [10], with a smoothing parameter of $\lambda = 0.10$ and separately calibrated control limits for each censoring scheme. Bold values in Table 11 indicate the lower ARL_1. The proposed chart achieves shorter ARL_1 values across all 36 evaluated configurations, with particularly pronounced improvements for increases in the Weibull shape parameter. For example, under early censoring, a 10% increase in the shape parameter yields ARL_1 values of 49.58 and 450.12 for the proposed and EWMA-MAX charts, respectively. Similar differences are observed under uniform and late censoring. Furthermore, the EWMA-MAX chart exhibits ARL-biased behavior in certain shape-increase scenarios, consistent with the findings of Yu et al. [10]. The proposed chart also yields lower $SDRL_1$ values across the evaluated configurations, indicating reduced run-length variability. Overall, these results provide external evidence of improved detection performance relative to the EWMA-MAX benchmark under the investigated conditions. Nevertheless, the findings do not establish universal superiority over other EWMA-type or likelihood-based monitoring procedures, including the joint WLRT chart.

Table 11. Comparison of out-of-control run-length performance between the proposed Shewhart–CUSUM chart and the EWMA-MAX chart of Yu et al.

ρ_α	ρ_β	Early Proposed	Early EWMA-MAX	Uniform Proposed	Uniform EWMA-MAX	Late Proposed	Late EWMA-MAX
0.80	1.00	9.80 (2.54)	15.46 (10.11)	9.69 (2.55)	15.19 (10.00)	10.53 (2.99)	14.88 (9.79)
0.90	1.00	23.87 (10.65)	107.80 (102.03)	24.17 (10.75)	115.15 (108.96)	27.05 (12.82)	106.10 (101.31)
1.10	1.00	23.42 (12.19)	59.09 (53.53)	25.23 (12.63)	57.81 (52.75)	28.84 (15.20)	56.68 (50.85)
1.20	1.00	9.93 (4.05)	13.06 (8.92)	11.03 (3.94)	13.30 (9.05)	11.89 (4.46)	12.75 (8.65)
1.00	0.80	12.74 (5.15)	21.77 (16.82)	14.90 (6.27)	29.51 (23.40)	19.03 (8.70)	28.15 (22.47)
1.00	0.90	39.44 (24.18)	89.94 (82.83)	45.01 (29.17)	108.09 (102.48)	51.62 (34.66)	97.67 (91.80)
1.00	1.10	49.58 (31.18)	450.12 (448.07)	64.01 (44.65)	551.16 (539.52)	79.92 (60.72)	546.84 (543.61)
1.00	1.20	19.45 (7.93)	251.99 (246.36)	23.61 (10.35)	376.24 (371.66)	29.07 (14.04)	313.46 (312.61)
0.90	0.90	22.36 (10.04)	65.24 (59.61)	20.21 (8.36)	43.41 (38.02)	22.26 (9.53)	35.07 (28.76)
0.90	1.10	21.65 (8.10)	118.21 (110.64)	24.73 (10.51)	253.18 (247.03)	28.94 (13.50)	330.93 (329.03)
1.10	0.90	18.15 (8.79)	25.35 (20.18)	23.51 (12.54)	40.08 (34.80)	27.70 (16.43)	48.25 (42.94)
1.10	1.10	23.85 (10.65)	87.59 (80.78)	22.86 (9.30)	58.27 (51.72)	26.03 (11.05)	43.26 (37.49)

5. Real Data Application

5.1. Kevlar 49/Epoxy Stress-Rupture Data and Weibull Model Adequacy

To illustrate the practical implementation of the proposed monitoring procedure, stress-rupture lifetime data for Kevlar 49/epoxy strands tested under two constant stress

levels were considered [35,36]. The dataset is a well-known reliability dataset originally associated with the Kevlar stress-rupture experiments and has subsequently been used in numerous lifetime-distribution studies. The lower-stress dataset consists of 49 complete failure-time observations obtained at the 70% stress level, whereas the higher-stress dataset contains 101 complete failure-time observations obtained at the 90% stress level. Failure times are expressed in hours. The 70% stress condition was adopted as the Phase-I reference state, while the 90% stress condition was treated as a changed operating condition for Phase-II monitoring.

The two stress levels exhibit a pronounced difference in lifetime behavior. Under the 70% stress condition, the observed lifetimes extend from approximately 1051 to 17,568 h, whereas the 90% stress data are concentrated at substantially shorter lifetimes, ranging from 0.01 to 7.89 h. Thus, increasing the applied stress from 70% to 90% produces a substantial downward displacement of the lifetime distribution, providing a practically meaningful setting for evaluating whether the proposed monitoring scheme can identify a large deterioration from an established reference state.

Because the proposed monitoring procedure is formulated in terms of a two-parameter Weibull model, the adequacy of the Weibull reference distribution was examined for the 49 observations obtained under the 70% stress condition. Maximum likelihood estimation under the parameterization used throughout this study yielded

$$\hat{\alpha}_0 = 9905.987, \quad \hat{\beta}_0 = 2.01494$$

where α and β denote the Weibull scale and shape parameters, respectively. The fitted model produced a Kolmogorov–Smirnov statistic of

$$D = 0.08769$$

Since the Weibull parameters were estimated from the same observations used for the goodness-of-fit assessment, a parametric-bootstrap procedure was additionally employed rather than relying solely on the conventional plug-in Kolmogorov–Smirnov p-value. Based on 5000 bootstrap replications, the resulting p-value was 0.4358 (Monte Carlo standard error = 0.0070). Hence, there was no evidence against the two-parameter Weibull model at the 5% significance level. For comparison, the conventional plug-in Kolmogorov–Smirnov calculation yielded $p = 0.8454$. The fitted Weibull model was therefore considered adequate for defining the Phase-I reference state.

This conclusion is also consistent with previous analyses of the same 70% stress dataset. For example, published Weibull fitting results for these data report a Kolmogorov–Smirnov statistic of approximately 0.088 and a p-value of approximately 0.843, together with Weibull parameter estimates close to those obtained here.

Accordingly, the 70% stress condition provides a statistically adequate Weibull reference model, whereas the 90% stress observations represent a substantially shortened lifetime regime. This contrast provides an appropriate real-data setting for evaluating both the immediate Shewhart response to a large departure and the subsequent accumulation of directional evidence through the CUSUM components.

As summarized in Table 12, the goodness-of-fit results indicate that the two-parameter Weibull distribution provides an adequate representation of the 70% stress data and can therefore be used to establish the Phase-I reference state. In contrast, the markedly shorter lifetimes observed under the 90% stress condition provide a natural Phase-II setting for evaluating whether the proposed monitoring procedure can detect a substantial deterioration in lifetime behavior. The construction of the Phase-I reference distributions, the retrospective progressive censoring scheme, and the calibration of the monitoring limits are described in the following subsection.

Table 12. Characteristics of the Kevlar 49/epoxy stress-rupture data and Weibull fit for the Phase-I reference condition.

Measure	70% Stress (Phase I)	90% Stress (Phase II)
Number of observations	49	101
Minimum lifetime (h)	1051	0.01
Maximum lifetime (h)	17,568	7.89
Weibull scale, $\hat{\alpha}$	9905.987	—
Weibull shape, $\hat{\beta}$	2.01494	—
KS statistic	0.08769	—
Conventional plug-in KS p-value	0.8454	—
Parametric-bootstrap KS p-value	0.4358	—

5.2. Phase-I Reference Model and Progressive Censoring Design

For implementation of the proposed monitoring procedure, the 49 complete Kevlar 49/epoxy stress-rupture lifetimes observed under the 70% stress condition were used to establish the Phase-I reference state. All 49 observations were retained for estimation of the reference Weibull model, yielding the Phase-I estimates $\hat{\alpha}_0 = 9905.987$ and $\hat{\beta}_0 = 2.01494$. The adequacy of this reference model was supported by the parametric-bootstrap goodness-of-fit analysis reported in Section 5.1.

To maintain consistency with the principal simulation framework, Phase-II monitoring was conducted using subgroups of $n = 20$ units under a progressive Type-II censoring design with $m = 10$ observed failures. The approximately uniform withdrawal scheme

$$\boldsymbol{R} = (1,1,1,\ldots,1,1)$$

was adopted, so that one surviving unit was withdrawn following each observed failure. Since $\sum_{i=1}^{10} R_i = 10$, the resulting design corresponds to an overall censoring proportion of 50%. For each Phase-II subgroup, progressive censoring was imposed retrospectively on the complete lifetime observations by randomly withdrawing one surviving unit after each observed failure. Thus, the censored samples reproduce the information structure that would have been observed under the corresponding progressive life-testing design.

The 101 observations obtained under the 90% stress condition were randomly permuted once using a fixed seed and divided into five non-overlapping subgroups of 20 observations. These five subgroups were subsequently treated as five sequential Phase-II monitoring periods. The single remaining observation was retained as a holdout and was not used in the monitoring calculations. The subgroup allocation and censoring randomization were fixed before evaluation of the monitoring outcomes and were kept unchanged throughout the analysis to ensure reproducibility and to avoid data-dependent selection of Phase-II samples.

For the uniform progressive censoring scheme, the monitoring configuration selected in the nominal simulation study was retained, with $\tau = 0$ and $k = 0.25$. Thus, estimation in the Kevlar application corresponds to maximum likelihood estimation, while the same component-wise signed normal-score transformation and four-directional CUSUM structure developed in Section 3 were preserved. Rather than transferring the control limits obtained under the standardized simulation model, the finite-sample reference distributions and monitoring limits were recalibrated specifically under the fitted Kevlar Phase-I Weibull model. This preserves the application-specific sampling behavior associated with the fitted lifetime distribution while maintaining the monitoring architecture used in the simulation study.

The finite-sample marginal reference distributions of the log-transformed scale and shape estimators were generated using 20,000 Monte Carlo replications under the fitted Phase-I Weibull model and the specified progressive censoring scheme. The Shewhart

and CUSUM limits were then calibrated to attain a target in-control average run length of $ARL_0 = 370$. The selected application-specific configuration was

$$H_S = 3.35 \text{ and } H_C = 10.7314$$

Independent verification based on 20,000 in-control monitoring sequences yielded

$$ARL_0 = 369.19, SDRL_0 = 359.30, MRL_0 = 260.$$

The verified ARL_0 differs from the nominal target by only approximately 0.22%, indicating close agreement with the desired in-control performance. The corresponding Shewhart-only in-control ARL was 666.67, which exceeds the overall target and therefore satisfies the feasibility requirement imposed before allocation of the CUSUM limit.

In the Kevlar application, the reference Weibull parameters were estimated from 49 Phase-I observations. Therefore, the reported application-specific in-control ARL is conditional on the fitted Phase-I reference model and should not be interpreted as an unconditional false-alarm guarantee. In particular, the uncertainty associated with estimating the Weibull parameters from a limited Phase-I sample is not incorporated into the reported calibration. The resulting control limits should consequently be regarded as model-based reference limits, and their unconditional performance remains to be investigated.

As summarized in Table 13, the application-specific calibration closely reproduced the target in-control performance while preserving the monitoring structure used in the simulation study. The five non-overlapping subgroups constructed from the 90% stress observations therefore provide a sequential Phase-II sequence for evaluating the proposed procedure under a substantially more severe stress condition. The resulting component-wise normal scores, directional CUSUM trajectories, and corresponding monitoring decisions are presented in the following subsection.

Table 13. Phase-I reference model and calibrated monitoring parameters for the Kevlar 49/epoxy application.

Design Component	Setting
Phase-I stress level	70%
Phase-I reference observations	49
Phase-II stress level	90%
Phase-II observations	101
Phase-II monitoring subgroups	5
Observations per subgroup, n	20
Observed failures per subgroup, m	10
Phase-II observations used	100
Holdout observations	1
Censoring proportion	50%
Censoring scheme, $\boldsymbol{R}$	$(1, 1, 1, 1, 1, 1, 1, 1, 1, 1)$
Tuning parameter, τ	0
CUSUM reference value, k	0.25
Shewhart limit, H_S	3.35
CUSUM limit, H_C	10.7314
Reference-distribution replications	20,000
Verification replications	20,000
Shewhart-only ARL_0	666.67
Verified combined ARL_0	369.19
$SDRL_0$	359.30
MRL_0	260

5.3. Phase-II Monitoring Results

The calibrated monitoring procedure was applied sequentially to the five non-overlapping 20-unit subgroups constructed from the 90% stress observations. The Phase-I Weibull reference model and the calibrated monitoring limits $H_S = 3.35$ and $H_C = 10.7314$ were held fixed throughout the Phase-II analysis. Progressive Type-II censoring with $m = 10$ observed failures and the uniform withdrawal scheme $\boldsymbol{R} = (1,\ldots,1)$ was applied retrospectively within each subgroup.

A pronounced departure from the Phase-I reference state was observed immediately in the first Phase-II period. The estimated Weibull parameters for the first subgroup were

$$\hat{\alpha}_1 = 0.827, \quad \hat{\beta}_1 = 0.640$$

with corresponding component-wise normal scores

$$Z_{\alpha,1} = -4.056, \quad Z_{\beta,1} = -4.056.$$

Consequently,

$$M = \{max|Z_\alpha|,|Z_\beta|\} = 4.056 > H_S = 3.35.$$

and the Shewhart component generated an immediate signal in the first monitoring period. Thus, the combined procedure detected the substantial deterioration in lifetime behavior as soon as the first 90% stress subgroup was evaluated.

The directional CUSUM statistics subsequently accumulated evidence that the departure persisted across monitoring periods. For the Weibull scale component, the negative-direction CUSUM evolved as

$$C_{\alpha,t}\text{: } 3.806, 7.611, 11.417, 15.223, 19.028, \quad t = 1,\ldots,5.$$

Since $H_C = 10.7314$, the scale CUSUM first crossed its decision limit in the third monitoring period,

$$C_{\alpha,3} = 11.417 > 10.7314.$$

The negative-direction shape CUSUM similarly increased as

$$C_{\beta,t}\text{: } 3.806, 6.312, 7.090, 10.896, 14.016,$$

and crossed the CUSUM limit in the fourth monitoring period,

$$C_{\beta,4} = 10.896 > 10.7314.$$

Thus, the real-data application exhibits the complementary behavior intended by the combined monitoring architecture. The Shewhart component responded immediately to the large abrupt departure observed in the first Phase-II subgroup, whereas the directional CUSUM components accumulated evidence over subsequent periods and confirmed that the change persisted. This behavior is consistent with the transient-shift simulation results, in which the Shewhart pathway was particularly important for large short-duration departures while the contribution of the CUSUM pathway increased as the departure persisted. The scale normal score reached the lower boundary of the finite-sample normal-score transformation in all five Phase-II periods, with

$$Z_{\alpha,t} = -4.056.$$

This boundary value should not be interpreted as indicating identical underlying scale departures across the five subgroups. Rather, it indicates that the estimated scale parameters under 90% stress were consistently more extreme than the lower tail represented by the B_0

= 20.000 Phase-I reference distribution. A similar boundary effect occurred for the shape score in the first and fourth monitoring periods. Accordingly, the score saturation itself provides evidence that these Phase-II estimates lie far outside the finite-sample Phase-I reference distribution.

As shown in Table 14, the combined procedure generated a formal signal immediately in the first Phase-II period through the Shewhart component. The negative scale CUSUM subsequently crossed its decision limit in period 3, followed by the negative shape CUSUM in period 4. Hence, the Kevlar application provides a real-data illustration of the complementary roles of the two monitoring mechanisms: immediate detection of a large abrupt departure by the Shewhart component and sequential confirmation of persistence through directional CUSUM accumulation.

Table 14. Component-wise Phase-II monitoring results for the Kevlar 49/epoxy stress-rupture data.

Phase II Period	$\hat{\alpha}$	$\hat{\beta}$	Z_α	Z_β	M	C_α^-	C_β^+	Shewhart Signal	CUSUM Signal	Combined Signal
1	0.827	0.640	−4.056	−4.056	4.056	3.806	3.806	Yes	No	Yes
2	0.903	1.171	−4.056	−2.756	4.056	7.611	6.312	Yes	No	Yes
3	1.108	1.721	−4.056	−1.028	4.056	11.417	7.090	Yes	Yes	Yes
4	0.638	0.695	−4.056	−4.056	4.056	15.223	10.896	Yes	Yes	Yes
5	1.640	1.059	−4.056	−3.371	4.056	19.028	14.016	Yes	Yes	Yes

5.4. Practical Interpretation of the Real-Data Application

The Kevlar 49/epoxy application provides a practical illustration of how the proposed component-wise monitoring framework responds to a substantial deterioration in lifetime behavior under progressive Type-II censoring. The 70% stress data provided an adequate Weibull Phase-I reference model, while the five non-overlapping subgroups constructed from the 90% stress observations formed a sequential Phase-II monitoring sequence. As shown in Table 14, a pronounced departure from the reference state was identified immediately after monitoring began.

The first Phase-II subgroup produced $Z_{\alpha,1} = -4.056$ and $Z_{\beta,1} = -4.056$, yielding $M = 4.056$, which exceeded the calibrated Shewhart limit $H_S = 3.35$. Consequently, the combined procedure generated a formal signal in the first monitoring period. This immediate response illustrates the role of the Shewhart component in detecting a large abrupt departure before substantial sequential evidence has accumulated. The result is consistent with the large transient-shift experiments in Section 4, where the Shewhart pathway provided a substantial contribution to detection for large short-duration changes.

The subsequent evolution of the directional CUSUM statistics provides complementary information about the persistence and parameter-specific nature of the departure. The negative scale CUSUM increased from 3.806 in the first monitoring period to 7.611 in the second and crossed the decision limit $H_C = 10.7314$ in the third period, reaching $C_{\alpha,3} = 11.417$. It continued to increase to 15.223 and 19.028 in periods 4 and 5, respectively. Similarly, the negative shape CUSUM accumulated more gradually and first exceeded the decision limit in period 4, with $C_{\beta,4} = 10.896$, before increasing to 14.016 in period 5. Thus, while the Shewhart component provided immediate detection, the directional CUSUM components subsequently confirmed that the departure persisted across successive monitoring periods.

The component-wise structure also provides diagnostic information that would be unavailable from a single omnibus monitoring statistic. The consistently negative scale scores indicate a pronounced downward displacement of the Weibull scale component under the higher stress condition, while the negative shape scores indicate an accompanying change in distributional shape. The repeated occurrence of the lower finite-sample normal-

score boundary in several Phase-II periods further indicates that the corresponding estimates lie beyond the lower tail represented by the Phase-I reference distribution. These boundary values should therefore be interpreted as evidence of extreme departures relative to the reference state rather than as identical changes across monitoring periods.

Taken together, the Kevlar application complements the simulation results in Section 4 and provides a real-data demonstration of the intended division of roles within the combined monitoring architecture. For a large abrupt departure, the Shewhart component can provide an immediate signal, whereas continued exposure to the changed process allows the directional CUSUM statistics to accumulate evidence and confirm persistence. At the same time, the separate scale and shape components retain information about the parameter and direction associated with the distributional change. The application therefore demonstrates not only the implementability of the proposed procedure under progressive Type-II censoring, but also its practical ability to provide immediate detection, sequential confirmation, and parameter-specific diagnostic information in a real lifetime-monitoring setting.

6. Conclusions

This study developed a component-wise combined Shewhart–CUSUM monitoring framework for progressively Type-II censored Weibull lifetime processes, incorporating minimum density power divergence estimation (MDPDE) as an optional robust estimation component. The proposed approach accounts for the progressive censoring mechanism through the joint density of the censored sample and transforms the estimated Weibull scale and shape parameters into finite-sample signed normal scores. By combining a two-sided Shewhart mechanism with four directional CUSUM statistics, the framework enables monitoring of both abrupt and persistent distributional changes while retaining parameter-specific diagnostic information.

The simulation results demonstrated that the proposed monitoring procedure can achieve the nominal in-control average run length of approximately 370 under early, uniform, and late progressive censoring. Detection performance depended on the magnitude and direction of parameter changes and the timing of progressive withdrawals. Scale changes were generally detected more rapidly than comparable shape changes. The combined procedure exhibited performance similar to CUSUM-only monitoring for moderate persistent shifts, whereas the Shewhart component provided an important additional signaling pathway for large transient departures. For example, under early censoring with a one-period scale multiplier of 2.00, the detection probability increased from 0.031 for CUSUM-only monitoring to 0.878 for the combined procedure. These findings demonstrate the complementary rather than uniformly superior performance of the combined monitoring structure.

An external comparison with the EWMA-MAX chart showed shorter out-of-control average run lengths for the proposed procedure across the 36 investigated configurations, with particularly pronounced differences for increases in the Weibull shape parameter. Under early censoring, a 10% shape increase yielded average run lengths of 49.58 and 450.12 for the proposed and EWMA-MAX charts, respectively. These results support the effectiveness of the proposed directional monitoring structure under the investigated conditions, although they do not establish universal superiority over other monitoring procedures.

The robustness analysis revealed that the benefits of positive MDPDE tuning are conditional on the censoring scheme, contamination direction, and affected Weibull parameter. At the estimator level, MDPDE reduced the mean squared error of shape estimation by approximately 16.6–35.6% across the investigated configurations, whereas improvements in scale estimation were not uniform. Similarly, chart-level contamination experiments showed that positive tuning could improve false-alarm protection or detection performance in selected settings but could also produce less favorable results than maximum likelihood estimation.

Therefore, robust estimation should be regarded as an optional component rather than an unconditional improvement in monitoring performance. The overall robustness of the procedure depends on the interaction between parameter estimation, progressive censoring, and sequential monitoring.

The practical applicability of the framework was illustrated using Kevlar 49/epoxy stress-rupture data under 70% and 90% stress conditions. The 70% stress observations provided an adequate Weibull Phase-I reference model, supported by a parametric-bootstrap Kolmogorov–Smirnov p-value of 0.4358. The application-specific monitoring procedure achieved a verified in-control average run length of 369.19. When applied to the 90% stress observations, the combined chart generated an immediate Shewhart signal in the first monitoring period, while the negative scale and shape CUSUM statistics subsequently crossed their decision limits in periods 3 and 4, respectively. These findings illustrate the complementary roles of immediate detection and sequential confirmation of persistent changes.

The proposed framework is particularly relevant to reliability monitoring when complete lifetime observations are unavailable. Nevertheless, its present formulation is restricted to two-parameter Weibull processes and prespecified progressive Type-II censoring schemes. Furthermore, the reported control-limit calibration is conditional on the specified or fitted Phase-I reference parameters, and the Kevlar application represents a retrospective rather than prospective monitoring study. Future research should address unconditional false-alarm performance incorporating Phase-I estimation uncertainty, adaptive selection of the divergence tuning parameter, alternative lifetime distributions and censoring mechanisms, and prospective validation using industrial lifetime data. Such extensions may further improve the flexibility and practical applicability of component-wise monitoring under censoring.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/math14214103/s1, Table S1. Sensitivity of the global design criterion to smaller CUSUM reference values under early, uniform, and late censoring. Table S2. Conditional Steady-State Run-Length Performance of the Proposed Combined Shewhart–CUSUM Chart Following a 50-Period In-Control Warm-Up. Table S3. Numerical results underlying Figure S1. Figure S1. Sensitivity of the combined Shewhart–CUSUM chart to the DPD tuning parameter $\tau \in \{0, 0.05, 0.10, 0.20, 0.30\}$ ($k = 0.25$, $H_S = 4.25$).

Author Contributions: Methodology, S.Ç.; Software, S.Ç. and A.B.; Investigation, S.Ç. and A.B.; Writing—Original Draft, S.Ç. and A.B.; Supervision, S.Ç. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by the Scientific Research Projects Coordination Unit of Fırat University under Project No. FF.26.61.

Data Availability Statement: The real lifetime data used in the application are publicly available from the source cited in Section 5.1.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Montgomery, D.C.; Jennings, C.L. *Introduction to Statistical Quality Control*; Wiley: New York, NY, USA, 1991; Volume 5.
2. Page, E.S. Continuous inspection schemes. *Biometrika* **1954**, *41*, 100–115. [CrossRef]
3. Roberts, S.W. Control chart tests based on geometric moving averages. *Technometrics* **2000**, *42*, 97–101. [CrossRef]
4. Lucas, J.M.; Saccucci, M.S. Exponentially weighted moving average control schemes: Properties and enhancements. *Technometrics* **1990**, *32*, 1–12. [CrossRef]
5. Westgard, J.O.; Groth, T.; Aronsson, T.; de Verdier, C.H. Combined Shewhart–CUSUM control chart for improved quality control in clinical chemistry. *Clin. Chem.* **1977**, *23*, 1881–1887. [CrossRef]

6. Haq, A.; Munir, W. New CUSUM and Shewhart–CUSUM charts for monitoring the process mean. *Qual. Reliab. Eng. Int.* **2021**, *37*, 3512–3528. [CrossRef]
7. Hu, X.; Xia, F.; Zhang, J.; Song, Z. Combined Shewhart–EWMA and Shewhart–CUSUM monitoring schemes for time between events. *Qual. Reliab. Eng. Int.* **2024**, *40*, 3352–3380. [CrossRef]
8. Hu, X.; Zhao, Y.; Zhang, J.; Zhong, J. Combined Shewhart–EWMA and Shewhart–CUSUM schemes for monitoring the ratio of two normal variables. *J. Stat. Comput. Simul.* **2025**, *95*, 258–304. [CrossRef]
9. Tyagi, D.; Shonik, H.; Saxena, V. Combined EWMA-mixed EWMA CUSUM schemes for monitoring the process mean. *Int. J. Product. Qual. Manag.* **2025**, *46*, 191–210. [CrossRef]
10. Yu, J.; Wang, Q.; Wu, C. Online monitoring of the Weibull distributed process based on progressive Type II censoring scheme. *J. Comput. Appl. Math.* **2024**, *443*, 115744. [CrossRef]
11. Zhao, W.; Yu, J.; Wu, C. Self-starting monitoring of the progressive Type II censoring data based on goodness-of-fit test. *Qual. Technol. Quant. Manag.* **2025**, *22*, 32–54. [CrossRef]
12. Enami, S.; Alamri, O.A. Phase-II monitoring of the lower truncated proportional hazard rate process based on progressive Type-II censoring. *PLoS ONE* **2025**, *20*, e0322996. [CrossRef] [PubMed]
13. Mukherjee, A.; McCracken, A.K.; Chakraborti, S. Control charts for simultaneous monitoring of parameters of a shifted exponential distribution. *J. Qual. Technol.* **2015**, *47*, 176–192. [CrossRef]
14. Mohammadipour, P.; Farughi, H.; Rasay, H.; Arkat, J. Designing exponentially weighted moving average control charts under failure censoring reliability tests. *Int. J. Eng.* **2021**, *34*, 2398–2407. [CrossRef]
15. Wang, L. Inference for Weibull competing risks data under generalized progressive hybrid censoring. *IEEE Trans. Reliab.* **2018**, *67*, 998–1007. [CrossRef]
16. Liu, M.; Zheng, H.; Yang, J. Reliability evaluation for Weibull distribution with heavily Type II censored data. *Qual. Reliab. Eng. Int.* **2024**, *40*, 3620–3642. [CrossRef]
17. Nassar, M.; Alotaibi, R.; Elshahhat, A. Reliability analysis at usual operating settings for Weibull constant-stress model with improved adaptive Type-II progressively censored samples. *AIMS Math.* **2024**, *9*, 16931–16965. [CrossRef]
18. Alotaibi, R.; Nassar, M.; Elshahhat, A. Rainfall data modeling using improved adaptive Type-II progressively censored Weibull-exponential samples. *Sci. Rep.* **2024**, *14*, 30484. [CrossRef] [PubMed]
19. Kundu, A.; Chowdhury, S.; Modok, B. Control chart for generalized Weibull quantiles under hybrid censoring. *arXiv* **2023**, arXiv:2309.11776.
20. Sanaullah, A.; Chaudhary, A.M.; Hanif, M.; Sharma, P. EWMA control charts based on robust estimators: A powerful tool for monitoring a process with a non-normal distribution. *J. Comput. Appl. Math.* **2024**, *449*, 115919. [CrossRef]
21. Saeed, N.; Bataineh, A.A.M.F.; Abu-Shawiesh, M.O.A.; Haddad, F. Sensitivity and reliability comparisons of EWMA mean control chart based on robust scale estimators under non-normal process: COVID data application. *Qual. Reliab. Eng. Int.* **2024**, *40*, 4513–4529. [CrossRef]
22. Haq, A. An adaptive EWMA mean chart in the presence of outliers. *Qual. Technol. Quant. Manag.* **2024**, *21*, 926–941. [CrossRef]
23. Wu, X.; Hong, L.; Gong, C. A new EWMA control chart for outlier detection in nonlinear profiles of automobile engines. *Qual. Reliab. Eng. Int.* **2025**, *41*, 3669–3685. [CrossRef]
24. AlQadi, H.; Safeer, A.; Alquraish, M.; Aljohani, H.M. Outlier-resistant process monitoring using an RCV-based EWMA control chart. *Qual. Reliab. Eng. Int.* **2026**, *42*, 1149–1158. [CrossRef]
25. Ajadi, J.O.; Raji, I.A.; Abbas, N.; Riaz, M. Detecting outliers in the multivariate control charts for dispersion monitoring. *Qual. Reliab. Eng. Int.* **2024**, *40*, 1904–1917. [CrossRef]
26. Capizzi, G.; Masarotto, G. An adaptive exponentially weighted moving average control chart. *Technometrics* **2003**, *45*, 199–207. [CrossRef]
27. Basu, A.; Harris, I.R.; Hjort, N.L.; Jones, M.C. Robust and efficient estimation by minimising a density power divergence. *Biometrika* **1998**, *85*, 549–559. [CrossRef]
28. Ghosh, A.; Basu, A. Robust estimation for independent non-homogeneous observations using density power divergence with applications to linear regression. *Electron. J. Stat.* **2013**, *7*, 2420–2456. [CrossRef]
29. Fujisawa, H.; Eguchi, S. Robust parameter estimation with a small bias against heavy contamination. *J. Multivar. Anal.* **2008**, *99*, 2053–2081. [CrossRef]
30. Song, J. Sequential change point test in the presence of outliers: The density power divergence based approach. *Electron. J. Stat.* **2021**, *15*, 3504–3550. [CrossRef]
31. Park, C.; Wang, M.; Ouyang, L. Novel robust g and h charts using the generalized Kullback–Leibler divergence. *Comput. Ind. Eng.* **2023**, *176*, 108951. [CrossRef]
32. Mangold, B.; Konopik, J. Introduction of a general class of entropy-based control charts: The Φ-chart. *Manag. Econ.* **2024**, *25*, 55–70. [CrossRef]

33. Atif, M.; Akhtar, N.; Ramzan, Q.; Abbas, S.; Nazir, H.Z.; Alballa, T. Development and validation of an adaptive quantile monitoring framework for Alpha Power Weibull process under censoring. *Qual. Reliab. Eng. Int.* **2026**, *42*, 1580–1593. [CrossRef]
34. Balakrishnan, N.; Aggarwala, R. *Progressive Censoring: Theory, Methods, and Applications*; Birkhäuser: Boston, MA, USA, 2000.
35. Olalude, G.A.; Ojurongbe, T.A.; Bello, O.A.; Bashiru, K.A.; Alamu, K.A. Alpha Power Harris-G family of distributions: Properties and application to Burr XII distribution. *arXiv* **2025**, arXiv:2512.10276.
36. Andrews, D.F.; Herzberg, A.M. *Data: A Collection of Problems from Many Fields for the Student and Research Worker*; Springer: New York, NY, USA, 2012.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.