mathematics

MDPI

Article

A Shape-Adaptive Gamma-Mixture Power Shanker Distribution: Boundary-Aware Inference, Censored Regression, and Predictive Validation

Muqrin A. Almuqrin

Department of Mathematics, College of Science in Zulfi, Majmaah University, Majmaah 11952, Saudi Arabia; m.almuqrin@mu.edu.sa

Abstract

The Gamma-Mixture Power Shanker distribution models positive lifetimes by allowing the gamma component of a Shanker mixture to change shape and then applying a power transformation. Its three parameters control the latent gamma shape, the rate and mixture weight, and the power transformation. The model contains Power Shanker as an interior submodel and Weibull at the boundary. We develop complete-data and right-censored likelihoods, covariate-dependent rate regression, and likelihood-based diagnostics for choosing among these nested models. A boundary argument establishes continuity, identifies the Weibull nuisance parameters, and verifies positive but small efficient information for the additional shape. A new near-boundary simulation reports bias, empirical standard deviation, root mean squared error, profile-set coverage, and boundary-test rejection rates, revealing appreciable finite-sample uncertainty and limited power. In the 228-patient lung-cancer application, the fitted closure selects Weibull; a multi-start parametric-bootstrap reanalysis supports the boundary reference at the five-percent threshold while showing a discrepancy in its zero mass. Across the complete-data applications, relative fit, evidence against Weibull, evidence against Power Shanker, and absolute goodness of fit lead to different conclusions. The rainfall fit fails two calibrated goodness-of-fit criteria, while the glass-fibre Kolmogorov–Smirnov result is borderline. Repeated predictive validation shows no material GMPS advantage for the lung data. The family therefore provides a testable extension whose additional shape should be retained only when the data support it.

Keywords: Shanker distribution; gamma mixture; boundary likelihood ratio; nonregular inference; Weibull limit; right censoring; survival regression; parametric bootstrap; predictive validation

MSC: 62E15; 62F10; 62F03; 62F40; 62N02; 62N05

Check for updates

Academic Editor: Bo Wang

Received: 19 September 2026
Revised: 21 October 2026
Accepted: 26 October 2026
Published: 29 October 2026

Copyright: © 2026 by the author. 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

Positive lifetime and reliability measurements can differ in skewness, tail weight, and the way failure risk changes with time. Mixtures combine simple probability laws to describe this variation. The Shanker distribution [1] combines an exponential component with a gamma component having shape two and positive rate θ. Here, a rate is the parameter governing the exponential decay of the transformed lifetime. Shanker and Shukla [2] applied a power transformation to obtain the two-parameter Power Shanker distribution.

Recent work has continued to enrich this family. Aforka et al. [3] proposed the three-parameter exponentiated Power Shanker (EPS) model and extended it to censored-data

regression. In parallel, recent Sujatha-family studies have introduced three-parameter power extensions [4]. These developments make the novelty question more demanding: a new lifetime model should not be justified merely by adding another free parameter, but by a mathematically distinct construction, interpretable nesting, documented likelihood inference, and evidence that the extra flexibility is used by at least some datasets.

A targeted audit of the closest literature was therefore carried out. Al-Omari and Dobbah [5] mix an entire Shanker law with a gamma law; this is different from changing the shape of the latent gamma component inside the Shanker mixture. Generalized Power Sujatha [6] and Power quasi Sujatha [7] are Sujatha-family constructions, while EPS [3] exponentiates the Power Shanker CDF. The mechanisms are summarized below; the comparison is a positioning audit of the closest located models, not a claim that no mathematically equivalent form could exist elsewhere in the literature.

The term shape-adaptive means that the latent gamma component is allowed to change its shape instead of being fixed at shape two. We write its shape as $\alpha + 1$, where α is the additional nonnegative mixture-shape parameter used in inference. The positive parameter θ determines both the component rate and its constrained mixing weight. A positive power parameter δ transforms the baseline variable and strongly affects endpoint hazard behavior. These roles overlap in the fitted distribution, so the three parameters should not be interpreted as independent physical mechanisms. The original positive-shape model has $\alpha > 0$; its continuous closure includes $\alpha = 0$. Setting $\alpha = 1$ recovers Power Shanker, whereas setting $\alpha = 0$ gives Weibull.

GMPS is appropriate for strictly positive continuous measurements, including failure times and survival durations, when a Weibull or Power Shanker baseline may be too restrictive. It can represent changing density concentration and hazard patterns through its constrained exponential–gamma mixture. It is not intended for observations with an atom at zero, negative values, or arbitrary dependence without further modeling. Its practical advantage is a directly testable extension with exact simulation, rather than an assurance of improved fit for every sample.

Right censoring occurs when follow-up ends before a failure is observed: the recorded time then supplies only a lower bound on the lifetime. This commonly arises when patients are still alive at study closure, or components remain operational when a test ends. Ordinary regression on the recorded times would treat these lower bounds as completed outcomes. Censored regression instead combines density contributions from observed failures with survival probabilities for censored observations, while allowing explanatory variables to change the fitted lifetime distribution. This is the motivation for the rate-regression extension developed in Section 4.

The comparisons concern the four complete-data applications and two censored regressions. We distinguish model construction, usefulness of an additional parameter, absolute model adequacy, and predictive performance; none implies the others. The literature comparison in Table 1 describes the closest located constructions and is not an exhaustive equivalence theorem.

For the specific Power Shanker comparison, the distinction can be seen from the tail. A positive-shape GMPS density has an exponential factor $\exp\left(-\theta x^{\delta}\right)$ and a leading polynomial power $x^{\delta-1+\alpha\delta}$. Equality with a Power Shanker density first forces the same power and exponential rate and then forces $\alpha = 1$ by matching the leading polynomial power. Thus, a member with $\alpha > 0$ and $\alpha \neq 1$ is not obtained merely by relabeling the two Power Shanker parameters. This argument does not establish inequivalence to every other mixture family.

The contributions are the constrained gamma-mixture construction, exact Power Shanker and Weibull nesting, distributional quantities and exact simulation, complete and censored likelihood inference, rate regression, and empirical model diagnostics. Bound-

ary theory and the new near-boundary stress study clarify when the additional shape can be estimated and when simpler models remain preferable. The distinction from earlier models is confined to the construction and comparisons documented in Table 1.

Table 1. Positioning of GMPS against the closest models located in the targeted literature audit.

Model	Year	Core Construction	Parameters	Censoring/Regression in Cited Work	Distinction from GMPS
Power Shanker [2]	2017	Power transform of Shanker	2	No regression reported	Exact submodel at $\alpha = 1$; tail comparison distinguishes positive shapes other than one.
Shanker-Gamma [5]	2023	Mixture of a Shanker distribution and a gamma distribution	2	No regression reported	Mixes whole distributions rather than altering Shanker's latent gamma shape
Three-parameter Power Sujatha [4]	2024	Power-type Sujatha extension	3	Complete-data inference	Different Sujatha baseline and nesting
EPS [3]	2025	Lehmann exponentiation of Power Shanker	3	Right censoring and regression	Changes CDF by exponentiation, not mixture-component shape
Generalized Power Sujatha [6]	2025	Power transform of generalized Sujatha	3	Complete-data inference	Different Sujatha-family mixture
Power quasi Sujatha [7]	2026	Power transform of quasi Sujatha	3	MLE, MPS and bootstrap intervals	Different baseline; no Power Shanker nesting
GMPS (this work)	2026	Shape-adaptive latent gamma component + power transform	3	Complete and right-censored likelihoods	Constrained latent shape; exact Power Shanker and Weibull cases; no exhaustive inequivalence claim.

2. The GMPS Distribution

2.1. Shape-Adaptive Shanker Baseline and Mixture Representation

Let $Y > 0$ denote the baseline variable, let $\alpha > 0$ control the gamma shape, and let $\theta > 0$ be the common component rate. Define the positive normalizing denominator

$$D(\alpha, \theta) = \theta^{\alpha+1} + \Gamma(\alpha + 1). \tag{1}$$

The denominator in (1) collects the integrated masses of the two unnormalized terms. Dividing by it yields the baseline probability density

$$g(y; \alpha, \theta) = \frac{\theta^{\alpha+1}}{D(\alpha, \theta)}(\theta + y^{\alpha})e^{-\theta y}, \qquad y > 0. \tag{2}$$

Integrating (2) over the positive half-line gives one: the constant-in-y term uses $\int_0^{\infty} e^{-\theta y} dy = 1/\theta$, and the power term uses the gamma integral $\int_0^{\infty} y^{\alpha} e^{-\theta y} dy = \Gamma(\alpha + 1)/\theta^{\alpha+1}$. Thus $D(\alpha, \theta)$ is required for normalization, rather than being an additional parameter.

Separating these normalized terms rewrites the same density as a two-component mixture:

$$g(y) = wf_{\text{Exp}(\theta)}(y) + (1-w)f_{\Gamma(\alpha+1,\text{rate}=\theta)}(y), \qquad w = \frac{\theta^{\alpha+1}}{D(\alpha,\theta)}. \tag{3}$$

Here w is the probability of choosing the exponential component, and $1-w$ is the probability of choosing the gamma component; $0 < w < 1$. Equation (3) explains the latent construction, verifies normalization by a second argument, and provides the exact sampler in Section 4.2.

When $\alpha = 1$, $w = \theta^2/(\theta^2+1)$ and the baseline reduces exactly to the Shanker distribution. This nesting is central to the inference developed below.

2.2. Power Transformation

To add a power parameter that changes density and hazard shape while preserving positivity, define the observed variable by

$$X = Y^{1/\delta}, \qquad \delta > 0. \tag{4}$$

The inverse transformation is $y = x^\delta$, with Jacobian $dy/dx = \delta x^{\delta-1}$. Substituting this into (2) gives the GMPS density

$$f(x;\alpha,\theta,\delta) = \frac{\delta\theta^{\alpha+1}}{D(\alpha,\theta)} x^{\delta-1} e^{-\theta x^\delta}\left(\theta + x^{\alpha\delta}\right), \qquad x > 0. \tag{5}$$

Put $z = \theta x^\delta$. The lower and upper incomplete gamma functions are denoted by $\gamma(s,z)$ and $\Gamma(s,z)$. Integrating the two mixture components up to x gives the cumulative probability $F(x) = \Pr(X \le x)$:

$$F(x) = \frac{\theta^{\alpha+1}(1-e^{-z}) + \gamma(\alpha+1,z)}{D(\alpha,\theta)}, \tag{6}$$

The complementary probability $S(x) = \Pr(X > x) = 1 - F(x)$ follows from $\gamma(s,z) + \Gamma(s,z) = \Gamma(s)$:

$$S(x) = \frac{\theta^{\alpha+1}e^{-z} + \Gamma(\alpha+1,z)}{D(\alpha,\theta)}. \tag{7}$$

The CDF in (6) supplies interval probabilities and numerical quantiles. The survival function in (7) gives the probability of remaining failure-free beyond a given time, enters the hazard denominator, and supplies each right-censored likelihood contribution. The density, CDF, and survival function therefore serve different but linked roles in inference.

The hazard rate is $h(x) = f(x)/S(x)$. Conditional on the first mixture component, X is Weibull with CDF $1 - e^{-\theta x^\delta}$; conditional on the second component, X is generalized gamma [8] with shape $\alpha+1$, rate θ, and power δ.

2.3. Special and Limiting Cases

The verified cases are: $\alpha = 1$ gives Power Shanker [2]; $\alpha = 1, \delta = 1$ gives Shanker; and $\delta = 1$ gives the baseline. The continuous extension to $\alpha = 0$ is Weibull, and additionally setting $\delta = 1$ gives exponential. No Sujatha or Lindley-type nesting (e.g., Power Lindley [9]) is asserted. For inference, the parameter space includes the boundary explicitly:

$$\overline{\mathcal{D}} = \{(\alpha,\theta,\delta) : \alpha \ge 0,\ \theta > 0,\ \delta > 0\}. \tag{8}$$

For rate regression, the corresponding space is $\alpha \ge 0$, $\delta > 0$, and a finite coefficient vector, with $\theta_i = \exp(\beta_0 + \mathbf{z}_i^T\boldsymbol{\beta})$. Only the shape coordinate is on a boundary; the nuisance coordinates remain interior. Section 7 verifies the consequences for boundary inference.

2.4. Endpoint Hazard Behavior

As $x \downarrow 0$,

$$h(x) \sim \frac{\delta\theta^{\alpha+2}}{D(\alpha,\theta)}x^{\delta-1}. \tag{9}$$

Consequently, the hazard diverges at the origin for $0 < \delta < 1$, approaches a positive finite value for $\delta = 1$, and starts from zero for $\delta > 1$. As $x \to \infty$,

$$h(x) \sim \delta\theta x^{\delta-1}. \tag{10}$$

Thus, the terminal hazard tends to zero, θ, or infinity according as $\delta < 1$, $\delta = 1$, or $\delta > 1$. These asymptotic statements are analytic; finer global shapes should be assessed numerically for specific parameter values.

As shown in Figure 1, the GMPS density can move from a sharply decreasing reverse-J form to clearly unimodal shapes as the parameter combination changes. The location and height of the mode, together with the concentration of right-tail mass, vary appreciably, confirming that the three parameters control more than a simple scale change.

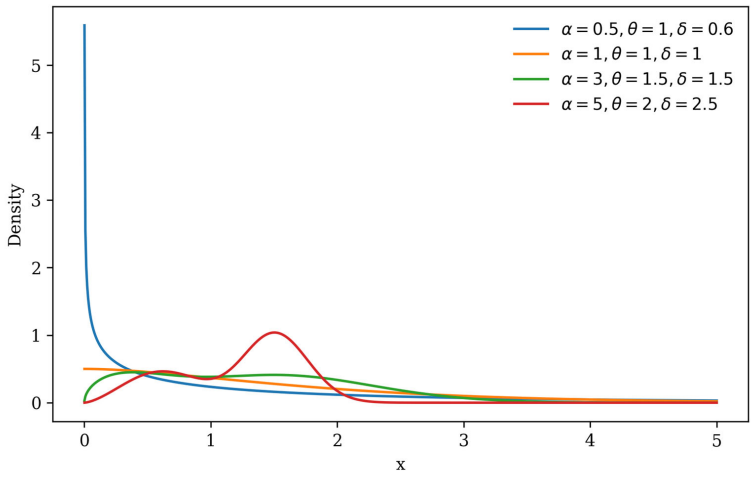


Figure 1. Representative GMPS density shapes for verified parameter combinations.

Figure 2 illustrates the corresponding hazard flexibility. Depending on the parameters, the failure rate may remain low over an initial interval, increase gradually, or rise very rapidly at larger times; these numerical curves are consistent with the analytic endpoint regimes derived in Equations (9) and (10).

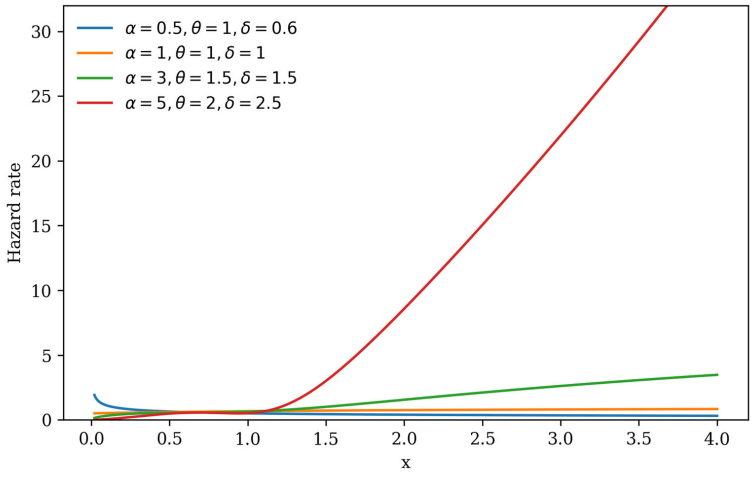


Figure 2. Representative GMPS hazard-rate curves. The curves are numerical evaluations of (5) and (7), and their endpoint tendencies agree with (9) and (10).

3. Distributional Properties

3.1. Raw Moments

Raw moments summarize magnitude and shape through expectations of powers of the lifetime. Conditional on the mixture component, $X^r = Y^{r/\delta}$. Applying exponential and gamma moments and then averaging with the weights in (3) gives, for $r > -\delta$,

$$\mu'_r = E(X^r) = \frac{\theta^{\alpha+1}\Gamma\left(1+\frac{r}{\delta}\right)+\Gamma\left(\alpha+1+\frac{r}{\delta}\right)}{\theta^{r/\delta}D(\alpha,\theta)}. \tag{11}$$

Hence $E(X) = \mu'_1$, $\mathrm{Var}(X) = \mu'_2 - (\mu'_1)^2$, and standardized skewness and kurtosis follow from the first four raw moments.

Table 2 evaluates (11) at $r = 1, 2, 3, 4$. The variance is the second central moment; the coefficient of variation is standard deviation divided by mean; skewness is the standardized third central moment; and kurtosis is the standardized fourth central moment, without subtracting three. The table is a numerical illustration of shape variation, not an estimation experiment. For example, changing the power from 0.7 to 1.5 at the same first two parameters greatly reduces relative dispersion and right skewness. The final row illustrates that left skewness is also possible within the examined parameter choices.

Table 2. Verified moment summaries obtained directly from (11).

α	θ	δ	Mean	Variance	CV	Skewness	Kurtosis
0.5	0.5	0.7	5.0412	40.8677	1.2681	2.9710	17.6799
0.5	1.5	0.7	0.8641	1.4210	1.3795	3.2416	20.4947
1.0	1.0	0.7	2.1700	7.5048	1.2624	2.8185	15.9715
3.0	1.5	0.7	2.6888	9.1552	1.1253	1.9147	8.4225
0.5	0.5	1.5	1.7945	1.1405	0.5951	0.8910	3.9224
1.0	1.0	1.5	1.2037	0.5356	0.6080	0.8217	3.6817
3.0	1.5	1.5	1.3301	0.6607	0.6111	0.4436	2.6108
5.0	2.0	2.5	1.2258	0.2335	0.3942	−0.4376	2.2238

3.2. Incomplete Moments, Lorenz and Bonferroni Curves

The lower incomplete moment $M_k(t) = E\left[X^k \mathbf{1}(X \le t)\right]$ measures the contribution below a threshold. Integrating each component only up to t, with the same substitution used for (11), gives for $k > -\delta$,

$$M_k(t) = \frac{\theta^{-k/\delta}}{D(\alpha,\theta)}\left[\theta^{\alpha+1}\gamma\left(1+\frac{k}{\delta},\theta t^\delta\right)+\gamma\left(\alpha+1+\frac{k}{\delta},\theta t^\delta\right)\right]. \tag{12}$$

For a numerically obtained quantile $Q(p)$, the Lorenz and Bonferroni curves are $L(p) = M_1(Q(p))/\mu'_1$ and $B(p) = L(p)/p$, respectively.

The Lorenz curve is the share of total expected magnitude contributed by the lowest proportion p of observations; the Bonferroni curve rescales that share by p. They are useful for concentration summaries of positive measurements, while the incomplete moment itself also supports truncated expectations.

3.3. Mean Residual Life

The mean residual life is the expected additional time conditional on having survived to x. Subtracting the elapsed time from $E[X\mathbf{1}(X > x)]/S(x)$, and evaluating the upper incomplete first moment with the mixture, gives

$$m(x) = \theta^{-1/\delta}\frac{\theta^{\alpha+1}\Gamma\left(1+\frac{1}{\delta},z\right)+\Gamma\left(\alpha+1+\frac{1}{\delta},z\right)}{\theta^{\alpha+1}e^{-z}+\Gamma(\alpha+1,z)} - x, \qquad z = \theta x^{\delta}. \tag{13}$$

Unlike the instantaneous hazard, mean residual life has the same units as the lifetime and answers a remaining-duration question. Its numerator weights the two components by their surviving tail mass, so the relevant mixture composition changes as the conditioning time increases.

3.4. Quantiles

A p-quantile is the time by which a proportion p has failed. Since the density is positive on the positive half-line, (6) is strictly increasing from zero to one. Setting $F(q) = p$ therefore defines a unique positive solution:

$$\theta^{\alpha+1}\left(1-e^{-\theta q^{\delta}}\right)+\gamma\left(\alpha+1,\theta q^{\delta}\right) = pD(\alpha,\theta), \qquad 0 < p < 1. \tag{14}$$

A bracketing solver expands a positive upper endpoint until $F(q) > p$ and then bisects or uses Brent iteration. This gives medians and other fitted lifetime summaries without assuming an elementary inverse CDF.

3.5. Moment-Generating and Characteristic Functions

The moment-generating function is $M_X(t) = E\left[e^{tX}\right]$, when finite; it can encode moments through derivatives at zero. The characteristic function replaces t by an imaginary argument and always exists. Because the right tail has exponential factor $e^{-\theta x^{\delta}}$ times a polynomial factor, comparison with e^{tx} yields

$$M_X(t) < \infty \quad \begin{cases} \text{for all real } t, & \delta > 1, \\ \text{for } t < \theta, & \delta = 1, \\ \text{only for } t \le 0, & 0 < \delta < 1. \end{cases} \tag{15}$$

For $\delta > 1$ the decay dominates every linear exponential; for $\delta = 1$ finiteness requires $t < \theta$; and for $0 < \delta < 1$ every positive t produces divergence. Thus, a moment series must not be used as an MGF outside its convergence domain. Characteristic-function evaluation by integration remains possible even when the MGF is unavailable.

3.6. Rényi Entropy

Rényi entropy measures concentration of a continuous density, with a larger order emphasizing high-density regions. For order $\rho > 0$, $\rho \neq 1$, and a finite integral, define

$$H_{\rho}(X) = \frac{1}{1-\rho}\log I_{\rho}, \qquad I_{\rho} = \int_0^{\infty} f(x)^{\rho}\,dx. \tag{16}$$

For a general real order, numerical quadrature is used. Near zero, integrability requires $\rho(\delta-1) > -1$; the exponential tail handles infinity. For a positive integer m, expanding $\left(\theta + x^{\alpha\delta}\right)^m$ by the finite binomial theorem and substituting $u = m\theta x^{\delta}$ gives

$$I_m = \left(\frac{\delta\theta^{\alpha+1}}{D}\right)^m \frac{1}{\delta}\sum_{j=0}^{m}\binom{m}{j}\theta^{m-j}(m\theta)^{-\nu_j/\delta}\Gamma\left(\frac{\nu_j}{\delta}\right), \qquad \nu_j = m(\delta-1)+\alpha\delta j+1. \tag{17}$$

The restrictions $\nu_j > 0$ ensure that every resulting gamma integral exists. Equation (17) is therefore an evaluable finite sum for integer orders satisfying these restrictions and does not justify a finite expansion for noninteger orders. Entropy is a distributional summary, not evidence of empirical superiority.

3.7. Order Statistics

Order statistics describe ordered failures, including the first and last failure in a test. For the kth smallest observation from n independent lifetimes, choose the observations falling below and above x, multiply their probabilities by the density at x, and obtain

$$f_{X_{(k)}}(x) = \frac{n!}{(k-1)!(n-k)!} f(x) F(x)^{k-1} S(x)^{n-k}, \qquad k = 1, \ldots, n. \tag{18}$$

The factors in (18) correspond to $k-1$ observations below x and $n-k$ above it; the combinatorial coefficient counts their arrangements. This is useful for component-order reliability questions and for deriving inference under order-based sampling schemes.

4. Maximum-Likelihood Inference

4.1. Log-Likelihood and Score Equations

Maximum likelihood chooses the parameter values that assign the greatest joint probability density to the observed data. For independent complete observations $x_1, \ldots, x_n > 0$, multiplying (5) and taking logarithms yields

$$\begin{aligned} \ell(\alpha, \theta, \delta) = n\log\delta + n(\alpha+1)\log\theta - n\log D + (\delta-1)\textstyle\sum_{i=1}^{n} \log x_i + \\ \textstyle\sum_{i=1}^{n} \log\left(\theta + x_i^{\alpha\delta}\right) - \theta \textstyle\sum_{i=1}^{n} x_i^{\delta}. \end{aligned} \tag{19}$$

The score is the vector of first derivatives of (19). With ψ denoting the digamma function, the derivatives with respect to the rate, mixture shape, and power are

$$U_\theta = \frac{n(\alpha+1)}{\theta} - \frac{n(\alpha+1)\theta^{\alpha}}{D} + \sum_{i=1}^{n} \frac{1}{\theta + x_i^{\alpha\delta}} - \sum_{i=1}^{n} x_i^{\delta}. \tag{20}$$

$$U_\alpha = n\log\theta - n\frac{\theta^{\alpha+1}\log\theta + \Gamma(\alpha+1)\psi(\alpha+1)}{D} + \delta \sum_{i=1}^{n} \frac{x_i^{\alpha\delta}\log x_i}{\theta + x_i^{\alpha\delta}}. \tag{21}$$

$$U_\delta = \frac{n}{\delta} + \sum_{i=1}^{n} \log x_i + \alpha \sum_{i=1}^{n} \frac{x_i^{\alpha\delta}\log x_i}{\theta + x_i^{\alpha\delta}} - \theta \sum_{i=1}^{n} x_i^{\delta}\log x_i. \tag{22}$$

At an interior maximum, the three scores vanish. One first evaluates (19) and its scores, optimizes over positive parameters, and compares the attained likelihood across several starting values. The original interior parameterization uses logarithms to enforce strict positivity. A boundary fit additionally requires the explicit Weibull branch $\alpha = 0$; this value cannot be reached by taking a finite logarithm of α. The operational distinction and the fully specified new stress-study implementation are described below.

Curvature of the log-likelihood measures local information. At a well-identified interior maximum, a covariance approximation is

$$\widehat{\mathrm{Var}}(\hat{\boldsymbol{\eta}}) = \left[-\frac{\partial^2 \ell}{\partial \boldsymbol{\eta}\, \partial \boldsymbol{\eta}^T} \Big|_{\hat{\boldsymbol{\eta}}} \right]^{-1}, \qquad \boldsymbol{\eta} = (\alpha, \theta, \delta)^T. \tag{23}$$

The inverse in (23) is meaningful only when the observed information is positive definite and sufficiently well-conditioned. A convergence flag alone does not establish identification. Flat profiles, highly correlated parameter directions, sensitivity to starting points, and very wide intervals require caution. Profile likelihood is therefore preferred for α, and symmetric Wald intervals are not used at its boundary.

For the new boundary stress study in Section 5.4, we optimize $(\alpha, \log\theta, \log\delta)$ using L-BFGS-B with analytic scores, $\alpha \in [0, 60]$, log-rate bounds $[-30, 30]$, and log-power bounds $[-5, 5]$. Starts for α are 0, 0.15, 1, and 5, with nuisance starts derived from the log-lifetime mean and variance. Function and gradient tolerances are 10^{-12} and 10^{-7}; the iteration limit is 1200. The released implementation minimizes the mean negative log-likelihood and its

mean analytic score for numerical scaling; summed log-likelihoods are used for likelihood comparisons and profiling. The exact Weibull nuisance fit is always a candidate, and the candidate with the largest finite likelihood is retained. The artificial upper limit is reported as a computational safeguard, not as part of the model definition; no selected fit reached it.

Each reported likelihood comparison should use the larger of the optimized interior and exact boundary likelihoods, with nuisance parameters refitted under every restriction. For right censoring, the same steps use (26), and regression uses (28) with an unconstrained coefficient vector. When starts disagree materially, additional starts or profiling are required. The original complete-data Monte Carlo procedure used a single start at the generating parameter, log-parameter bounds $[-4, 3]$, $[-6, 4]$, and $[-4, 2.2]$, 400 iterations, and function tolerance 10^{-10}. Those settings differ from the new multi-start stress study and can make the original experiment optimistic about numerical initialization. For the lung boundary reanalysis, zero, 0.5, and 2 are used as alternative-shape starts from the fitted null nuisance vector; the best finite likelihood is compared with the exact null.

4.2. Exact Random Generation

The mixture representation gives an exact simulation algorithm: generate $B \sim \text{Bernoulli}(w)$; if $B = 1$, generate $Y \sim \text{Exp}(\theta)$, otherwise generate $Y \sim \Gamma(\alpha + 1, \text{rate} = \theta)$; then set $X = Y^{1/\delta}$. No envelope constant or rejection step is required.

4.3. Testing Whether the Extra Shape Parameter Is Needed

Since Power Shanker is nested at $\alpha = 1$, test $H_0 : \alpha = 1$ against $H_1 : \alpha \neq 1$ by

$$\Lambda = 2\left\{\ell\left(\hat{\alpha}, \hat{\theta}, \hat{\delta}\right) - \ell\left(1, \tilde{\theta}, \tilde{\delta}\right)\right\}, \tag{24}$$

The ordinary one-degree-of-freedom chi-square reference applies to this interior null under regular identification. This test asks whether the data require departing from Power Shanker. It is distinct from testing the Weibull boundary $\alpha = 0$ in Section 7, and a failure to reject $\alpha = 1$ is not proof that the full three-parameter model is necessary.

4.4. Profile Likelihood for $\boldsymbol{\alpha}$

Because the complete-data simulation indicates weak finite-sample identification of α, Wald intervals based only on a local Hessian can be misleading. We therefore use the profile-deviance function

$$R_p(\alpha) = 2\left\{\ell\left(\hat{\alpha}, \hat{\theta}, \hat{\delta}\right) - \sup_{\theta,\delta>0} \ell(\alpha, \theta, \delta)\right\}. \tag{25}$$

For a regular positive value, a two-sided 95% profile set uses the one-degree-of-freedom 95th percentile. Such a set is a diagnostic of likelihood shape and is preferable to a local Wald interval when the profile is flat. Near zero, its nominal coverage need not be accurate. Boundary tests use a different reference; a two-sided set reaching zero can coexist with rejection by a one-sided boundary test because their cutoffs differ.

4.5. Independent Right Censoring

Let $T_i = \min(X_i, C_i)$ be the recorded time and $d_i = \mathbf{1}(X_i \leq C_i)$ indicate an observed failure. We assume that censoring is independent of lifetime, conditional on the modeled covariates in the regression, with distinct censoring parameters. An observed failure contributes $f(T_i)$, while a censored observation contributes the probability $S(T_i)$ of surviving past its recorded time. Omitting the parameter-distinct censoring factor gives

$$\ell_c(\alpha, \theta, \delta) = \sum_{i=1}^{n} [d_i \log f(T_i; \alpha, \theta, \delta) + (1 - d_i) \log S(T_i; \alpha, \theta, \delta)]. \tag{26}$$

Maximizing (26) combines both types of information without imputing the unobserved failure times. Logarithms preserve positivity for nuisance parameters and for an interior shape; the exact zero-shape branch is evaluated separately. Closed-form survival evaluation avoids numerical integration, although incomplete gamma evaluation and numerical optimization are still required.

4.6. Covariate-Dependent Rate Regression

For subject i, let $\mathbf{z}_i$ denote a vector of centered/scaled covariates and replace the common rate by

$$\theta_i = \exp(\beta_0 + \mathbf{z}_i^T \boldsymbol{\beta}). \tag{27}$$

The global shape parameters α and δ remain common across subjects. With independent right censoring, the regression log-likelihood is

$$\ell_R(\alpha, \delta, \beta_0, \boldsymbol{\beta}) = \sum_{i=1}^{n} [d_i \log f(T_i; \alpha, \theta_i, \delta) + (1 - d_i) \log S(T_i; \alpha, \theta_i, \delta)]. \tag{28}$$

This parameterization is deliberately direct: covariates act on the positive GMPS rate rather than being introduced through an unrelated transformation. Because increasing θ_i generally compresses the survival-time scale while also changing the mixture weight, regression effects are interpreted through fitted survival/quantile contrasts rather than as constant hazard ratios. Nested hypotheses such as $H_0 : \beta_j = 0$ can be assessed by likelihood-ratio or profile-likelihood inference.

4.7. Predictive Validation Under Censoring

For each subject, the fitted regression yields a predicted survival curve. The revised validation uses five repetitions of five stratified folds. Within each training fold, continuous predictors are median-imputed and standardized, and the learned transformations are applied to its held-out fold. All likelihood fits now use initialization calculated from that training fold, with multiple starts; no whole-sample fitted coefficients are used for initialization. The GMPS fit compares nonnegative-shape candidates with the exact Weibull branch. The repeat seeds are 2026091900 through 2026091904, and the evaluation grid is 60 through 540 days in 30-day increments.

For the reported validation, a single scalar risk score is obtained from each fitted curve as $r_i = 1 - \hat{S}_i(365)$. Harrell concordance compares the ordering of these scores among pairs for which an earlier observed failure establishes the survival-time ordering [10]. It is not a cumulative/dynamic area under the curve at 365 days: the horizon defines the risk score, while comparable pairs use the observed follow-up. For unequal recorded times, the comparable-pair definition is given in (30) below. Identical risk scores receive half credit. The audited implementation excludes pairs with equal recorded times, including event–censor ties, and gives half credit only for tied predicted risks among unequal-time comparable pairs. Calibration and overall prediction error use IPCW Brier scores [11]. Let $\hat{G}$ be the censoring survival estimate learned from the training fold. Then

$$\widehat{\mathrm{BS}}(t) = \frac{1}{n} \sum_{i=1}^{n} \left[\frac{I(T_i \le t, d_i = 1) \hat{S}_i(t)^2}{\hat{G}(T_i^-)} + \frac{I(T_i > t) \{1 - \hat{S}_i(t)\}^2}{\hat{G}(t)} \right]. \tag{29}$$

The integrated Brier score (IBS) is the trapezoidal integral of $\mathrm{BS}(t)$ over 60–540 days divided by the interval length. Parametric models are additionally compared by the held-out censored negative log-likelihood per subject. Calibration at 365 days is assessed by grouping cross-validated predicted survival into quintiles and comparing each group's mean prediction with its Kaplan–Meier estimate. The full validation uses five repetitions of stratified five-fold cross-validation (25 train/test splits), consistent with the principle that

prognostic discrimination and calibration should be assessed on data not used to fit the model [10].

The 180- and 365-day Brier scores provide six-month and one-year summaries, while the 60–540-day IBS describes the stated medium-term evaluation window. These horizons are descriptive choices and do not define an independently validated clinical decision threshold. IPCW evaluation requires positive censoring-survival estimates over the relevant window; predictions beyond supported follow-up should be excluded. The between-repeat standard deviations describe resampling variability and are not independent-sample standard errors. No formal paired test or multiplicity adjustment was applied to the small model differences in the predictive comparison of Section 6.7.2.

For unequal follow-up times, the concordance definition used to interpret the reported scalar risks is

$$C = \frac{\sum_{i,j:\, T_i<T_j,\, d_i=1}\left[\mathbf{1}\left(r_i > r_j\right) + 1/2\mathbf{1}\left(r_i = r_j\right)\right]}{\sum_{i,j}\mathbf{1}\left(T_i < T_j,\, d_i = 1\right)}, \qquad r_i = 1 - \hat{S}_i(365). \tag{30}$$

5. Monte Carlo Study

The numerical implementation was checked using five parameter settings: PDF normalization, differentiation of the CDF, agreement of moments with quadrature, and finite-difference checking of the complete-data score. The reported maximum discrepancies were 2.95×10^{-12}, 3.43×10^{-10}, 8.42×10^{-13}, and 6.14×10^{-8}, respectively. For the new experiment in Section 5.4, the analytic score was independently checked against finite differences, and the efficient-information identity in Section 7.1 was verified.

5.1. Complete-Data Design

The original complete-data experiment used the exact sampler with three positive-shape settings: $(0.5, 1.0, 0.7)$, $(1.5, 1.2, 1.0)$, and $(3.0, 1.5, 1.6)$, with sample sizes 50, 100, 250, and 500 and 500 replications per cell. Table 3 retains the reported RMSE summaries. The original estimation procedure initializes each sample at its generating parameter. The multi-start experiment in Section 5.4 uses different scenarios and is not a matched sensitivity analysis of Table 3. It adds bias, empirical standard deviation, and profile-set coverage for its own design, without establishing the effect of initialization on the original RMSE summaries.

Table 3 shows that the absolute RMSE of α falls with sample size but remains substantial. Direct comparison of raw RMSEs across parameters also reflects their different scales, so the conclusion concerns practical uncertainty in α, rather than a scale-invariant ranking.

The decrease in RMSE in Table 3 is descriptive of this generating-parameter initialization scheme. Since these same scenarios have not been refitted from data-derived or dispersed starts, the magnitude and direction of any change in RMSE under alternative initialization remain unquantified. Favorable starts may reduce numerical error relative to fitting data with unknown parameters. Table 3 therefore does not establish the robustness of a practical estimation routine to starting values, and its reported convergence rates do not demonstrate that different starts reach the same optimum.

Table 3. Monte Carlo RMSEs from 500 replications per cell, conditional on initialization at the true generating parameters. All reported optimizer convergence rates were 1.000; these flags do not establish attainment of a global maximum or insensitivity to initialization.

Scenario	n	RMSE $(\hat{\boldsymbol{\alpha}})$	RMSE $(\hat{\boldsymbol{\theta}})$	RMSE $(\hat{\boldsymbol{\delta}})$
S_1	50	1.2416	0.3811	0.1087
	100	0.9355	0.2923	0.0720
	250	0.6588	0.2149	0.0452
	500	0.5270	0.1767	0.0323
S_2	50	1.2625	0.3965	0.1310
	100	1.0052	0.3159	0.0927
	250	0.7222	0.2250	0.0571
	500	0.5574	0.1746	0.0424
S_3	50	1.2816	0.4118	0.1539
	100	1.0162	0.3168	0.0946
	250	0.5327	0.1653	0.0583
	500	0.4050	0.1274	0.0417

5.2. Right-Censoring Stress Test

To verify (26), data were generated from GMPS$(1.5, 1.2, 1.0)$ and independently censored by exponential censoring times calibrated to approximately 20% and 40% censoring. For each censoring level and $n = 100, 250, 200$ datasets were generated and refitted by censored MLE. All 800 fits converged in this experiment.

Table 4 shows lower RMSE at the larger sample size and larger RMSE with heavier censoring in the examined cells. The value 1.3682 for α at 40% censoring and sample size 100 is substantial relative to the generating shape 1.5. Convergence of all reported fits does not remove this uncertainty.

Table 4. Right-censored MLE stress test (200 replications per cell).

Target Censoring	n	Actual Censoring	RMSE $(\hat{\boldsymbol{\alpha}})$	RMSE $(\hat{\boldsymbol{\theta}})$	RMSE $(\hat{\boldsymbol{\delta}})$
20%	100	20.88%	1.1233	0.3394	0.1062
20%	250	20.26%	0.7951	0.2506	0.0653
40%	100	39.28%	1.3682	0.4293	0.1386
40%	250	39.80%	0.9208	0.2754	0.0865

5.3. Regression Simulation Under Right Censoring

The regression likelihood (28) was evaluated under $\alpha = 1.5$, $\delta = 1.2$, $\beta_0 = -0.4$, and a standardized-covariate effect $\beta_1 = 0.45$. A standard-normal covariate was generated for each subject, and the censoring rate was calibrated to approximately 20% or 40%. For $n = 100$ and 250, 150 independent datasets were generated per cell. Every one of the 600 regression fits used in the summary converged and produced an invertible observed information matrix.

The covariate coefficient is estimable in this design, but nominal Wald coverage is below 95% in several cells. The regression-simulation summaries are reported in Table 5.

Profile-based uncertainty and sensitivity to weak shape identification remain important in applications.

Table 5. GMPS rate-regression simulation under independent right censoring (150 replications per cell). The intercept is omitted here for compactness.

Target Censoring	n	Parameter	Bias	RMSE	Wald 95% Coverage
20%	100	α	0.0129	0.7590	0.953
20%	100	β_1	0.0467	0.1663	0.900
20%	100	δ	0.0606	0.1912	0.907
20%	250	α	0.0288	0.4473	0.940
20%	250	β_1	0.0058	0.0843	0.933
20%	250	δ	0.0149	0.0732	0.973
40%	100	α	0.0744	0.8596	0.920
40%	100	β_1	0.0469	0.1630	0.907
40%	100	δ	0.0578	0.1652	0.920
40%	250	α	−0.0447	0.5663	0.940
40%	250	β_1	0.0202	0.1025	0.900
40%	250	δ	0.0385	0.0960	0.940

5.4. Near Boundary Complete Data Stress Study

A new independent experiment addresses the unexamined region near zero shape. The exact sampler generates 1000 samples per cell at $\alpha = 0, 0.1, 0.5$, $\theta = 1.2$, $\delta = 1$, and $n = 100, 500$. Each fit explicitly compares the Weibull branch with four constrained alternative starts using the settings in Section 4.1. Seed 20261014 and the scenario-specific seed sequences specify the random-number initialization used in this experiment. Table 6 summarizes shape estimation and testing. These runs add a separate near-boundary experiment and do not reconstruct the original complete-data scenarios. For each scenario, the NumPy SeedSequence input is [20261014, round(10000 a), n, r], where a is the generating mixture shape, n is the sample size, and r is the replication index from 0 to 999.

Table 6. New near-boundary experiment with 1000 replications per cell. Rejection uses LR greater than 2.7055. At zero shape, it estimates type I error; at positive shapes, it estimates power. The Monte Carlo standard error of a rejection fraction near 0.05 is approximately 0.007.

True Shape	n	Bias Shape	SD Shape	RMSE Shape	LR Rejection	Zero Fit	Converged
0.0	100	1.501	1.286	1.976	0.064	0.002	1.000
0.0	500	0.887	0.831	1.215	0.063	0.028	1.000
0.1	100	1.384	1.308	1.904	0.067	0.006	1.000
0.1	500	0.731	0.830	1.106	0.068	0.026	1.000
0.5	100	0.772	1.183	1.412	0.076	0.007	1.000
0.5	500	0.219	0.655	0.690	0.073	0.055	1.000

At true zero shape, rejection fractions are 0.064 and 0.063 for sample sizes 100 and 500, respectively. The selected zero-shape fractions are only 0.002 and 0.028, far from the asymptotic atom of one half in these finite-sample searches. An alternative interior

likelihood mode can dominate even when a local boundary mode exists, so the finite-sample global maximum need not exhibit the local asymptotic zero mass. At positive shapes 0.1 and 0.5, rejection fractions remain only 0.067–0.076 in this design. The shape RMSE at true zero is 1.976 and 1.215, confirming substantial estimation uncertainty. These findings do not refute the conditional local theorem, but they do rule out an unqualified finite-sample claim.

Profile-set coverage is evaluated by likelihood-ratio inversion at the generating parameter: cutoff 3.8415 for interior coordinates and 2.7055 for zero shape. Coverage is empirical, not asserted from the cutoff. Shape coverage is 0.936–0.937 at true zero and 0.953–0.972 at positive shapes. In the verified public-release rerun, all 6000 selected likelihood values were finite, and all selected optimizer flags indicated success after mean-objective scaling. Table 6 reports these rerun convergence fractions; the statistical summaries are unchanged at the displayed precision. No selected shape reaches the artificial bound 60. The experiment supports cautious profiling and design-specific bootstrap calibration; it does not establish a universally calibrated boundary test.

6. Real-Data Applications

6.1. Datasets and Competitors

Four positive datasets were analyzed. The 63 strengths of 1.5-cm glass fibres were originally examined by Smith and Naylor [12]. The carbon-fibre breaking-stress dataset contains 100 observations and is associated with Nichols and Padgett [13]. The bladder-remission dataset contains 128 uncensored remission times from Lee and Wang [14]. The Los Angeles series contains 85 annual rainfall totals for the Civic Center over 1922–2006 and has been used in lifetime-distribution comparisons [15].

Competitors include the nested Power Shanker model, the recent three-parameter EPS model [3], generalized gamma, Power Sujatha, Weibull, and gamma. All fits below were recomputed by maximum likelihood under a common numerical workflow. KS and Anderson-Darling values are reported as descriptive statistics; no unadjusted plug-in p-values are asserted because parameters were estimated from the same data. The recomputed fits are reported in Table 7 for the glass-fibre and Los Angeles rainfall data and in Table 8 for the carbon-fibre and bladder-remission data.

Table 7. Recomputed fits for the glass-fibre and Los Angeles rainfall data.

Dataset	Model	LogLik	AIC	BIC	KS	AD
Glass fibres	GMPS	−12.454	30.909	37.338	0.1125	0.5531
	EPS	−14.924	35.848	42.278	0.1477	1.0944
	Generalized gamma	−14.588	35.175	41.605	0.1453	1.0675
	Power Shanker	−15.797	35.595	39.881	0.1626	1.4213
	Power Sujatha	−14.351	32.702	36.989	0.1415	1.0505
	Weibull	−15.207	34.414	38.700	0.1522	1.2408
	Gamma	−23.952	51.903	56.189	0.2164	3.0873
LA rainfall	GMPS	−277.100	560.201	567.529	0.1054	0.6681
	EPS	−276.241	558.482	565.810	0.0744	0.4283
	Generalized gamma	−276.390	558.781	566.109	0.0898	0.4995
	Power Shanker	−279.683	563.365	568.251	0.1260	1.0369
	Power Sujatha	−279.112	562.224	567.109	0.1236	0.9640

Table 7. *Cont.*

Dataset	Model	LogLik	AIC	BIC	KS	AD
	Weibull	−282.274	568.548	573.433	0.1362	1.4244
	Gamma	−277.993	559.987	564.872	0.1169	0.8313

Table 8. Recomputed fits for the carbon-fibre and bladder-remission data.

Dataset	Model	LogLik	AIC	BIC	KS	AD
Carbon fibres	GMPS	−141.338	288.675	296.491	0.0650	0.4087
	EPS	−141.355	288.709	296.525	0.0666	0.4120
	Generalized gamma	−141.346	288.692	296.507	0.0642	0.4090
	Power Shanker	−141.356	286.711	291.922	0.0674	0.4149
	Power Sujatha	−141.568	287.135	292.346	0.0610	0.3958
	Weibull	−141.529	287.059	292.269	0.0605	0.4177
	Gamma	−143.234	290.467	295.678	0.0934	0.7589
Bladder remission	GMPS	−412.255	830.511	839.067	0.0639	0.6001
	EPS	−410.413	826.826	835.382	0.0426	0.2376
	Generalized gamma	−410.854	827.708	836.264	0.0468	0.2991
	Power Shanker	−412.256	828.512	834.216	0.0640	0.5999
	Power Sujatha	−413.293	830.585	836.289	0.0685	0.7548
	Weibull	−414.087	832.174	837.878	0.0700	0.9577
	Gamma	−413.368	830.736	836.440	0.0733	0.7716

For glass fibres, GMPS has the smallest AIC, KS, and AD, although Power Sujatha has a slightly smaller BIC. For Los Angeles rainfall, EPS has the smallest AIC, KS, and AD, while gamma has the smallest BIC. For carbon fibres, Power Shanker has the smallest AIC and BIC. For bladder remission, EPS provides the strongest AIC, KS, and AD, whereas Power Shanker has the smallest BIC. These results do not support a blanket claim that GMPS is universally best.

6.2. Does the Data Use the New Parameter?

Table 9 reports likelihood-ratio tests of the nested Power Shanker submodel (unit shape) for the four complete-data applications.

Table 9. Likelihood-ratio tests comparing GMPS with nested Power Shanker.

Dataset	Λ	p-Value	Interpretation at 5%
Glass fibres	6.6860	0.0097	Unit shape rejected
LA rainfall	5.1647	0.0231	Unit shape rejected
Carbon fibres	0.0361	0.8492	Unit shape not rejected
Bladder remission	0.0013	0.9715	Unit shape not rejected

This nested comparison is necessary but not sufficient: it asks whether α improves Power Shanker, not whether the fitted GMPS is itself adequate. We therefore add profile-likelihood and refitted bootstrap diagnostics.

6.3. Profile-Likelihood Identification of $\boldsymbol{\alpha}$

Table 10 reports 95% profile-likelihood intervals for the additional gamma-shape parameter in the four complete-data applications.

Table 10. Profile-likelihood intervals for the additional gamma-shape parameter. Here 0^+ means that the 95% profile set reaches arbitrarily close to the positive boundary; a finite lower crossing was not obtained.

Dataset	$\hat{\boldsymbol{\alpha}}$	95% Profile Interval	LR p-Value for $\boldsymbol{\alpha} = 1$
Glass fibres	4.8916	[2.0808, 8.1830]	0.0097
LA rainfall	14.1853	[1.6096, 28.1844]	0.0231
Carbon fibres	0.5542	$(0^+, 3.3763]$	0.8492
Bladder remission	0.9759	$(0^+, 2.4431]$	0.9715

Glass-fibre and rainfall profile sets exclude one, consistent with the Power Shanker tests. Carbon-fibre and bladder-remission profile sets reach the lower boundary under the ordinary two-sided cutoff. This indicates uncertainty about the magnitude of the shape. For bladder remission, the boundary LR of 3.6629 lies between the one-sided boundary cutoff 2.7055 and the ordinary two-sided cutoff 3.8415, so these two findings are compatible.

Figure 3 makes the identification differences across the four complete-data applications explicit. The glass-fibre and Los Angeles rainfall profiles have clear interior minima and cross the 95% cutoff on both sides, whereas the carbon-fibre and bladder-remission profiles extend toward the lower boundary, indicating much weaker information about α in those data.

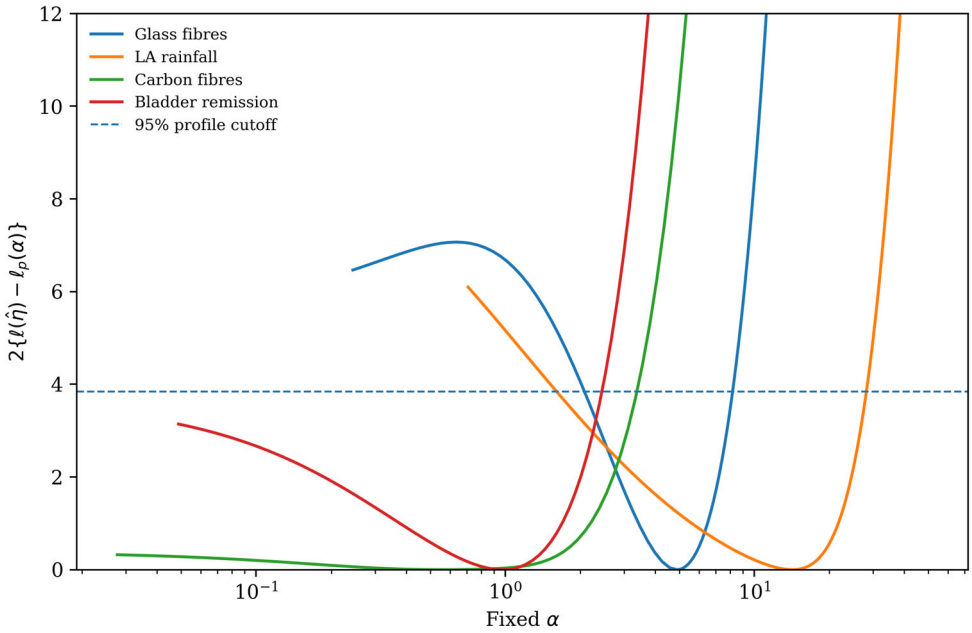


Figure 3. Profile likelihood deviance for alpha across the four real datasets.

6.4. Local Uncertainty for All Three GMPS Parameters

Table 11 retains local observed-information intervals computed on the log-parameter scale and back-transformed. In carbon fibres, the shape interval from 0.00094 to 328.074 is too wide to support a precise scientific interpretation of the point estimate. This is practical weak identification, even though the model is locally identifiable in principle. Profile likelihood, the nested alternatives, and uncertainty in derived predictions should take priority over that point estimate.

Table 11. Local observed-information 95% intervals. The extremely wide carbon-fibre intervals agree with the boundary-reaching profile and indicate weak practical identification.

Dataset	$\hat{\alpha}$ (95% Wald)	$\hat{\theta}$ (95% Wald)	$\hat{\delta}$ (95% Wald)
Glass fibres	4.892 [2.656, 9.008]	1.491 [0.755, 2.946]	2.888 [2.454, 3.399]
LA rainfall	14.185 [4.288, 46.930]	3.490 [0.536, 22.709]	0.554 [0.331, 0.928]
Carbon fibres	0.554 [0.00094, 328.074]	0.151 [0.00090, 25.415]	2.210 [0.711, 6.871]
Bladder remission	0.976 [0.252, 3.776]	0.345 [0.115, 1.034]	0.787 [0.608, 1.018]

6.5. Parametric-Bootstrap Goodness of Fit After Refitting

Because plug-in EDF p-values ignore parameter estimation, each GMPS fit was recalibrated at submission-scale replication by simulating from the fitted model, refitting all three parameters, recomputing the EDF statistic, and comparing with the observed statistic. The target was 1000 refitted samples per dataset. There were 999 successful fits for glass fibres and 1000 for each of the other datasets. The resulting refitted parametric-bootstrap p-values are reported in Table 12.

Table 12. Refitted parametric-bootstrap p-values for GMPS using approximately 1000 successful refits per dataset.

Dataset	KS p	CvM p	AD p	Successful Refits
Glass fibres	0.0470	0.0690	0.0870	999
LA rainfall	0.0170	0.0400	0.0599	1000
Carbon fibres	0.3167	0.2258	0.2657	1000
Bladder remission	0.1908	0.0809	0.0799	1000

The higher-replication calibration materially sharpens the earlier interpretation. Glass fibres remain strong in relative model comparison and support $\alpha \neq 1$, but the bootstrap KS value is 0.047, just below 5%, so absolute adequacy is borderline rather than unqualified. Los Angeles rainfall is rejected by KS and CvM at 5%, with AD borderline ($p = 0.0599$). Carbon fibres and bladder remission are not rejected by any of the three calibrated EDF criteria. The manuscript therefore separates parameter usefulness, relative information-criterion performance, and absolute calibration.

As shown in Figure 4, the refitted bootstrap distributions place the observed KS statistics in very different parts of their null distributions. The glass-fibre result is borderline, the Los Angeles rainfall statistic lies in the extreme tail, and the carbon-fibre and bladder-remission statistics are well inside their bootstrap distributions; this supports the dataset-specific adequacy conclusions reported in the text.

Figure 5 shows that the fitted density curves reproduce the main unimodal structure of the glass-fibre measurements around their central mass. The competing models are visually close over the bulk of the data, so the advantage of any one model should be judged jointly with the likelihood and goodness-of-fit statistics rather than from the histogram alone.

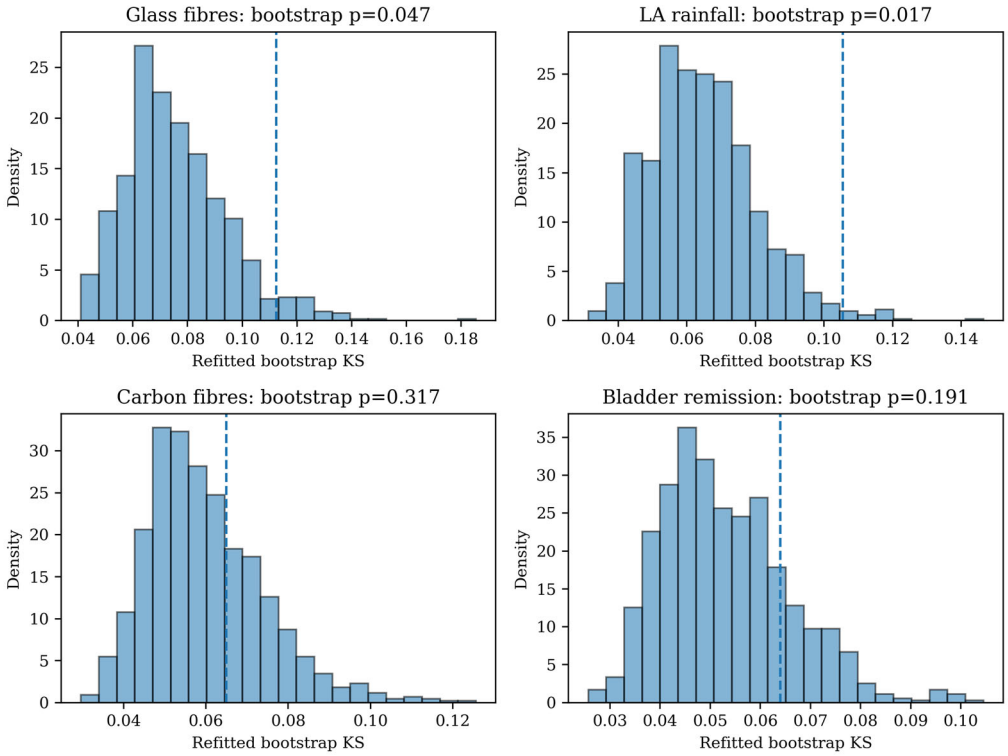


Figure 4. Refitted bootstrap distributions of the KS statistic; vertical dashed lines show observed KS values.

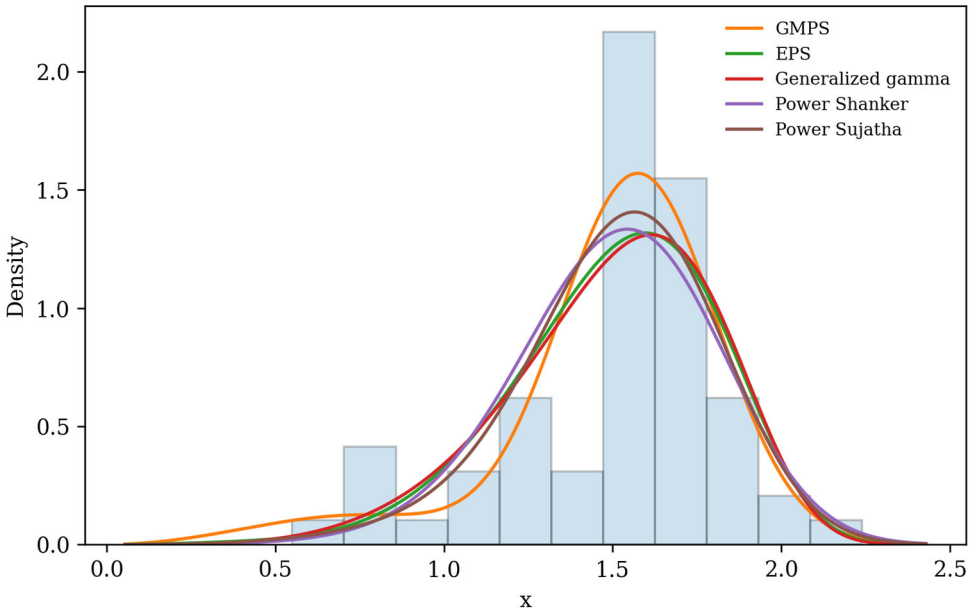


Figure 5. Glass-fibre histogram with fitted densities.

As shown in Figure 6, the fitted CDFs track the empirical glass-fibre distribution closely over most of the support. The largest visible differences occur around the middle transition region, while the fitted curves converge in both tails.

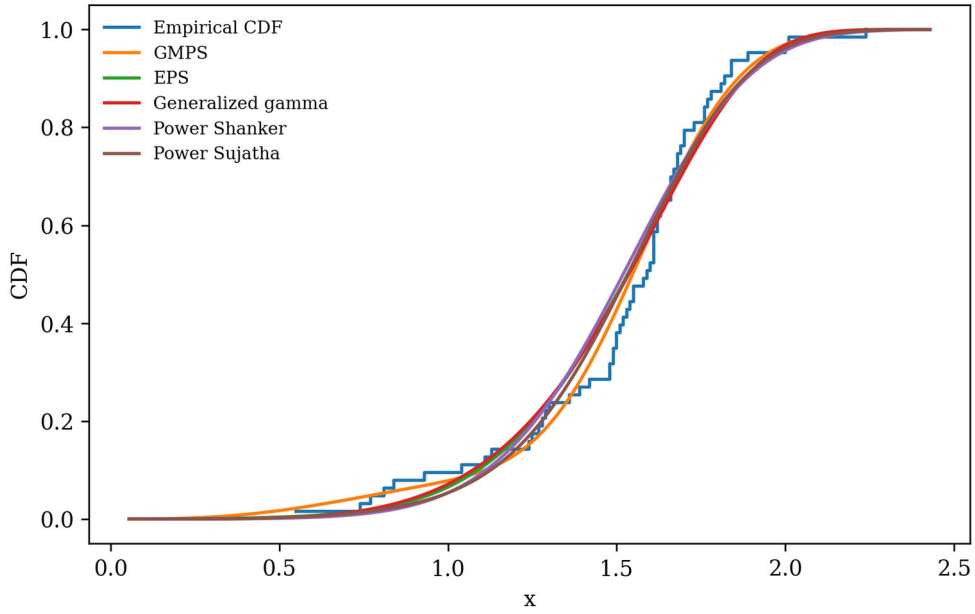


Figure 6. Empirical and fitted CDFs for the glass-fibre data.

Figure 7 displays the strong right-skewness of the bladder-remission times. The fitted densities agree closely over the principal probability mass but differ most near the origin and in the long right tail, which are precisely the regions where survival models with different hazard structures can separate.

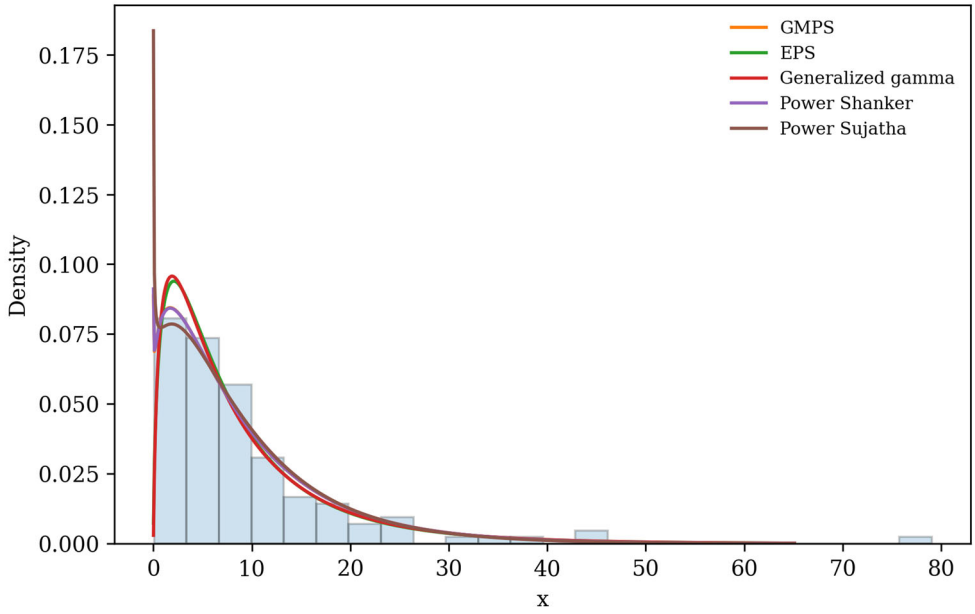


Figure 7. Bladder-remission histogram with fitted densities.

Figure 8 shows that all fitted CDFs follow the empirical bladder-remission CDF closely through the central range. The small separation among curves at intermediate and large times explains why model-selection differences are modest relative to the very similar overall cumulative fit.

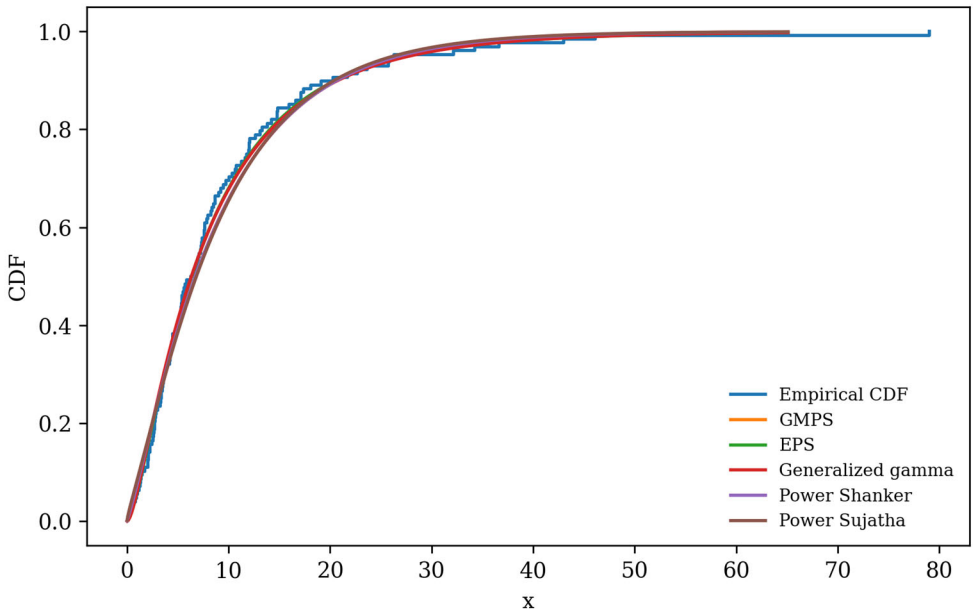


Figure 8. Empirical and fitted CDFs for the bladder-remission data.

6.6. Real Right-Censored Regression: Stanford Heart-Transplant Data

A real censored regression analysis was conducted using the Stanford heart-transplant survival data distributed with statsmodels [16] and traced to Miller's censored-regression study [17]. The data contain 69 patients, with 45 observed deaths and 24 right-censored times (34.8% censoring), together with age at surgery. Age was standardized before fitting the rate regression in (27) and (28).

The fitted GMPS regression gives $\hat{\alpha} = 0.523$, $\hat{\delta} = 0.448$, and $\hat{\beta}_{age} = 0.378$ per one standard deviation of age (8.81 years). The likelihood-ratio statistic for no age effect is 7.005 ($p = 0.0081$). Because the local Wald interval is unstable in this small sample, profile likelihood is emphasized: the 95% profile interval for β_{age} is $[0.075, 0.908]$, excluding zero. The model-implied median survival decreases from approximately 566 days at age 40 to 213 days at age 50 and 79 days at age 60. These are model-based contrasts, not causal effects. Table 13 compares the GMPS rate regression with competing parametric regression models.

Table 13. Parametric regression comparison for the right-censored heart-transplant data. Smaller AICc/BIC are preferred.

Model	k	LogLik	AICc	BIC
Lognormal AFT	3	−311.628	629.625	635.958
Log-logistic AFT	3	−311.750	629.869	636.202
Power Shanker rate regression	3	−312.020	630.410	636.743
Weibull PH regression	3	−312.108	630.585	636.918
GMPS rate regression	4	−311.923	632.470	640.782

The heart data do not require an additional shape beyond Power Shanker: the reported LR is 0.196 with $p = 0.658$. Lognormal and log-logistic regressions also have smaller AICc. The boundary test does not reject Weibull. Thus the age association does not establish a need for the full GMPS baseline, and model-based survival contrasts should be interpreted alongside simpler alternatives.

As shown in Figure 9, the younger-age group has consistently higher Kaplan-Meier survival than the older-age group, and the GMPS regression predictions preserve this ordering. The fitted smooth survival curves follow the main empirical separation, providing a direct visual counterpart to the significant age effect found by likelihood inference.

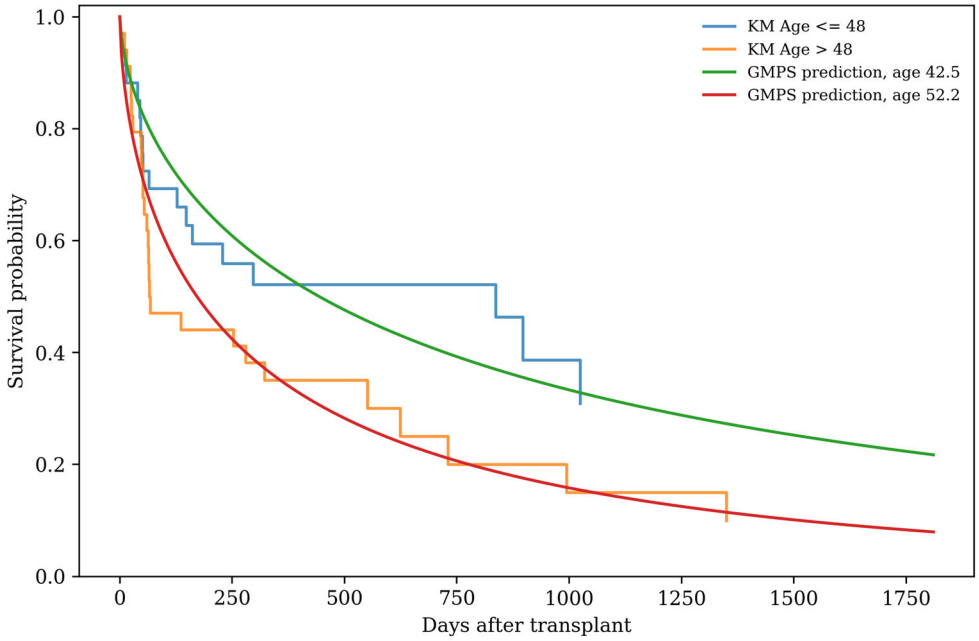


Figure 9. Kaplan-Meier curves by age half together with GMPS regression predictions at the first and third age quartiles.

Figure 10 shows an asymmetric profile-deviance curve with its minimum near the estimated standardized age coefficient. The 95% cutoff is crossed away from zero, so the profile supports a nonzero age effect while also revealing the asymmetry that a symmetric Wald interval would miss.

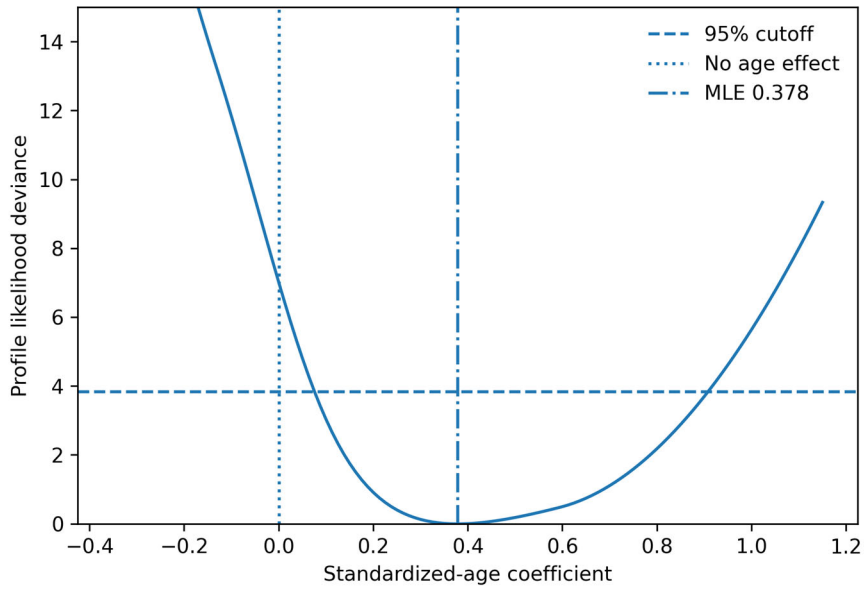


Figure 10. Profile-likelihood deviance for the standardized age coefficient in the heart-transplant GMPS regression.

Figure 11 compares the Cox-Snell residual cumulative hazard with the unit-slope reference. Agreement is close over most of the residual range, with greater departure only

in the sparsely populated upper tail, suggesting no major systematic lack of fit in the central portion of the censored heart-transplant data.

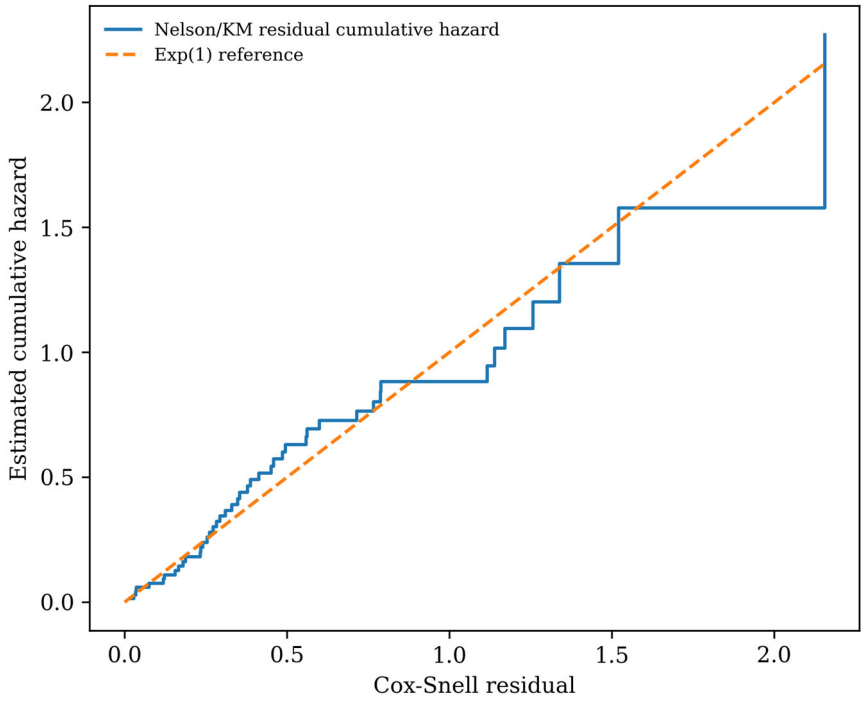


Figure 11. Cox-Snell residual cumulative-hazard diagnostic for the heart-transplant GMPS regression.

6.7. Large Multivariable Censored Validation: NCCTG Lung Cancer Data

To test the regression framework under a larger and more realistic covariate structure, the GMPS model was applied to the North Central Cancer Treatment Group advanced lung-cancer data distributed with the survival package and originating from the prognostic study of Loprinzi et al. [18]. The dataset contains 228 patients, 165 observed deaths, and 63 right-censored observations (27.6% censoring). Four clinically interpretable predictors were pre-declared for this stress test: age, sex, physician-rated ECOG performance score, and physician-rated Karnofsky performance score. Only one ECOG value and one Karnofsky value are missing. Full-sample descriptive fitting uses median imputation, while cross-validation performs median imputation and standardization inside each training fold before applying those transformations to its test fold.

6.7.1. Full-Data Regression and Boundary Behavior

After selection of the Weibull boundary, the reduced Weibull regression gives a covariate-block LR statistic of 32.286 on four degrees of freedom ($p = 1.67 \times 10^{-6}$). The fitted log-rate coefficients are 0.106 for standardized age, −0.585 for female sex, 0.461 for standardized ECOG, and 0.158 for standardized Karnofsky. Within this reduced model, local two-sided Wald p-values are 0.209, 0.00051, 0.00018, and 0.163, respectively. These are conditional summaries under the selected Weibull baseline and do not incorporate baseline-selection uncertainty. Female sex and ECOG have the clearest adjusted associations. In the unrestricted GMPS family, coefficients should be interpreted using fitted survival contrasts; at the Weibull boundary the rate link has a proportional-hazards interpretation. These associations are not causal effects.

The continuous model closure at zero is exactly Weibull. The reported constrained optimum is at $\alpha = 0$, so the closure and Weibull must have identical maximized likelihoods. Table 14 now uses the reported Weibull value for both closure fits and recalculates the

nominal information criteria from that common value; the earlier, slightly smaller GMPS interior likelihood is not a constrained maximum. Because a boundary parameter changes effective model complexity, these nominal AICc/BIC values are descriptive rather than a substitute for boundary inference. Section 7.2 presents the boundary profile.

Table 14. Full-data likelihood comparison for the 228-patient lung-cancer application. Smaller AICc/BIC are preferred.

Model	k	LogLik	AICc	BIC
Weibull	6	−1137.708	2287.797	2307.993
GMPS	7	−1137.708	2289.926	2313.422
Power Shanker	6	−1139.160	2290.700	2310.896
Log-logistic	6	−1143.502	2299.384	2319.580
Lognormal	6	−1152.804	2317.988	2338.184

The fitted multivariable model still produces substantial survival separation between favorable and adverse covariate profiles, but this separation is driven by the covariates and the simpler boundary baseline rather than by evidence for an interior α.

Figure 12 illustrates the clinical impact of the fitted covariates: the favorable profile maintains substantially higher predicted survival, whereas the adverse profile declines much more rapidly. The separation grows over the follow-up period, showing that the covariate block is informative even though the baseline shape collapses to the Weibull boundary.

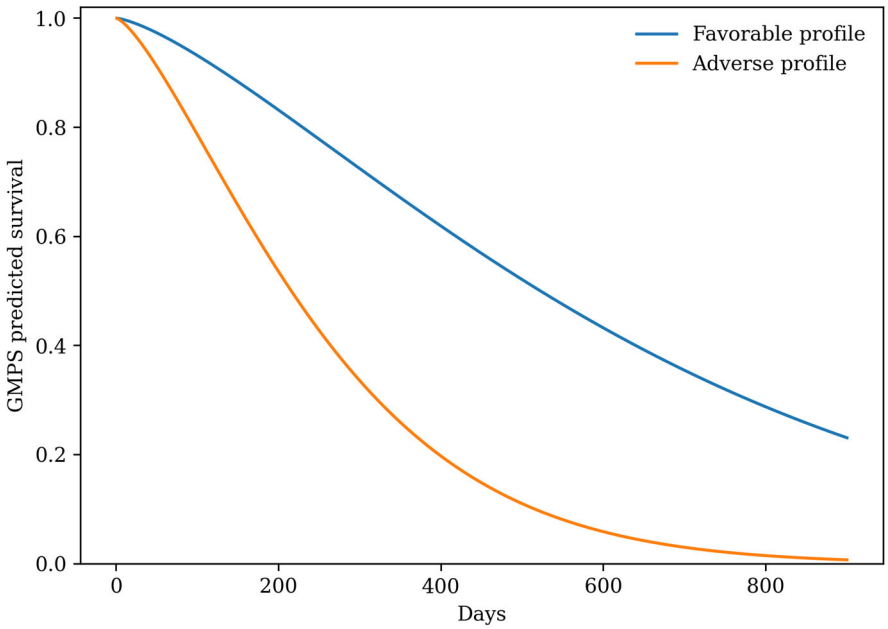


Figure 12. Illustrative GMPS predicted survival curves for favorable and adverse multivariable lung-cancer profiles.

6.7.2. Repeated Cross-Validated Predictive Comparison

Predictive validation was rerun with training-only initialization and an explicit zero-shape branch in every GMPS fit. All 150 model fits across 25 training splits converged. Table 15 reports means across the five repeated cross-validation runs; parentheses give between-repeat standard deviations. The smallest training-fold censoring-survival estimate on the evaluation grid is 0.537, so the stated evaluation window is supported in these folds. The numerical weight floor is 10^{-6} and is inactive on that grid. The numerical reanalyses

used Python 3.12.14, NumPy 2.3.5, pandas 2.2.3, SciPy 1.17.0, statsmodels 0.14.6, and scikit-learn 1.8.0 (the verified public-release environment). The repeat seeds, fold design, fitting rules, and metric definitions are stated in Sections 4.1 and 4.7.

Table 15. Repeated 5-by-5-fold predictive validation. Larger C-index and smaller Brier/IBS/negative-log-score values are preferred. Differences are small and should not be overinterpreted.

Model	Harrell C	Brier 180 d	Brier 365 d	IBS 60–540 d	Held-Out Neg. Log Score
GMPS	0.6187 (0.0016)	0.1866 (0.0015)	0.2362 (0.0014)	0.1963 (0.0009)	5.0280 (0.0086)
Power Shanker	0.6206 (0.0017)	0.1859 (0.0020)	0.2368 (0.0014)	0.1963 (0.0010)	5.0299 (0.0053)
Weibull	0.6187 (0.0016)	0.1865 (0.0017)	0.2361 (0.0014)	0.1963 (0.0010)	5.0244 (0.0035)
Lognormal	0.6200 (0.0045)	0.1896 (0.0023)	0.2388 (0.0013)	0.1983 (0.0009)	5.1025 (0.0234)
Log-logistic	0.6254 (0.0036)	0.1839 (0.0023)	0.2381 (0.0011)	0.1959 (0.0009)	5.0487 (0.0095)
Cox PH [19]	0.6192 (0.0015)	0.1867 (0.0013)	0.2375 (0.0009)	0.1970 (0.0007)	–

The numerical predictive differences in Table 15 are small. Log-logistic has the largest mean concordance and the smallest 180-day Brier score and IBS, while Weibull has the smallest 365-day Brier score and parametric held-out negative log score. These are descriptive comparisons from the repeated splits, without a formal paired superiority or equivalence test. GMPS has no demonstrated predictive advantage. Table 15 reports concordance; Figures 13 and 14 present the Brier curves and calibration plot, respectively.

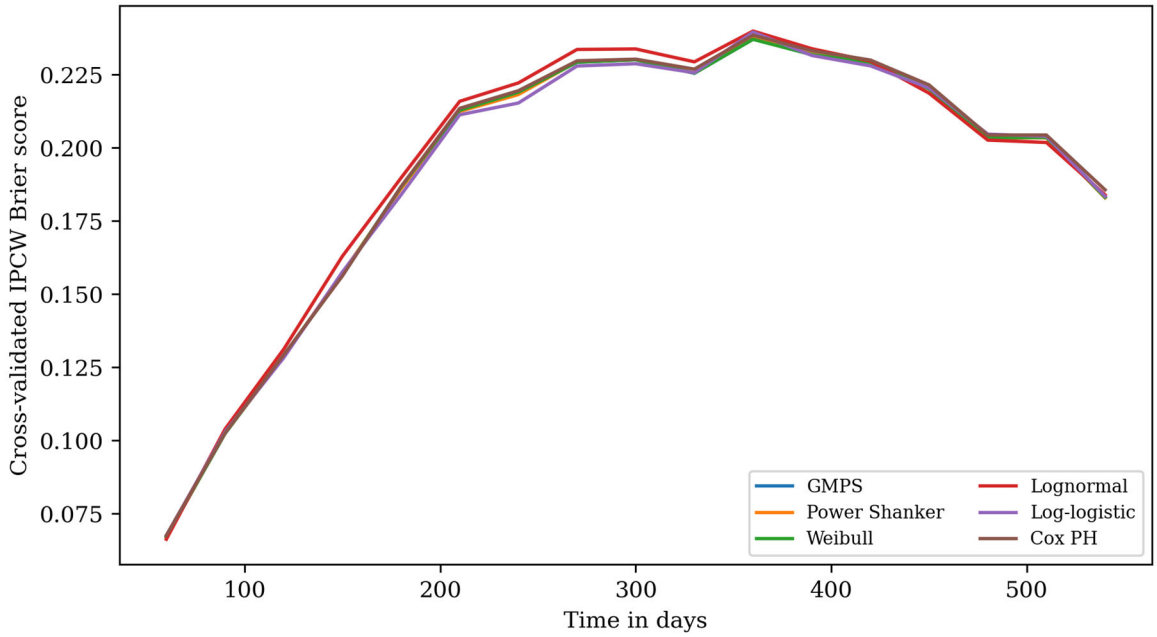


Figure 13. Repeated cross-validated IPCW Brier-score curves for the lung-cancer regression models.

As shown in Figure 13, the IPCW Brier-score curves for the competing models are tightly grouped over the prediction horizon. Prediction error rises toward the middle of follow-up and later declines, while no model separates decisively; this agrees with the conclusion that GMPS does not provide a material predictive advantage in this dataset.

Calibration at 365 days is broadly ordered but not perfect. For GMPS, the five predicted-survival quintiles have mean predicted survival of approximately 0.207, 0.326, 0.424, 0.535, and 0.658, versus Kaplan–Meier estimates of 0.268, 0.318, 0.398, 0.510, and 0.547. The highest

predicted-survival quintile is therefore optimistic by about 0.11. This is precisely the type of discrepancy that in-sample AIC alone cannot reveal [10,11].

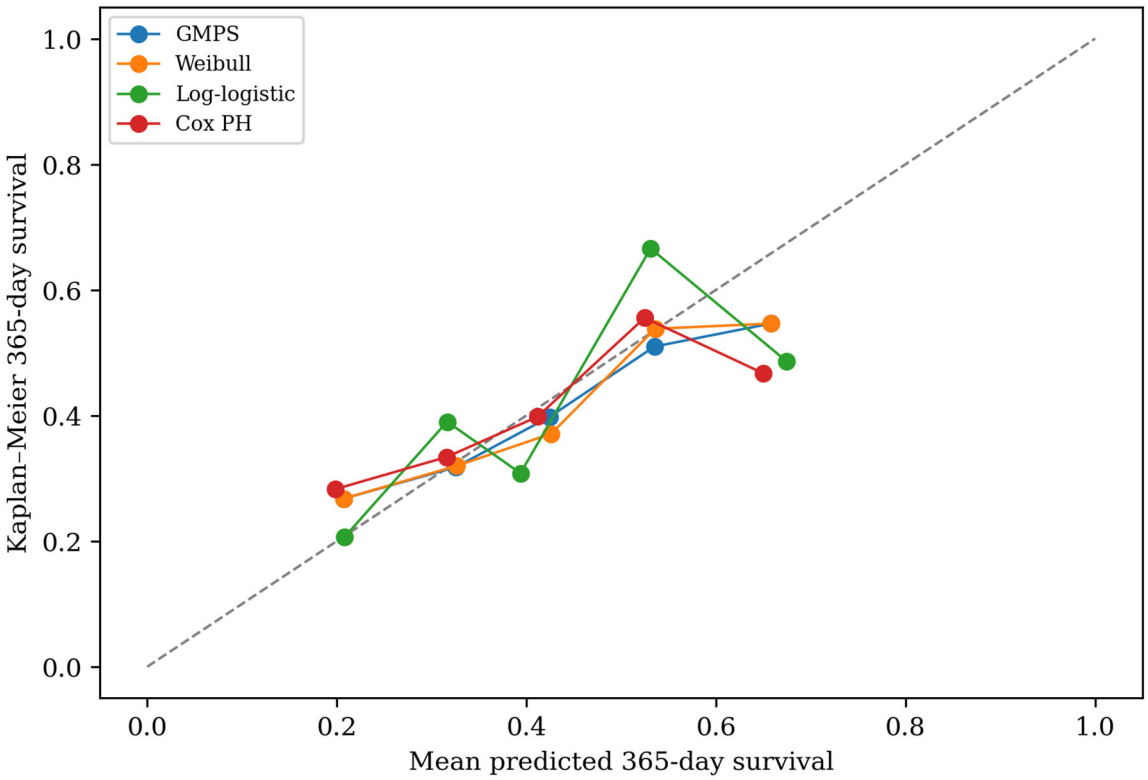


Figure 14. Cross-validated 365-day calibration by prediction quintile for selected lung-cancer models.

Figure 14 compares cross-validated predicted 365-day survival with Kaplan-Meier estimates across prediction quintiles. The points follow the 45-degree line reasonably well overall, although the highest-survival groups show some overprediction, indicating that calibration is acceptable but not perfect at the upper end of predicted survival.

7. Boundary-Aware Inference for the Weibull Limit

The lung application motivates inference on the continuous closure. The mathematical issue is whether a one-sided shape constraint is compatible with identifiable nuisance parameters and nonzero efficient information. This section establishes the local structure and states the additional sampling assumptions needed for the standard boundary limit; the new stress study shows why a finite-sample calibration is still needed.

Proposition 1. *Continuity and local boundary information*

For each fixed $x > 0$, the density and survival function extend continuously to $\alpha = 0$ on the parameter space (8), with the Weibull forms

$$f(x; 0, \theta, \delta) = \delta\theta x^{\delta-1} e^{-\theta x^{\delta}}, \qquad S(x; 0, \theta, \delta) = e^{-\theta x^{\delta}}. \tag{31}$$

At the Weibull boundary, the rate and power parameters are identifiable. For complete observations, the nuisance Fisher information is nonsingular, and the efficient information for the mixture shape is strictly positive, with the expression in (34). In rate regression, coefficient identification additionally requires a design matrix of full column rank.

Proof. To verify this, use $x^{\alpha\delta} \to 1$, $D(0, \theta) = \theta + 1$, and $\Gamma(1, z) = e^{-z}$ in (5) and (7). At the null, equality of two survival functions implies $\theta_1 x^{\delta_1} = \theta_2 x^{\delta_2}$ for every positive x; hence

the rate and power are identifiable. Under regression, the same argument identifies the power and the fitted rates, and identifies the coefficients when the design matrix has full column rank. There is no freely varying mixture weight to become unidentified when the components coincide: the weight is fixed by the same rate and shape parameters.

For complete observations, put $Z = \theta X^{\delta}$, so Z is unit exponential under the null, and let γ_E be Euler's constant. Direct differentiation gives the individual null scores

$$s_{\alpha} = \frac{\log Z + \gamma_E}{1+\theta}, \quad s_{\log\theta} = 1 - Z, \quad s_{\log\delta} = 1 + (1-Z)(\log Z - \log\theta). \tag{32}$$

These functions have finite second moments and are linearly independent. More explicitly, set $v = \pi^2/6, a = 1-\gamma_E-\log\theta$, and $\nu = (\log\theta, \log\delta)$. Unit-exponential logarithmic moments give

$$I_{\nu\nu} = \begin{pmatrix} 1 & a \\ a & a^2+v \end{pmatrix}, \quad I_{\alpha\nu} = \frac{(-1,\ 1-a)}{1+\theta}, \quad I_{\alpha\alpha} = \frac{v}{(1+\theta)^2}. \tag{33}$$

$$I_{\alpha\cdot\nu} = I_{\alpha\alpha} - I_{\alpha\nu} I_{\nu\nu}^{-1} I_{\nu\alpha} = \frac{v - 1 - v^{-1}}{(1+\theta)^2} > 0. \tag{34}$$

The numerator in (34) is only 0.037006965. Thus, the nuisance information is nonsingular, but profiling out the nuisance parameters removes most of the apparent shape information. At rate 1.2, the efficient shape information per observation is approximately 0.00764607. This provides an analytic explanation for the wide shape intervals and slow finite-sample estimation, without claiming that positive information guarantees useful precision. □

Regularity conditions for the subsequent likelihood-ratio theory. Local differentiability can be checked by extending the formula a small distance to $\alpha < 0$ with $\alpha > -1$: its gamma integrals remain finite, and the density stays positive. On compact positive rate and power neighborhoods and a sufficiently small neighborhood of zero shape, derivatives are bounded by integrable combinations of powers of Z, $|\log Z|$, and an exponential tail. This gives differentiability in quadratic mean and finite information at the null. In regression, bounded covariates and a full-rank limiting design are required. Under independent, parameter-distinct censoring, observed scores are the conditional expectations of complete-data scores given the recorded time and event indicator. Nondegeneracy additionally requires failures with positive probability over a nonempty time interval and covariate support sufficient to identify the regression coefficients; censoring may reduce information substantially. A singular design or effectively absent failures invalidates the usual reference.

7.1. Likelihood-Ratio Theory at $\boldsymbol{\alpha = 0}$

Write $\boldsymbol{\psi}$ for all nuisance parameters, including δ and the regression coefficients. The test is

$$H_0 : \alpha = 0 \qquad \text{versus} \qquad H_1 : \alpha > 0. \tag{35}$$

The boundary likelihood-ratio statistic is

$$\Lambda = 2\{\ell(\hat{\alpha}, \hat{\boldsymbol{\psi}}) - \ell(0, \hat{\boldsymbol{\psi}}_0)\}, \qquad \hat{\alpha} \geq 0. \tag{36}$$

With independent observations, the local differentiability and positive efficient information just described, identifiable interior nuisance parameters, and a consistent constrained likelihood maximizer in a regular neighborhood, the one-boundary-coordinate likelihood-ratio theorem [20,21] gives

$$\Lambda \xrightarrow{d} 1/2\chi_0^2 + 1/2\chi_1^2, \tag{37}$$

Here χ_0^2 is a point mass at zero. The tangent cone restricts one otherwise regular score coordinate to a half-line; its efficient normal score is positive with probability one half,

producing (37). The asymptotic 5% critical value is 2.7055. This is a local asymptotic result: it also requires the selected likelihood maximizer not to escape to a remote or diverging parameter sequence. The proof of local information alone does not establish uniform finite-sample calibration over all rates, powers, censoring designs, or numerical searches. Table 6 directly examines this distinction.

An upper profile endpoint may be summarized by finding U on the descending-likelihood side of the estimate, where the profile deviance reaches a stated cutoff:

$$\mathcal{U} = [0, U], \qquad R_p(U) = c, \quad U \geq \hat{\alpha}. \tag{38}$$

The reported application uses either the asymptotic cutoff 2.7055 or the fitted-null bootstrap cutoff. These are upper profile summaries. A fixed boundary cutoff does not define a uniformly calibrated two-sided 95% confidence set for positive shapes, and a bootstrap performed only at zero does not by itself establish 95% coverage of an upper endpoint at every positive shape. Nominal upper-bound coverage would require an appropriate signed-root or hypothesis-by-hypothesis calibration.

7.2. Parametric-Bootstrap Calibration Under the Lung Regression

The fitted Weibull null uses the same four standardized lung covariates, with seed 2026091805 and independently simulated censoring from a reverse-status Weibull fit. The reanalysis retains the same 600 simulated samples but refits each alternative from three shape starts, 0, 0.5, and 2, initialized at the null nuisance estimate. L-BFGS-B uses a shape bound of 50, log-power bounds $[-4, 4]$, intercept bounds $[-18, 8]$, coefficient bounds $[-5, 5]$, 500 iterations, function tolerance 10^{-10}, and gradient tolerance 5×10^{-6}. The exact null is always retained as a candidate. A numerical LR below 10^{-6} is counted as zero. Of 600 samples, 599 have successful null and alternative optimizer flags. This search changes 72 LR values by more than 10^{-4} relative to the earlier single-start bootstrap; the largest change is 0.6504.

Thus, 12.0% of the 600 likelihood-ratio values changed by more than the stated numerical tolerance when the starting values were expanded. Together with the change in zero mass from 0.521 to 0.439, this demonstrates sensitivity of the fitted likelihood-ratio distribution to initialization. The unchanged selected cutoff and reported 5% rejection fraction are limited aggregate findings; they do not establish stability of individual likelihood ratios, insensitivity of the original complete-data RMSEs, or attainment of a global maximum. These bootstrap results concern the lung fitted-null design and cannot quantify initialization effects in Table 3. The resulting calibration quantities are summarized in Table 16.

Table 16. Lung fitted-null bootstrap reanalysis using three alternative starts in each simulated sample. The rejection rate at the theoretical 5% threshold is stable, but the empirical zero mass differs from one half. The linear-interpolated 95th percentile and the higher-order-statistic calibration cutoff are reported separately.

Calibration Quantity	Value
Requested bootstrap replicates	600
Successful null refits	599
Empirical mass at LR = 0	0.439
Theoretical boundary mass at LR = 0	0.500
Empirical 90th percentile	1.647
Theoretical 90th percentile	1.642

Table 16. *Cont.*

Calibration Quantity	Value
Empirical 95th percentile (linear interpolation)	2.709
Theoretical 95th percentile	2.706
Empirical rejection at theoretical 5% cutoff	0.050
Calibration cutoff (higher order statistic)	2.769

Figure 15 shows the reanalyzed bootstrap LR distribution with its positive histogram scaled to unconditional density; its missing probability mass is at zero. The empirical zero mass is 0.439 rather than the earlier single-start value 0.521. The selected calibration cutoff remains 2.769, close to the theoretical 2.706, and the rejection fraction remains 0.050. Threshold agreement therefore does not imply that every part of the finite-sample distribution matches the asymptotic reference.

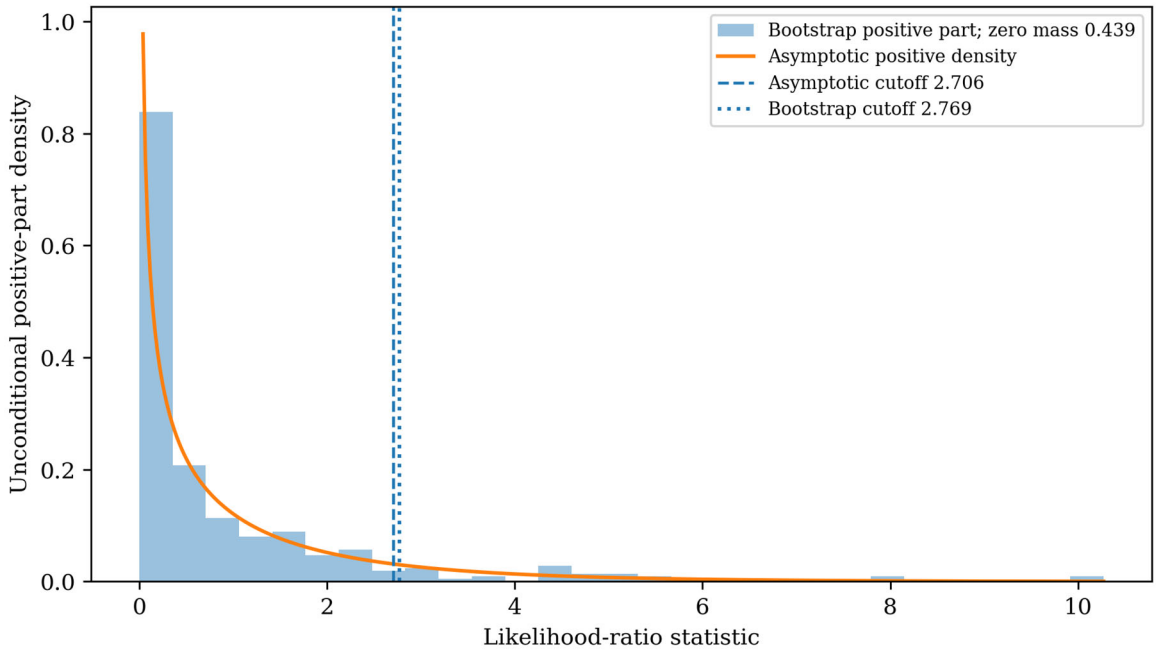


Figure 15. Parametric-bootstrap calibration of the boundary likelihood-ratio null distribution.

The reported lung optimum is $\alpha = 0$, with LR zero and boundary and bootstrap $p = 1.000$. The upper profile endpoints are 0.939 and 0.958 with the asymptotic and fitted-null bootstrap cutoffs, respectively. They summarize likelihood support under the chosen calibration and are not presented as a uniform coverage guarantee. The data support the simpler Weibull fit, while allowing some positive shapes in the high-likelihood region.

The boundary profile in Figure 16 is nearly flat close to zero and rises into the interior. The two upper endpoints are similar. This is a likelihood-support statement for this application, not evidence of uniformly precise shape estimation or uniform confidence coverage.

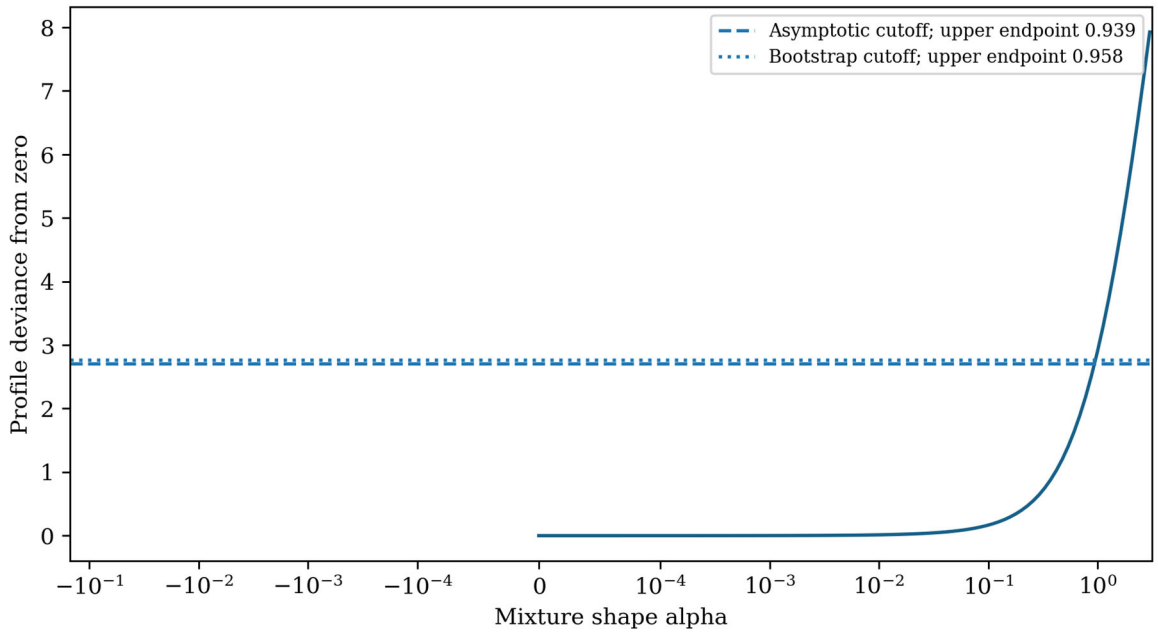


Figure 16. Boundary profile deviance for the NCCTG lung regression, with upper endpoints obtained from the asymptotic and fitted-null bootstrap cutoffs.

7.3. What the Boundary Test Says Across the Applications

Table 17 reports asymptotic boundary tests across the six applications as exploratory comparisons. Dataset-specific bootstrap calibrations are preferable when finite-sample shape information is weak; the lung calibration cannot be transferred automatically to the other datasets.

Table 17. Boundary likelihood-ratio tests of the exact Weibull submodel. The asymptotic p-values use $0.5\chi^2_0 + 0.5\chi^2_1$.

Dataset	n	$\hat{\alpha}$	LR for $\alpha = 0$	Boundary p-Value
Glass fibres	63	4.892	5.5051	0.0095
LA rainfall	85	14.19	10.3468	0.0006
Carbon fibres	100	0.5542	0.3834	0.2679
Bladder remission	128	0.9759	3.6629	0.0278
Heart transplant	69	0.523	0.3701	0.2715
NCCTG lung	228	0	0.0000	1.0000

The asymptotic comparisons indicate departures from Weibull for glass fibres, rainfall, and bladder remission. Only glass fibres and rainfall also reject Power Shanker in Table 9. Bladder remission is therefore compatible with the simpler Power Shanker model, despite evidence against Weibull. Moreover, rainfall fails the calibrated KS and CvM adequacy tests, and glass fibres are borderline by KS. Detectable departure from a nested baseline is not evidence of adequate absolute fit or predictive improvement.

8. Discussion

The boundary analysis supplies a model-reduction question that complements complete-data model comparison and censored regression. The exact Weibull closure permits an explicit constrained fit. Proposition 1 shows that nuisance parameters remain identifiable and the efficient shape information is positive, but small. These facts

explain why theoretical identifiability can coexist with a flat likelihood and substantial finite-sample uncertainty.

The reanalyzed 599 successful lung bootstrap samples give a zero mass of 0.439, linear-interpolated 95th percentile of 2.709, higher-order-statistic calibration cutoff of 2.769, and rejection fraction of 0.050 at the theoretical cutoff. At a rejection probability near 0.05, 599 replications give a Monte Carlo standard error of approximately 0.009. The earlier zero mass of 0.521 came from a single-start alternative search. Repeating the search from three starts changes the atom without changing the reported 5% decision. This supports design-specific calibration and numerical sensitivity checks, rather than treating agreement with the asymptotic atom as established.

The reported lung closure maximum occurs at zero, and the additional shape is unsupported there. Its upper profile endpoint describes the likelihood support under a chosen cutoff. Across the complete-data applications, the shape estimate is most weakly determined for carbon fibres, whose very broad local interval should preclude a precise substantive interpretation. Model selection should be paired with uncertainty in fitted survival or quantiles, rather than with the shape point estimate alone.

Testing zero shape asks about departure from Weibull. Testing unit shape asks whether Power Shanker is sufficient within GMPS. Absolute goodness of fit and out-of-sample prediction are separate questions. Bladder remission illustrates the distinction: it rejects the Weibull boundary under the asymptotic reference, does not reject Power Shanker, and provides no justification for retaining the full three-parameter baseline on that comparison alone.

Limitations include weak finite-sample shape information, model-dependent bootstrap calibration, and lack of a uniform coverage result for the reported upper profile endpoints. The cross-application boundary p-values are asymptotic and exploratory. The near-boundary study, multi-start bootstrap, and training-only cross-validation analyses are documented by their design, optimization settings, and reported diagnostic summaries. The original complete-data simulation uses favorable generating-parameter starts. Its RMSE conclusions are conditional on that scheme because no matched refitting of those scenarios from alternative starts has been performed. The separate multi-start stress study does not remove this limitation, and the lung bootstrap directly demonstrates sensitivity to starting values.

The executable code, analysis datasets, pinned dependencies, and numerical outputs for the simulations, bootstrap calculations, profile analyses, and predictive validation are available in public release 1.0.0 at https://doi.org/10.6084/m9.figshare.34718166. The release includes a single command for the full analysis and checks against the reported numerical tables. Re-execution was verified in the stated environment; this does not constitute independent third-party reproduction or guarantee bitwise identical results across platforms. The initialization, identification, calibration, and coverage limitations described above remain applicable.

9. Conclusions

GMPS extends the Shanker mixture by varying its latent gamma shape and applying a power transformation. Exact nesting, tractable likelihood evaluation, simulation, and rate regression make it possible to compare the full family with its Power Shanker member and Weibull boundary. The boundary derivation identifies positive but small efficient shape information, and the new stress study demonstrates the resulting uncertainty and limited detection power.

The multi-start lung bootstrap is consistent with the asymptotic reference at the reported 5% threshold, but its zero mass is 0.439 instead of one half. The near-boundary study

likewise shows that the local asymptotic reference should not be generalized to every finite sample. The lung fit selects Weibull, and training-only repeated cross-validation shows no demonstrated predictive gain from GMPS. Upper profile endpoints remain likelihood-based summaries under stated cutoffs, without a uniform coverage claim.

The empirical recommendation is conditional: check absolute adequacy, compare both nested hypotheses, inspect the shape profile, and assess held-out prediction before retaining an additional parameter. Rainfall and glass fibres require adequacy qualifications; bladder remission can retain Power Shanker; and the lung application favors Weibull. These conclusions describe the evidence in the examined samples and do not assert universal superiority.

Funding: This research was funded by the Deanship of Postgraduate Studies and Scientific Research at Majmaah University, grant number R-2026-514.

Institutional Review Board Statement: The Institutional Review Board (IRB) approval was waived because this study involved only secondary analysis of previously published, publicly available, de-identified datasets, including the Stanford heart-transplant and NCCTG lung-cancer datasets. No new clinical data or biological materials were collected, and no intervention involving human participants was performed.

Informed Consent Statement: Informed consent was waived because the study used only publicly available, de-identified data. No participants were recruited or contacted, and no attempts were made to identify individuals or link the datasets to identifiable personal records.

Data Availability Statement: The complete data applications use the published sources cited in Section 6.1 and References [12–15]. The Stanford heart-transplant dataset is publicly accessible through statsmodels at https://www.statsmodels.org/stable/datasets/generated/heart.html (accessed on 25 October 2026). The NCCTG lung dataset is publicly accessible through the R survival package at https://cran.r-project.org/package=survival (accessed on 25 October 2026). The distributional definitions, likelihoods, simulation designs, optimization constraints, initialization rules, convergence tolerances, random seeds, and software versions used in the new analyses are described in Sections 2–7. The exact numerical input files used in this study, with source attribution and event coding, are included in the code release cited below. Public release 1.0.0 is available at https://doi.org/10.6084/m9.figshare.34718166. It contains the GMPS likelihoods and exact sampler; complete-data, censored-data, regression, and near-boundary simulations; refitted goodness-of-fit and multi-start boundary bootstraps; profile likelihoods; training-only repeated cross-validation; the six analysis datasets; pinned software requirements; seeds and optimization settings; figure-generation scripts; numerical outputs; and a README with the full-run command and a table-to-output map. The release records the verified environment and comparison checks so readers can execute and inspect the calculations.

Acknowledgments: The author extends the appreciation to the Deanship of Postgraduate Studies and Scientific Research at Majmaah University for funding this research work through the project number (R-2026-514).

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Shanker, R. Shanker distribution and its applications. *Int. J. Stat. Appl.* **2015**, *5*, 338–348.
2. Shanker, R.; Shukla, K.K. Power Shanker distribution and its application. *Turk. Klin. J. Biostat.* **2017**, *9*, 175–187. [CrossRef]
3. Aforka, K.F.; Semary, H.E.; Onyeagu, S.I.; Etaga, H.O.; Obulezi, O.J.; Al-Moisheer, A.S. A new exponentiated power distribution for modeling censored data with applications to clinical and reliability studies. *Symmetry* **2025**, *17*, 1153. [CrossRef]
4. Prodhani, H.R.; Shanker, R. A three-parameter Power Sujatha distribution with properties and application. *Int. J. Stat. Reliab. Eng.* **2024**, *11*, 70–78.
5. Al-Omari, A.I.; Dobbah, S.A. On the mixture of Shanker and gamma distributions with applications to engineering data. *J. Radiat. Res. Appl. Sci.* **2023**, *16*, 100533. [CrossRef]
6. Prodhani, H.R.; Shanker, R. A generalized power Sujatha distribution with properties and applications. *Reliab. Theory Appl.* **2025**, *20*, 373–387. [CrossRef]

7. Prodhani, H.R.; Shanker, R. Power quasi Sujatha distribution with properties and applications to real lifetime data. *Stat. Transit. New Ser.* **2026**, *27*, 53–70. [CrossRef]
8. Stacy, E.W. A generalization of the gamma distribution. *Ann. Math. Stat.* **1962**, *33*, 1187–1192. [CrossRef]
9. Ghitany, M.E.; Al-Mutairi, D.K.; Balakrishnan, N.; Al-Enezi, L.J. Power Lindley distribution and associated inference. *Comput. Stat. Data Anal.* **2013**, *64*, 20–33. [CrossRef]
10. Harrell, F.E., Jr.; Lee, K.L.; Mark, D.B. Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Stat. Med.* **1996**, *15*, 361–387. [CrossRef]
11. Graf, E.; Schmoor, C.; Sauerbrei, W.; Schumacher, M. Assessment and comparison of prognostic classification schemes for survival data. *Stat. Med.* **1999**, *18*, 2529–2545. [CrossRef]
12. Smith, R.L.; Naylor, J.C. A comparison of maximum likelihood and Bayesian estimators for the three-parameter Weibull distribution. *J. R. Stat. Soc. Ser. C* **1987**, *36*, 358–369. [CrossRef]
13. Nichols, M.D.; Padgett, W.J. A bootstrap control chart for Weibull percentiles. *Qual. Reliab. Eng. Int.* **2006**, *22*, 141–151. [CrossRef]
14. Lee, E.T.; Wang, J.W. *Statistical Methods for Survival Data Analysis*, 3rd ed.; John Wiley & Sons: Hoboken, NJ, USA, 2003.
15. Hashem, A.F.; Alyami, S.A. Inference on a new lifetime distribution under progressive Type II censoring for a parallel-series structure. *Complexity* **2021**, *2021*, 6684918. [CrossRef]
16. Seabold, S.; Perktold, J. Statsmodels: Econometric and statistical modeling with Python. In Proceedings of the 9th Python in Science Conference 2010, Austin, TX, USA, 28 June–3 July 2010; pp. 92–96.
17. Miller, R.G. Least squares regression with censored data. *Biometrika* **1976**, *63*, 449–464. [CrossRef]
18. Loprinzi, C.L.; Laurie, J.A.; Wieand, H.S.; Krook, J.E.; Novotny, P.J.; Kugler, J.W.; Bartel, J.; Law, M.; Bateman, M.; Klatt, N.E. Prospective evaluation of prognostic variables from patient-completed questionnaires. North Central Cancer Treatment Group. *J. Clin. Oncol.* **1994**, *12*, 601–607. [CrossRef] [PubMed]
19. Cox, D.R. Regression models and life-tables. *J. R. Stat. Soc. Ser. B* **1972**, *34*, 187–220. [CrossRef]
20. Self, S.G.; Liang, K.-Y. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *J. Am. Stat. Assoc.* **1987**, *82*, 605–610. [CrossRef]
21. Chernoff, H. On the distribution of the likelihood ratio. *Ann. Math. Stat.* **1954**, *25*, 573–578. [CrossRef]

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.