AppliedMath

MDPI

Article

Exact Beta-P Margins for Dependent Competing Risks: Partial Identification and Sensitivity Bounds

Yousef Al-Zalzalah

Department of Information Systems and Operations Management, College of Business Administration, Kuwait University, Safat 13055, Kuwait; yousef.alzalzalah@ku.edu.kw

Abstract

Competing-risks data identify the earliest failure time and its cause, but generally do not identify the dependence among latent failure times. We construct a shared-gamma frailty model that preserves distinct Beta-P marginal survival distributions exactly. Within this model, we derive the joint survival function, cause-specific subdensities, and net, crude, and partial-crude risk functionals. We establish positive upper-orthant dependence and an ordering result for nested active-risk sets. We also define model-specific structural sensitivity measures and likelihood-supported probability regions. Using a publicly available transplantation dataset ($n = 177$), we compare the Aalen–Johansen estimates with the independent Beta-P boundary model and examine dependence through a multi-start profile analysis that re-estimates the marginal parameters at each of 25 grid values. The best likelihood found occurs at the independence boundary (log-likelihood approximately -561.2107); the numerical search found no improvement for positive dependence. We report descriptive profile-grid probability ranges at 12 and 24 months and separate fixed-margin sensitivity ranges. Seeded Monte Carlo studies check the probability calculations and the complete estimator. The formula check has a maximum absolute discrepancy of 0.0016; a separate 300-replicate study obtains reproducible constrained optimization but reveals substantial boundary selection and weak complete-estimator identification at $n = 177$. The profile-grid ranges are neither confidence intervals nor certified sharp bounds. The model illustrates how assumptions about latent dependence affect latent and observable risk functionals; it does not demonstrate improved clinical prediction.

Keywords: shared gamma frailty; latent failure times; active-risk functionals; profile likelihood; cumulative incidence

1. Introduction

In a competing-risks experiment, an individual is exposed to several mutually exclusive causes of failure. If the latent failure time associated with cause j is T_j, only

$$T = \min_{1 \le j \le k} T_j \tag{1}$$

Only the earliest failure time and its cause are observed; the remaining latent failure times are censored by the occurrence of that first event. This feature distinguishes competing-risks analysis from ordinary multivariate survival analysis and makes dependence among the latent failure times difficult to identify.

Observable cumulative-incidence functions can be estimated without specifying a joint distribution for the latent failure times. Net failure probabilities, cause-elimination probabilities, and other latent quantities require additional assumptions. Classical results

Check for updates

Received: 21 October 2026
Revised: 6 November 2026
Accepted: 11 November 2026
Published: 13 November 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.

show that the joint distribution of latent failure times is not nonparametrically identifiable from the observed minimum and its cause alone [1,2]. Parametric models provide an interpretable joint structure, but the data may contain little information about their dependence parameters. Accordingly, a fitted point estimate should be accompanied by an assessment of how substantive probabilities vary across plausible dependence values.

The Beta-P distribution considered here is the Burr type XII survival family

$$S_j(t) = \left[1 + \left(\frac{t}{\beta_j}\right)^{\theta_j}\right]^{-\alpha_j}, \qquad t > 0, \tag{2}$$

where α_j, β_j, and θ_j are positive. The family accommodates non-exponential tail behavior and a range of hazard shapes. Earlier work used this family under independent competing risks. The present study addresses a different question: how to introduce latent dependence while preserving these marginal distributions exactly and how to report the resulting uncertainty without overstating point identification.

Recent dependent competing-risks models with Burr-family margins have focused mainly on Marshall–Olkin constructions, specialized censoring schemes, and likelihood or Bayesian estimation [3,4]. These studies demonstrate the relevance of flexible Burr margins, but they do not remove the classical limitation: the observed minimum and its cause provide limited information about dependence among latent failure times. Our aim is not to propose another bivariate Burr distribution. Instead, we study probability sensitivity within a correlated-frailty family whose Beta-P margins remain fixed as dependence varies.

We introduce dependence through a static shared-gamma frailty decomposition. Additive gamma decompositions and frailty models for competing and multivariate failure times are established methods [5–8]. Our construction specializes this framework to preserve distinct Beta-P margins exactly and to derive explicit active-risk probabilities and sensitivity measures. It is a static model: it does not include a time-evolving latent process, such as a Markov chain, physiological-age trajectory, degradation process, jump mechanism, or longitudinal latent measurement.

The next subsection summarizes the paper's specific contributions and identifies the results that support each claim. It also distinguishes the established gamma-frailty mechanism from the exact-margin and active-risk analyses developed here.

Main Contributions

Exact marginal preservation and dependence characterization. Section 2 presents a shared-component gamma construction that preserves potentially distinct Beta-P survival margins for every admissible dependence value (Proposition 1). Proposition 2 establishes positive upper-orthant dependence. Our novelty claim concerns this exact-margin specialization and its integration with competing-risks functionals, not the gamma-frailty decomposition itself.

Unified active-risk probability framework. Section 3 derives cause-specific subdensities and interval probabilities for arbitrary active-risk subsets, covering net, crude, and partial-crude risks within a single formulation (Sections 3.1–3.3). Proposition 3 establishes an ordering for nested active-risk sets. Section 3.5 derives the dependence derivative and explains why its mixed-sign terms alone establish neither monotonicity nor nonmonotonicity of an interval probability.

Model-specific sensitivity and partial-identification analysis. Section 4 distinguishes classical nonparametric nonidentification from the likelihood information available within the specified parametric family. It defines fixed-margin structural sensitivity extrema and likelihood-supported probability regions. Section 6 reports numerical one-dimensional fixed-margin extrema with endpoint and dense-grid checks, together with descriptive

profile-grid ranges. These quantities are not sharp distribution-free bounds or calibrated confidence regions.

Empirical scope. Sections 5 and 6 present the independent Beta-P benchmark, numerical fitting and multi-start profiling of the shared-gamma model, fixed-margin sensitivity analyses, and reestimated-parameter probability ranges. The best likelihood found for the dependent model occurs at the independence boundary. The supplementary files document the computations and diagnostics. We do not claim global optimality or calibrated profile-likelihood confidence coverage.

2. The Exact-Margin Dependence Model

2.1. Beta-P Marginal Distributions

For cause j, define

$$H_j(t) = \left(\frac{t}{\beta_j}\right)^{\theta_j} \tag{3}$$

and

$$\dot{H}_j(t) = \frac{\theta_j}{\beta_j}\left(\frac{t}{\beta_j}\right)^{\theta_j - 1}. \tag{4}$$

The Beta-P survival and density functions are

$$S_j(t) = \left[1 + H_j(t)\right]^{-\alpha_j} \tag{5}$$

and

$$f_j(t) = \alpha_j \dot{H}_j(t)\left[1 + H_j(t)\right]^{-(\alpha_j+1)}. \tag{6}$$

The cumulative hazard is

$$R_j(t) = \alpha_j \log[1 + H_j(t)]. \tag{7}$$

2.2. Static Gamma Decomposition

Let

$$V_0 \sim \text{Gamma}(\delta,1) \tag{8}$$

and, independently,

$$V_j \sim \text{Gamma}(\alpha_j - \delta,1), \qquad j = 1, \ldots, k, \tag{9}$$

where

$$0 \le \delta \le \min_j \alpha_j. \tag{10}$$

At the boundary, zero-shape gamma components are defined as point masses at zero: $V_0 \equiv 0$ when $\delta = 0$, and $V_j \equiv 0$ when $\delta = \alpha_j$. These cases are also the weak limits of the corresponding positive-shape gamma distributions. Thus, the model does not invoke an improper $\text{Gamma}(0,1)$ distribution.

Define the cause-specific frailty

$$U_j = V_0 + V_j. \tag{11}$$

Conditional on the full frailty vector $(U_1, \ldots, U_k)$, the latent failure times $T_1, \ldots, T_k$ are mutually independent.

Given U_j, the conditional survival function is

$$\Pr(T_j > t | U_j) = \exp[-U_j H_j(t)]. \tag{12}$$

The common component V_0 induces positive association. The variables V_j retain cause-specific heterogeneity.

Because $U_j \sim \text{Gamma}(\alpha_j, 1)$,

$$\text{Var}(U_j) = \alpha_j, \qquad \text{Cov}(U_j, U_\ell) = \delta \quad (j \neq \ell), \tag{13}$$

and hence

$$\text{Corr}\left(U_j, U_\ell\right) = \frac{\delta}{\sqrt{\alpha_j \alpha_\ell}}. \tag{14}$$

Thus, each U_j reflects both a gamma-component shape and the dependence among frailty multipliers, subject to the admissibility constraint above. For distinct causes i and j, $\text{Cov}(U_i, U_j) = \delta$ and $\text{Corr}(U_i, U_j) = \delta/(\alpha_i \alpha_j)^{\wedge}(1/2)$. This is the correlation between frailty multipliers, not the rank correlation between latent failure times. We characterize the induced failure-time association through the joint survival function and the positive upper-orthant dependence result.

2.3. Joint Survival Function

Conditioning on the frailties and using gamma Laplace transforms gives

$$S_\delta(t_1, \ldots, t_k) = \left[1 + \sum_{j=1}^{k} H_j\left(t_j\right)\right]^{-\delta} \prod_{j=1}^{k} \left[1 + H_j\left(t_j\right)\right]^{-(\alpha_j - \delta)}. \tag{15}$$

When $\delta = 0$,

$$S_0(t_1, \ldots, t_k) = \prod_{j=1}^{k} S_j\left(t_j\right), \tag{16}$$

so independence is a boundary case.

2.4. Exact Preservation of Margins

Proposition 1 (exact Beta-P margins). *For every admissible δ and each cause j, the marginal survival function induced by S_δ is $S_j(t) = \left[1 + H_j(t)\right]^{-\alpha_j}$.*

Proof. Setting all arguments except t_j equal to zero yields

$$S_\delta\left(0, \ldots, t_j, \ldots, 0\right) = \left(1 + H_j\left(t_j\right)\right)^{-\delta} \left(1 + H_j\left(t_j\right)\right)^{-(\alpha_j - \delta)} = \left(1 + H_j\left(t_j\right)\right)^{-\alpha_j}. \tag{17}$$

Thus, every marginal survival function is exactly Beta-P and is independent of δ. Dependence sensitivity can therefore be studied without changing the selected margins. □

2.5. Positive Quadrant Dependence

Proposition 2 (positive upper-orthant dependence). *For all nonnegative $(t_1, \ldots, t_k)$ and every $\delta \geq 0$, $S_\delta(t_1, \ldots, t_k) \geq \prod_j S_j\left(t_j\right)$.*

Proof. The ratio of joint survival to independent joint survival is

$$\frac{S_\delta(t_1, \ldots, t_k)}{\prod_j S_j\left(t_j\right)} = \left\{ \frac{\prod_j \left[1 + H_j\left(t_j\right)\right]}{1 + \sum_j H_j\left(t_j\right)} \right\}^{\delta}. \tag{18}$$

Since

$$\prod\nolimits_j \left(1 + x_j\right) \geq 1 + \sum\nolimits_j x_j \qquad \left(x_j \geq 0\right), \tag{19}$$

the ratio is at least one. The construction therefore produces positive upper-orthant dependence. It is not intended to represent arbitrary negative dependence. □

2.6. Relation to Standard Dependence Constructions

The proposed model represents dependence through a shared gamma component rather than through an unrestricted copula. The shared-component parameter is constrained by the marginal frailty shapes, and the construction yields closed-form cause-specific subdensities for nested active-risk sets. These features distinguish the parameterization and its interpretation; they do not show that the resulting joint distribution is mathematically distinct from every previously studied copula or frailty family. The proposed model and benchmark constructions are compared in Table 1.

Table 1. Comparison of the proposed dependence model with benchmark constructions.

Feature	Correlated-Gamma Beta-P Model	Independent Beta-P	Clayton-Copula Beta-P
Margins	Exact Beta-P	Exact Beta-P	Exact Beta-P
Dependence	Shared gamma component	None	Copula parameter
Independence	Boundary $\delta = 0$	Assumed	Boundary $\eta = 0$
Multi-cause active sets	Closed form	Closed form	Requires a chosen multivariate copula
Main role here	Proposed structural and sensitivity model	Nested parsimony benchmark	Alternative-family robustness benchmark

Therefore, we make a deliberately narrow novelty claim: the contribution is the exact-margin active-risk and sensitivity framework derived from the correlated-gamma construction, not the general use of gamma frailties or Burr-type margins.

3. Active-Risk Probability Functionals

3.1. Survival Under an Active Set

Let $K = \{1, \ldots, k\}$ and let $R \subseteq K$ be a nonempty set of causes remaining active. Define

$$A_R(t) = 1 + \sum_{\ell \in R} H_\ell(t). \tag{20}$$

The survival probability when only causes in R remain active is

$$S_R(t;\delta) = A_R(t)^{-\delta} \prod_{\ell \in R} [1 + H_\ell(t)]^{-(\alpha_\ell - \delta)}. \tag{21}$$

3.2. Cause-Specific Subdensity

For $j \in R$, differentiation gives

$$g_{j|R}(t;\delta) = S_R(t;\delta)\dot{H}_j(t)\left[\frac{\delta}{A_R(t)} + \frac{\alpha_j - \delta}{1 + H_j(t)}\right]. \tag{22}$$

At independence,

$$g_{j|R}(t;0) = f_j(t) \prod_{\ell \in R, \ell \neq j} S_\ell(t). \tag{23}$$

The identity

$$-\frac{d}{dt} S_R(t;\delta) = \sum_{j \in R} g_{j|R}(t;\delta) \tag{24}$$

implies

$$\sum_{j \in R} Q_{j|R}(a,b;\delta) = S_R(a;\delta) - S_R(b;\delta). \tag{25}$$

3.3. Unified Interval Probability

Define

$$Q_{j|R}(a,b;\delta) = \int_a^b g_{j|R}(t;\delta)\, dt. \tag{26}$$

This expression covers three probability concepts:

1. **Net probability:** $R = \{j\}$.
2. **Crude probability:** $R = K$.
3. **Partial-crude probability:** $\{j\} \subset R \subset K$.

For the net probability,

$$q_j(a,b) = S_j(a) - S_j(b), \tag{27}$$

which is independent of δ because the margins are fixed.

3.4. Ordering Under Cause Addition

Proposition 3 (nested active-risk ordering). *If* $\{j\} \subseteq R_1 \subseteq R_2$, *then for every interval* $[a,b]$ *and admissible* δ,

$$Q_{j|R_2}(a,b;\delta) \le Q_{j|R_1}(a,b;\delta) \le q_j(a,b). \tag{28}$$

Proof. Couple the three systems using the same latent vector $(T_1, \ldots, T_k)$. The event that the target cause occurs in $[a,b]$ before every other cause in the larger active set is contained in the corresponding event for the smaller active set, which is itself contained in the event that the target cause occurs in $[a,b]$. Taking probabilities proves the result. □

No general ordering is asserted between active sets that are not nested.

3.5. Dependence Derivatives

For fixed margins,

$$D_R(t) = \frac{\partial}{\partial \delta} \log S_R(t;\delta) = \sum\nolimits_{\ell \in R} \log[1 + H_\ell(t)] - \log A_R(t) \ge 0. \tag{29}$$

Thus, $S_R(t;\delta)$ increases with δ. Define

$$C_{j|R}(t;\delta) = \frac{\delta}{A_R(t)} + \frac{\alpha_j - \delta}{1 + H_j(t)}. \tag{30}$$

Then

$$\frac{\partial}{\partial \delta} \log g_{j|R}(t;\delta) = D_R(t) + \frac{A_R(t)^{-1} - \left[1 + H_j(t)\right]^{-1}}{C_{j|R}(t;\delta)}. \tag{31}$$

The first term is nonnegative and the second is nonpositive, so this decomposition does not determine the sign of the derivative. It does not show that the derivative changes sign for any admissible parameter configuration. Nor does pointwise ambiguity about the derivative of a subdensity establish nonmonotonicity of the corresponding interval probability. We therefore make no nonmonotonicity claim on this basis; such a claim would require a verified numerical example or a formal proof.

3.6. Mathematical Consistency Checks

For each active set R, the joint survival S_R(t;δ) is the Laplace transform E[exp{−Σ_{ℓ∈R} U_ℓ H_ℓ(t)}] of a nonnegative frailty vector. It is therefore a valid survival function, including at δ = 0 and δ = $\min_j \alpha_j$ when zero-shape components are interpreted as point masses at zero. In particular, S_R(0;δ) = 1, and S_R(t;δ) tends to zero as t→∞ when all α_j, β_j, θ_j are positive.

Writing A_R(t) = 1 + Σ_{ℓ∈R}H_ℓ(t), differentiation yields − d log S_R(t;δ)/dt = Σ_{j∈R} H_j(t){δ/A_R(t) + (α_j − δ)/(1 + H_j(t))}. Each summand is nonnegative for $0 \le \delta \le \min_j \alpha_j$. Consequently, g_{j|R}(t;δ) = S_R(t;δ)H_j(t){δ/A_R(t) + (α_j−δ)/(1 + H_j(t))} is nonnegative, and Σ_{j∈R}g_{j|R}(t;δ) = −S′_R(t;δ). Integrating over [a,b) proves Σ_{j∈R}Q_{j|R}(a,b;δ) = S_R(a;δ) − S_R(b;δ).

For nested sets $\{j\}\subseteq R_1\subseteq R_2$, couple the latent failure times using the same gamma components and the same conditional survival construction. The event $\{a \le T_j < b, T_j = \min_{\ell\in R_2} T_\ell\}$ is contained in the corresponding event for R_1, which is contained in $\{a \le T_j < b\}$. Absolute continuity implies ties have probability zero. This establishes Proposition 3 without assuming independence between latent failure times.

These checks prove validity and ordering within the specified shared-gamma model. They do not prove that δ is globally identified from competing-risks observations, that fixed-margin extrema are sharp over unrestricted dependence models, or that a cause-elimination intervention has a causal interpretation without additional assumptions.

4. Partial Identification, Estimation, and Sensitivity

4.1. Observed-Data Likelihood

For observation (X_i,Δ_i), where $\Delta_i = 0$ denotes right censoring and $\Delta_i = j$ denotes failure from cause j, the likelihood contribution is

$$L_i(\Theta) = S_K(X_i;\Theta)^{I(\Delta_i=0)} \prod_{j=1}^{k} g_{j|K}(X_i;\Theta)^{I(\Delta_i=j)}. \tag{32}$$

The log-likelihood is

$$\ell(\Theta) = \sum_i \log S_K(X_i;\Theta) + \sum_{i:\Delta_i>0} \{\log \dot{H}_{\Delta_i}(X_i) + \log C_{\Delta_i|K}(X_i;\delta)\}. \tag{33}$$

The parameter space is

$$\Omega = \left\{\alpha_j,\beta_j,\theta_j > 0;\ 0 \le \delta \le \min_j \alpha_j\right\}. \tag{34}$$

For numerical optimization, write

$$\alpha_j = \delta + \kappa_j, \qquad \kappa_j > 0, \tag{35}$$

The parameter space can be written as $\alpha_j = \delta + \kappa_j$, with $\delta \ge 0$ and $\kappa_j \ge 0$; equivalently, $0 \le \delta \le \min_j \alpha_j$. The baseline profile uses logarithmic transformations and therefore restricts κ_j to positive values. In the expanded numerical audit, we separately optimized the exact faces $\kappa_1 = 0$, $\kappa_2 = 0$, and $\kappa_1 = \kappa_2 = 0$ at each positive grid value, and evaluated $\delta = 0$ as the independent model. These checks address the component boundaries on the evaluated grid but do not establish a continuous-domain or global maximum.

4.2. Weak Information About Latent Dependence

The observable likelihood depends on $S_K(t)$ and $g_{j|K}(t)$, not on the complete joint distribution away from the observed minimum-time structure. Parametric restrictions may provide local identification, yet the likelihood can remain nearly flat in the dependence direction. Optimizer convergence is therefore not sufficient evidence that δ is precisely estimated.

For fixed δ, let

$$\hat{\boldsymbol{\psi}}_\delta = \arg\max_{\boldsymbol{\psi}} \ell(\boldsymbol{\psi},\delta), \tag{36}$$

where $\boldsymbol{\psi} = (\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\theta})$. The profile log-likelihood is

$$\ell_p(\delta) = \ell(\hat{\boldsymbol{\psi}}_\delta,\delta). \tag{37}$$

A likelihood-supported set is

$$C_\delta(c) = \left\{\delta : 2\left[\ell_p(\hat{\delta}) - \ell_p(\delta)\right] \le c\right\}. \tag{38}$$

For each fixed δ, a profile analysis must re-estimate the nuisance parameters subject to $\alpha_j \geq \delta$ and positivity of the remaining parameters. The admissible domain is therefore determined jointly by these constraints, not by the marginal-shape estimates from the independent model. We do not claim a numerically verified continuous profile endpoint.

Three distinct ideas must not be conflated. The mapping $\Theta \mapsto (S_K, g_{1|K}, \ldots, g_{k|K})$ may be formally identifiable inside the parametric family; the sample likelihood may nevertheless contain little curvature in the δ direction; and the unrestricted competing-risks problem remains nonparametrically unidentified. The regions below quantify the second issue under the stated model and are not asserted to be sharp identified sets over all joint latent-time distributions.

4.3. Structural Sensitivity Bounds

With margins fixed at $\hat{\boldsymbol{\psi}}$, define

$$\hat{Q}^{S,L}_{j|R} = \min_{0 \leq \delta \leq \hat{\delta}_{\max}} Q_{j|R}\left(a, b; \delta, \hat{\boldsymbol{\psi}}\right) \tag{39}$$

and

$$\hat{Q}^{S,U}_{j|R} = \max_{0 \leq \delta \leq \hat{\delta}_{\max}} Q_{j|R}\left(a, b; \delta, \hat{\boldsymbol{\psi}}\right), \tag{40}$$

where $\hat{\delta}_{\max} = \min_j \hat{\alpha}_j$. These are sensitivity endpoints within the proposed frailty family, not distribution-free bounds over every joint distribution having the same margins.

4.4. Profile-Supported Probability Regions

The profile-supported endpoints are

$$\hat{Q}^{PL,L}_{j|R} = \min_{\delta \in C_\delta(c)} Q_{j|R}\left(a, b; \delta, \hat{\boldsymbol{\psi}}_\delta\right) \tag{41}$$

and

$$\hat{Q}^{PL,U}_{j|R} = \max_{\delta \in C_\delta(c)} Q_{j|R}\left(a, b; \delta, \hat{\boldsymbol{\psi}}_\delta\right). \tag{42}$$

We do not report continuous-domain extrema for the reestimated profile. Section 6.4.3 reports only an exploratory range over the evaluated grid points that pass the stated likelihood-ratio screen.

4.5. Latent Sensitivity Index

Define

$$\mathrm{LSI}_{j|R}(a, b) = \frac{Q^{S,U}_{j|R}(a, b) - Q^{S,L}_{j|R}(a, b)}{q_j(a, b)}. \tag{43}$$

The LSI is not a correlation coefficient. When $q_j(a,b) > 0$, it is nonnegative because its numerator is the range of the fixed-margin probabilities, and it equals zero when the two endpoints coincide. The index is dimensionless and specific to the cause, interval, active set, marginal parameters, and admissible dependence range. It measures the width of the sensitivity range relative to the independence net probability; it does not indicate the direction or monotonicity of $Q_j|R(a,b;\delta)$. It has no universal finite upper bound because its denominator may be arbitrarily small. When $q_j(a,b) = 0$, the ratio is undefined, and the index is not reported.

4.6. Diagnostics and Boundary Inference

Let $\mathbf{m}(\Theta)$ stack $S_K(t)$ and $g_{j|K}(t)$ over an empirically supported time grid and let

$$\mathcal{D}(\Theta) = \frac{\partial \mathbf{m}(\Theta)}{\partial \Theta^\top}. \tag{44}$$

Rank and conditioning of the observable map, information eigenvalues, and profile curvature may help diagnose local weak identification. We do not report numerical values for these diagnostics, and none of them alone establishes global identifiability.

For testing H_0:$\delta = 0$ against H_1:$\delta > 0$, use

$$\Lambda = 2\left[\ell_p\left(\hat{\delta}\right) - \ell_p(0)\right]. \tag{45}$$

Because $\delta = 0$ is a boundary point, the usual Wilks χ^2_1 reference distribution does not generally apply. Under standard regularity conditions and positive efficient information for δ, the one-sided likelihood-ratio statistic has the asymptotic distribution $\frac{1}{2}\chi^2_0 + \frac{1}{2}\chi^2_1$ [9]. This result depends on those assumptions. We do not claim a calibrated test or a numerical profile-likelihood confidence set for the clinical application.

Among the numerical solutions examined, the fitted model reaches its best recorded likelihood at the boundary $\delta = 0$. Thus, the numerical search found no improvement in likelihood for positive dependence within the specified family. This result does not establish biological independence or certify a global optimum. Section 6 presents the multi-start dependence profile and discusses its numerical limitations.

4.7. Alternative-Family Benchmark

For two causes, the robustness benchmark joins the same Beta-P survivals with the Clayton survival copula [10],

$$S_C(t_1, t_2; \eta) = \left\{S_1(t_1)^{-\eta} + S_2(t_2)^{-\eta} - 1\right\}^{-1/\eta}, \qquad \eta > 0, \tag{46}$$

with independence obtained as $\eta \downarrow 0$. Its observable cause-specific subdensity is

$$g_j^C(t) = f_j(t) S_C(t, t; \eta)^{1+\eta} S_j(t)^{-(1+\eta)}. \tag{47}$$

We include the Clayton survival copula only as a theoretical robustness benchmark. We do not report empirical Clayton likelihood rankings or fitted curves because an independently executable optimization procedure is not available.

5. Computational Validation and Reproducibility

5.1. Monte Carlo Validation Design

We conducted a reproducible Monte Carlo check of the exact Beta-P marginal survival and active-set interval-probability formulas. For each of the three dependence settings ($\delta = 0$, 0.08, and 0.24), we generated 500,000 latent-time pairs using shared and cause-specific unit-rate gamma frailties and conditional exponential thresholds. The common marginal parameters were $\alpha = (0.8, 1.2)$, $\beta = (12, 9)$, and $\theta = (1.2, 0.9)$. We compared empirical marginal survival probabilities at 6, 12, and 24 months, as well as empirical net and crude event probabilities over [0,12) and [12,24), with the exact Beta-P survival probabilities and one-dimensional numerical integrals of the cause-specific subdensities. The random seeds and executable code are provided in Supplementary File S2.

5.2. Monte Carlo Results and Scope

Across 36 comparisons of active sets, causes, intervals, and dependence settings, the largest absolute difference between a Monte Carlo frequency and its corresponding integrated probability was 0.00157, or 2.25 Monte Carlo standard errors. Across 18 marginal-survival comparisons, the largest absolute difference was also 0.00157. These results support the numerical implementation of the marginal and active-risk probability formulas under the specified data-generating model.

This Monte Carlo study checks the probability calculations. It does not assess estimation of δ, finite-sample bias or RMSE, likelihood-ratio test size or power, or confidence-region

coverage. It also does not resolve the weak-identification and optimizer-convergence limitations of the clinical profile analysis.

5.3. Fixed-Margin Monte Carlo Assessment of Dependence Estimation

We conducted an oracle Monte Carlo experiment to assess estimation of δ when the six Beta-P marginal parameters are known. For each scenario, we generated 2000 replicated datasets of $n = 177$ competing-risks observations from the shared-gamma model, using $\alpha = (0.02558137, 0.13958528)$, $\beta = (1.37908308, 1.78020412)$, and $\theta = (6.866957, 1.83601277)$, with independent administrative censoring at 60 months. The true values of δ were 0, $0.5\delta\max$, and $0.9\delta\max$, where $\delta\max = \min(\alpha_1, \alpha_2)$. For each dataset, we maximized the observed-data likelihood over the closed interval $[0, \delta\max]$ using bounded one-dimensional optimization and explicitly compared the two endpoints. The random seed, data generator, settings, replicate-level estimates, and summary code are provided in Supplementary File S2.

When $\delta = 0$, the mean estimate was 0.00584, the bias was 0.00584, and the RMSE was 0.01004; 50.7% of estimates were at the independence boundary. When $\delta = 0.5\delta\max$, the bias was 0.00009, and the RMSE was 0.00970, with 20.5% and 19.3% of estimates at the lower and upper boundaries, respectively. When $\delta = 0.9\delta\max$, the bias was -0.00453, and the RMSE was 0.00975, with 41.8% of estimates at the upper boundary. These results show substantial boundary behavior even when the marginal parameters are known. Because the experiment fixes the marginal parameters at their true values, it does not assess joint estimation of δ and the marginal parameters, confidence-interval coverage, or the performance of the full seven-parameter optimizer.

5.4. Complete-Estimator Validation and Optimization Stability

The unrestricted full likelihood exhibited a nonregular numerical path in pilot samples: alpha approached zero while theta increased toward the numerical upper bound and the likelihood continued to improve. Because increasing the number of starts cannot remove this boundary ridge, we defined the validation target as a compactly constrained complete seven-parameter profile-grid estimator. Positive free kappa components were restricted to $[10^{-4}, 2]$, β components to $[0.10, 120]$, and θ components to $[0.20, 20]$; exact $\kappa_1 = 0$, $\kappa_2 = 0$, and joint-zero faces remained admissible. The generating values lie strictly inside these limits.

For each of the three true dependence values, $\delta = 0$, $0.5\delta\max$, and $0.9\delta\max$, we simulated 100 datasets of $n = 177$ and re-estimated all six nuisance parameters over a 26-point δ grid. At each positive grid value, the interior and the three exact κ-boundary faces were optimized. Every face received two independent batches of 12 starts. If the batch minima differed by more than 10^{-6}, two additional independent batches were run. Four initially discordant replicates were subsequently rerun with 24 starts per batch. Analytic gradients were verified against centered finite differences, and convergence was assessed using the KKT projected-gradient norm.

After the doubled-start reruns, the two best independent batches agreed within 10^{-6} in all 300 replicates. Optimizer success was 100%, the median projected-gradient norms were 7.27×10^{-7}, 1.00×10^{-6}, and 8.69×10^{-7}, and the largest selected projected-gradient norm was 8.46×10^{-5}. For true δ values 0, 0.012791, and 0.023023, the biases were 0.011700, -0.003581, and -0.008533, and the RMSE values were 0.018519, 0.012792, and 0.018948. The independence boundary was selected in 43%, 50%, and 35% of replicates, respectively. The upper θ_1 limit was reached in 19%, 13%, and 25% of replicates.

The constrained optimization is therefore reproducible, but the complete estimator remains weakly identified at the clinical sample size. The boundary-selection rates, theta-limit

contacts, and non-negligible bias and RMSE do not support claims of precise dependence estimation, unconstrained estimator reliability, or calibrated confidence-region coverage. The clinical profile and probability envelopes are consequently retained as exploratory sensitivity analyses.

6. Application to Transplantation Data

6.1. Data and Nonparametric Analysis

The publicly available casebase::bmtcrr dataset contains 177 patients who underwent hematopoietic stem-cell transplantation for acute leukemia [11]. The dataset documentation codes Status = 0 as right-censored, Status = 1 as relapse, and Status = 2 as a competing event. The source vignette gives transplant-related death as an example of a competing event; it does not identify a more specific cause for every observation coded Status = 2. Follow-up is measured in months.

The dataset includes 56 relapses (31.6%), 75 competing events (42.4%), and 46 right-censored observations (26.0%). Median observed follow-up was 6.63 months (range: 0.13–131.77), and the mean age was 30.47 years (SD: 13.04). These descriptive estimates are summarized in Table 2.

Table 2. Aalen–Johansen estimates at selected follow-up times.

Months	Number at Risk	Relapse CIF	Competing-Event CIF	Event-Free Survival
12	56	0.2714	0.4017	0.3269
24	42	0.3020	0.4200	0.2780
36	31	0.3167	0.4266	0.2567
60	20	0.3253	0.4266	0.2481

6.2. Parametric Fit

The independent Beta-P model had a maximized log-likelihood of −561.210720. With six fitted parameters, its AIC was 1134.42, and its BIC was 1153.48 ($n = 177$). We then fitted the shared-gamma model by multi-start numerical optimization, re-estimating the six marginal parameters at each of the 25 fixed dependence values (the Dependence-Parameter Likelihood Profile subsection). The best likelihood found occurs at the independence boundary ($\delta = 0$); the numerical search found no improvement for positive dependence. We independently re-evaluated the likelihood values (Section 7), but these checks do not establish global optimality or calibrated dependence-parameter confidence coverage. We retain the independent-model predictions as the boundary benchmark. Fit statistics and parameter estimates are summarized in Table 3.

Table 3. Independent Beta-P model fit statistics and parameter estimates.

Model	Log-Likelihood	AIC	BIC
Independent Beta-P (descriptive benchmark)	−561.2107	1134.42	1153.48
Cause	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\theta}$
Relapse	0.02558	1.37908	6.86695
Competing event	0.13959	1.78020	1.83601

Computational and writing assistance. ChatGPT (OpenAI) assisted with manuscript drafting and organization, numerical diagnostic scripts, tabulation, and plotting. The author reviewed the reported estimates and diagnostics against the available data, code, and numerical records; this assistance is not independent scientific validation. The included seeded studies validate the probability formulas, assess oracle fixed-margin estimation of delta,

and evaluate the explicitly constrained complete estimator. The complete-estimator study documents numerical reproducibility and finite-sample limitations; it does not establish unconstrained global optimality or confidence-region coverage. The complete-estimator validation was run with Python 3.12.14, NumPy 2.3.5, SciPy 1.17.0, and pandas 2.2.3.

Descriptive Assessment of Observable Event-Free Survival

We compared observed event-free survival with the fitted independent Beta-P model at nine reporting times: 3, 6, 9, 12, 18, 24, 36, 48, and 60 months. At 12 months, the Aalen–Johansen estimate was 0.3269, compared with 0.4176 under the independent model. Because we selected the reporting times after exploratory inspection, these differences are descriptive and are not prespecified goodness-of-fit tests. The comparison is shown in Figure 1.

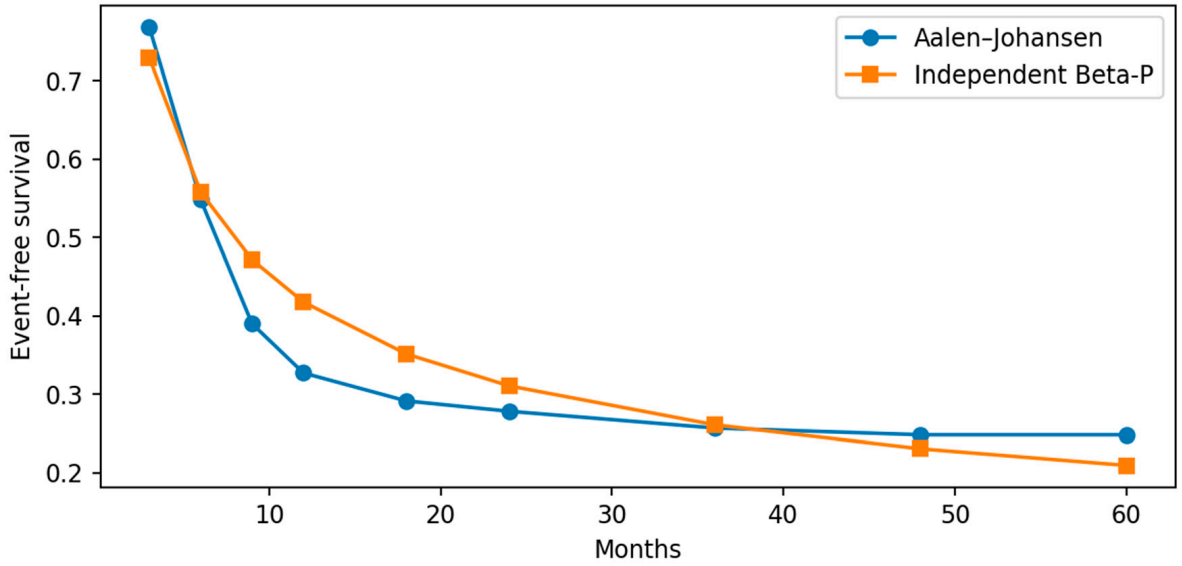


Figure 1. Observed and independent Beta-P event-free survival (descriptive).

6.3. Net and Crude Probabilities

For the fitted independent Beta-P benchmark, the crude probability was lower than the corresponding net probability in every reported interval, consistent with the active-risk ordering. The interval-specific values are reported in Table 4.

Table 4. Net and crude probabilities by cause and interval.

Cause	Interval	Net	Crude ($\delta = 0$)	Crude/Net
Relapse	0–12	0.3162	0.2513	0.7949
Relapse	12–24	0.0784	0.0440	0.5615
Relapse	24–36	0.0416	0.0203	0.4875
Relapse	36–60	0.0484	0.0210	0.4340
Competing event	0–12	0.3893	0.3311	0.8504
Competing event	12–24	0.0979	0.0631	0.6448
Competing event	24–36	0.0503	0.0294	0.5847
Competing event	36–60	0.0566	0.0306	0.5398

Cause-Specific Incidence Diagnostics

At 12 months, the Aalen–Johansen cumulative-incidence estimates were 0.2714 for relapse and 0.4017 for the competing event. The independent Beta-P model predicted 0.2513 and 0.3311, respectively. Thus, the model underestimated first-year competing-event

incidence by approximately 0.0706. These differences are descriptive and should not be interpreted as causal effects. The cumulative-incidence comparison is shown in Figure 2.

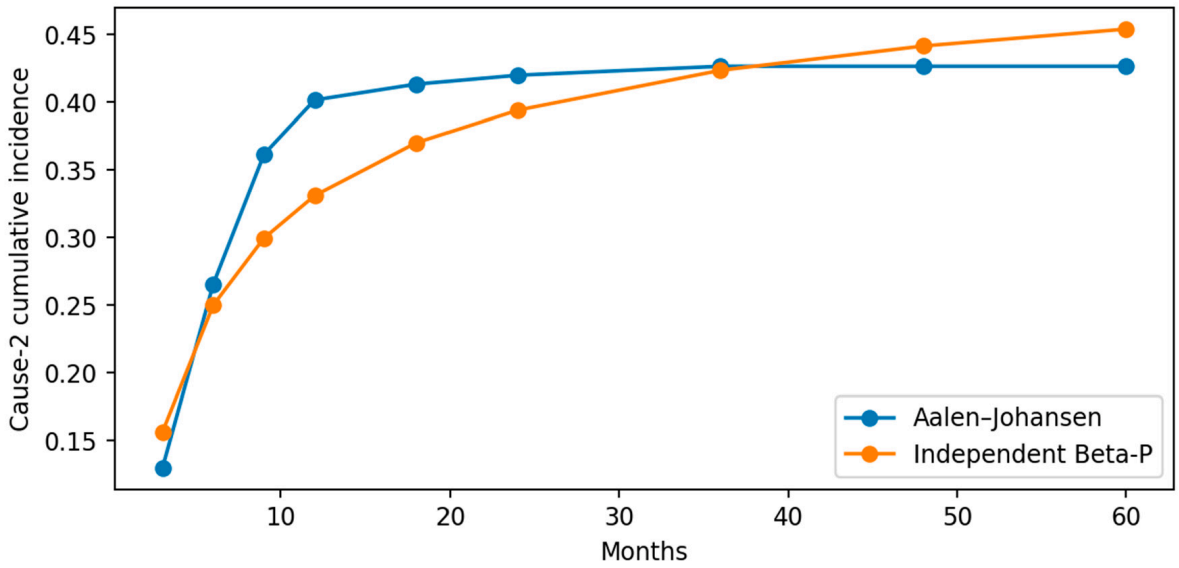


Figure 2. Observed and independent Beta-P cause-2 cumulative incidence (descriptive).

6.4. Scope of Sensitivity Analysis

Section 4 defines structural sensitivity and likelihood-supported probability regions. Section 6.4.1 reports numerical fixed-margin extrema over the full admissible one-dimensional δ interval. Section 6.4.3 reports exploratory probability ranges from the reestimated profile grid, and Section 7 summarizes the independent likelihood audit. The profile-grid ranges cover only evaluated points; they are not confidence intervals, continuous-domain profile extrema, distribution-free identification bounds, or evidence of global likelihood optimality.

6.4.1. Fixed-Margin Numerical Sensitivity Illustration

To illustrate the structural sensitivity measures defined in Section 4.3, we hold the independent Beta-P marginal estimates in Table 3 fixed and vary δ over its admissible interval, $[0, \delta\max]$, where $\delta\max = \min(\alpha_1, \alpha_2) = 0.025581$. At each value of δ, we calculate event-free survival and cause-specific cumulative incidence by one-dimensional numerical integration. We then optimize each probability separately at 12 and 24 months over the full interval. The procedure evaluates both endpoints, screens 2001 equally spaced values, refines any detected interior extrema over neighboring brackets, and repeats bounded scalar optimization over the full interval with an x-tolerance of 10^{-12}. All resulting extrema occur at the endpoints and agree with the dense-grid extrema to the reported precision. They are numerical extrema within this fixed-margin frailty family, not analytically certified bounds over unrestricted dependence models. This analysis neither estimates δ from the data nor demonstrates improved predictive fit.

Table 5 and Figure 3 show five deterministic points on the sensitivity curves. Separate optimization over the full interval gives the following ranges. At 12 months, event-free survival ranges from 0.417605 to 0.457107, relapse incidence from 0.248162 to 0.251321, and competing-event incidence from 0.294731 to 0.331074. At 24 months, the corresponding ranges are 0.310474–0.350897, 0.295339–0.297081, and 0.352021–0.394187. All six minima and maxima occur at $\delta = 0$ or $\delta = \delta\max$. These are numerical extrema for the fixed margins, not universal sharp bounds. Small differences from the displayed grid values reflect parameter rounding and integration precision.

Table 5. Fixed-margin structural sensitivity of event-free survival and cumulative incidence at 12 months (illustrative; not a fitted dependence estimate).

δ/δmax	δ	S(12,12)	Relapse CIF(12)	Competing CIF(12)
0.00	0.000000	0.417607	0.251308	0.331085
0.25	0.006395	0.427150	0.250494	0.322356
0.50	0.012790	0.436910	0.249696	0.313394
0.75	0.019185	0.446894	0.248914	0.304192
1.00	0.025580	0.457106	0.248148	0.294745

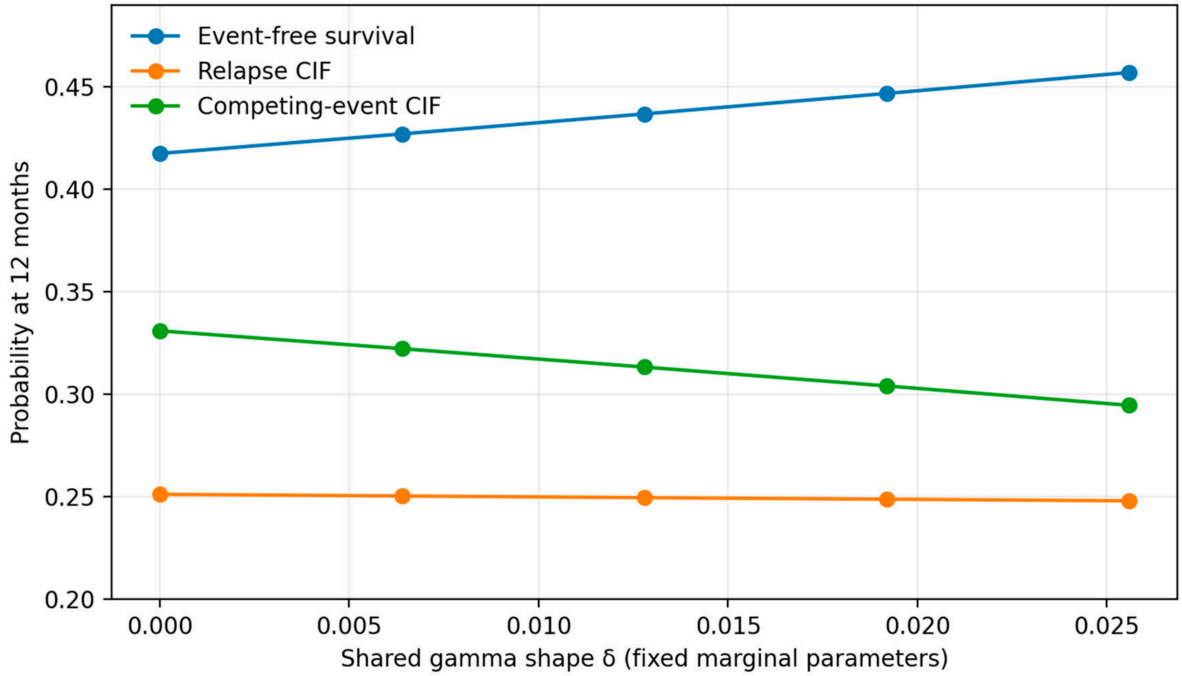


Figure 3. Dependence sensitivity at 12 months with fixed Beta-P marginal parameters. Lines connect five deterministic evaluations, not estimates with uncertainty bands.

6.4.2. Clinical Benchmark Comparison and Interpretation

Table 6 compares the 12-month Aalen–Johansen estimates with the independent Beta-P benchmark ($\delta = 0$) and the illustrative fixed-margin endpoint ($\delta = 0.02558$). The observed estimates describe the data nonparametrically. Both model columns are conditional predictions based on the same fitted marginal parameters; the positive-dependence endpoint is not estimated from the data.

Table 6. Twelve-month observable probability comparison (signed differences are model minus Aalen–Johansen).

Quantity	Aalen–Johansen	Independent δ = 0	Sensitivity δ = 0.02558	Model − AJ (δ = 0 → Endpoint)
Event-free survival	0.326930	0.417607	0.457106	+0.090677 → +0.130176
Relapse CIF	0.271375	0.251308	0.248148	−0.020067 → −0.023227
Competing-event CIF	0.401695	0.331085	0.294745	−0.070610 → −0.106950

The independent benchmark exceeds the observed event-free survival by 0.090677 and underestimates relapse and competing-event incidence by 0.020067 and 0.070610, respectively. At the illustrative positive-dependence endpoint, the signed differences are +0.130176, −0.023227, and −0.106950. Thus, with the margins held fixed, this dependence perturbation increases the discrepancy from the Aalen–Johansen estimate for all three

observable probabilities at 12 months. This comparison does not constitute a formal model-comparison test.

In both the nonparametric and model calculations, event-free survival plus the two cumulative-incidence functions equals one, up to rounding. The three signed differences are therefore algebraically linked and are not independent measures of fit. Because the sensitivity calculations retain the parameters fitted under independence, they are not predictions from a refitted dependent model. The five-point display is illustrative; the fixed-margin extrema are calculated separately as described in Section 6.4.1. Neither analysis provides uncertainty coverage or sharp bounds over unrestricted dependence models.

Dependence-Parameter Likelihood Profile

In the baseline numerical profile, we re-estimated the six marginal parameters at each of 25 fixed values of δ. The best log-likelihood found was -561.210720 at $\delta = 0$. Expanded profile checks separately optimized the exact $\kappa_1 = 0$, $\kappa_2 = 0$, and joint-zero faces at every positive grid value; the best likelihood remained at $\delta = 0$. The largest absolute numerical-gradient component among the selected solutions was 1.14×10^{-4}. Forward and reverse continuation agreed within 10^{-4} at 23 of the 25 grid points; at two points, the reverse path reached inferior local solutions. Across-start objective gaps remained large, so these diagnostics do not establish a global optimum. Because $\delta = 0$ is a boundary point, ordinary interior-parameter likelihood-ratio cutoffs do not automatically provide calibrated confidence regions. The finite grid and optimizer instability also preclude claims of continuous-domain or global optimization.

6.4.3. Reestimated-Gamma Clinical Comparison at 12 and 24 Months

Table 7 compares the Aalen–Johansen estimates with the independent Beta-P boundary fit ($\delta = 0$) and the ranges of fitted probabilities across the 21 profile-grid points that pass an exploratory likelihood-ratio screen ($2\Delta\ell \leq 3.84146$). We re-estimated all six marginal parameters at each fixed δ. These ranges are descriptive envelopes over the evaluated solutions; they are not confidence intervals or sharp identified bounds.

Table 7. Reestimated-parameter probability ranges at 12 and 24 months.

Month/Outcome	Aalen–Johansen	Independence Fit	Profile-Grid Range	Ind. − AJ
12/Event-free survival	0.3269	0.4176	0.4129–0.4177	+0.0907
12/Relapse CIF	0.2714	0.2513	0.2513–0.2543	−0.0201
12/Competing-event CIF	0.4017	0.3311	0.3309–0.3350	−0.0706
24/Event-free survival	0.2780	0.3105	0.2990–0.3105	+0.0324
24/Relapse CIF	0.3020	0.2953	0.2952–0.3046	−0.0066
24/Competing-event CIF	0.4200	0.3942	0.3942–0.3963	−0.0258

The shared-gamma model has its best recorded likelihood at the independence boundary. Accordingly, the dependent model does not improve the maximized likelihood over the independent model. Because the dependent model includes δ as an additional parameter, its AIC and BIC are larger at this boundary fit. This result applies to the numerical solutions examined; it does not establish global optimality or true biological independence.

At 12 months, Aalen–Johansen event-free survival is 0.3269, compared with 0.4176 under the independent model; the exploratory reestimated-profile range is 0.4129–0.4177. Thus, the profile considered here does not eliminate the observed survival discrepancy. Table 7 reports the corresponding relapse and competing-event incidence comparisons and the 24-month results. We present the fitted model as a framework for sensitivity and partial-identification analysis, not as a clinically superior predictor.

The fixed-margin analysis in Section 6.4.1 holds the independent-model margins constant while varying δ. By contrast, Table 7 reports ranges from profiles that re-estimate the margins at each grid value. These analyses answer different questions, and their ranges should not be conflated. We assign no nominal confidence level to the exploratory cut-off because δ is a boundary parameter and the likelihood-ratio calibration has not been established.

6.5. Numerical Verification

The exact-margin identity follows analytically from the gamma Laplace transform (Proposition 1). The numerical comparisons in this section are descriptive. We do not report a dense-grid discrepancy statistic because none has been independently verified. Additional reproducibility details are provided in Appendix A.

7. Discussion

The model provides an exact separation between marginal Beta-P behavior and a positive static frailty dependence mechanism. This separation makes structural sensitivity interpretable: changes in probability across δ are not produced by changing the fixed margins.

The active-risk formulation brings together probability concepts that are often treated separately and clarifies the difference between observable crude incidence and hypothetical net or partial-crude probabilities. Adding causes to nested active sets yields an ordering result, but the dependence parameter does not determine a universal direction of change for each cause-specific interval probability.

The application distinguishes structural sensitivity from empirical fit. Observing only the minimum failure time and its cause does not identify unrestricted latent dependence. Section 6.4.1 varies δ while holding the independent-model margins fixed; Section 6.4.2 shows that the positive-dependence endpoint increases the descriptive discrepancies from the Aalen–Johansen estimates for all three observable probabilities at 12 months. the Dependence-Parameter Likelihood Profile subsection reports the multi-start dependence profile, and Section 6.4.3 presents exploratory probability ranges based on reestimated parameters. These analyses establish neither improved predictive fit nor global optimality, sharp identification bounds, or calibrated confidence coverage.

The independent Beta-P model is a boundary submodel of the correlated-gamma construction. The Aalen–Johansen estimator provides a nonparametric benchmark for observable cumulative incidence. We discuss the Clayton copula as a theoretical alternative dependence family, but do not use an unverified numerical Clayton fit to rank the models.

The independent-model curve does not fully reproduce the observed first-year competing-event incidence. This descriptive discrepancy limits the clinical predictive claims that can be made, but it does not invalidate the analytic probability results.

Expanded optimization audit. At each positive δ grid point, we optimized the interior and the exact faces $\kappa_1 = 0$, $\kappa_2 = 0$, and $\kappa_1 = \kappa_2 = 0$. Each candidate face used 12 L-BFGS-B starts in the forward continuation and 12 in the reverse continuation, for 24 starts per face across both runs. The independent model at $\delta = 0$ was evaluated separately. We used maxiter = 2500, ftol = 10^{-13}, gtol = 10^{-7}, maxls = 60, and maxcor = 30; transformed free parameters were bounded by $[-15, 8]$. Forward and reverse continuation gave the same best objective to within 10^{-4} at 23 of 25 grid points. At $\delta = 0.008733$ and 0.010667, reverse continuation reached inferior local solutions; the expanded multi-start search recovered lower objectives. The best log-likelihood found was -561.21071984 at $\delta = 0$, and the largest absolute component of the numerical gradient at the selected solution was 1.14×10^{-4}. Across-start objective gaps nevertheless remained large, and the $\kappa_1 = 0$ face was selected

at several positive grid values. These checks improve reproducibility but do not establish global optimality or reliable full-estimator performance.

The contribution is not the generic use of gamma frailty or its additive decomposition, both of which are established [6–8,12,13]. Rather, it is the exact Beta-P marginal specialization, the resulting cause-specific subdensities and active-risk functionals, the nested-set ordering result, and the model-specific sensitivity summaries. The profile-grid ranges are descriptive; they do not establish continuous-domain extrema or calibrated profile regions. Unlike dynamic frailty or longitudinal joint models, this model samples the dependence variables once per individual and has an explicit observed-data likelihood.

8. Limitations and Future Work

The construction permits positive dependence but not arbitrary negative association. Dependence is model-based because the complete latent failure-time vector is not observed. A single common parameter imposes a parsimonious association structure and may be restrictive for three or more causes. Omitted covariates may be absorbed partly by frailty, so δ must not be interpreted as residual biological dependence after complete adjustment.

The fixed-margin sensitivity analysis holds the marginal parameters constant and does not estimate δ by likelihood. Its numerical extrema apply only to the proposed family; they are not universal Fréchet-type bounds. Net and partial-crude probabilities are model-based cause-elimination functionals. A causal interpretation requires additional assumptions about how removing one cause affects the remaining hazards. Misspecification of the marginal distributions may distort both dependence estimates and sensitivity regions.

The comparisons of absolute fit are descriptive, not prespecified confirmatory tests. The reporting times were selected after exploratory inspection of the empirical curves. Profile-grid ranges are not bootstrap confidence bands, and optimizer convergence flags do not establish global optimality. We report neither a fitted Clayton comparison nor a validated full-model operating-characteristic study. The fixed-margin Monte Carlo experiment assesses δ estimation only when the marginal parameters are known; it is not a goodness-of-fit test for the clinical data. Comparisons at 12 and 24 months do not establish overall calibration or predictive superiority.

Future work may introduce regression effects in β_j or other marginal parameters, pair-specific shared components for multiple causes, robustness comparisons across several dependence families, and nonregular asymptotic theory for sensitivity extrema and the LSI. These extensions should be developed separately to preserve the transparency of the present model.

9. Conclusions

We developed a dependent competing-risks model with exact Beta-P margins using an additive shared gamma frailty. The model yields a closed-form joint survival function, explicit cause-specific subdensities, and a unified formulation of net, crude, and partial-crude probabilities. It also establishes an ordering result when causes are added to nested active sets. The mixed-sign derivative decomposition does not determine whether cause-specific interval probabilities are monotone.

We distinguish model-specific structural sensitivity from empirically identified latent dependence. The LSI is defined for a selected cause, interval, and active-risk set. We report numerical fixed-margin extrema with endpoint and dense-grid checks; ranges from the reestimated-parameter profile remain grid-based and exploratory. Neither analysis provides sharp distribution-free bounds or calibrated confidence regions.

The shared-gamma dependence mechanism and the general nonidentifiability of competing-risks dependence are established concepts. Our contribution is the exact Beta-P

marginal construction, its active-risk probability formulas, and the resulting model-specific sensitivity analysis. The clinical application documents a numerical optimum at the independence boundary and presents model-based probability ranges; it does not show that positive dependence improves prediction.

In the transplantation example, the independent Beta-P benchmark underestimates first-year competing-event incidence and overestimates event-free survival. The best shared-gamma likelihood found occurs at the independence boundary, and the reestimated positive-dependence profile does not establish improved fit. With the margins held fixed, the positive-dependence perturbation increases the discrepancy from the Aalen–Johansen estimates at 12 months. The contribution is a transparent exact-margin sensitivity framework, not demonstrated clinical predictive superiority.

Subject to its parametric assumptions, the framework separates marginal failure-time behavior, observable crude incidence, and sensitivity to latent dependence.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/appliedmath6110196/s1, Supplementary File S1 contains the public bmtcrr data and descriptive clinical calculations. Supplementary File S2 contains Python scripts, data, instructions, and machine-readable results for the independent Beta-P fit, multi-start shared-gamma fit, dependence profile, fixed-margin sensitivity calculations, reestimated-parameter grid ranges, clinical comparisons, and independent numerical audit. The code documents the numerical solutions examined; it does not certify global optimization, calibrated likelihood-ratio coverage at the boundary, or the unverified simulation and Clayton results.

Funding: This research received no external funding.

Institutional Review Board Statement: This methodological study uses a previously published, publicly available secondary dataset. It involved no recruitment or intervention and no access to directly identifying patient records. No new ethical approval was sought for this secondary analysis.

Informed Consent Statement: Not applicable. This secondary analysis used an existing, publicly available, de-identified dataset; no participants were enrolled or contacted.

Data Availability Statement: The bmtcrr dataset is publicly available in the R package casebase (https://sahirbhatnagar.com/casebase/reference/bmtcrr.html (accessed on 10 November 2026)). Supplementary File S1 contains the data used for this secondary analysis and selected derived results. The dataset documentation page was accessed on 11 November 2026.

Acknowledgments: The author acknowledges using ChatGPT (OpenAI; the specific model version was not recorded) during manuscript preparation for English-language editing, manuscript organization, mathematical and computational cross-checks, illustrations, and presentation of numerical results. Where feasible, the author checked AI-assisted computational outputs against available numerical records and supporting calculations. Exploratory procedures that could not be fully reproduced independently were excluded from reported inferential findings, and their limitations are disclosed. The author reviewed and revised the AI-assisted material and accepts full responsibility for the scientific content, methods, analyses, interpretations, and conclusions. ChatGPT is not an author.

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

Appendix A. Reproducibility Notes

Reproducibility note. Supplementary File S2 includes archived public data, source scripts and outputs for independent and dependent likelihood fitting, the 25-point clinical dependence profile, descriptive comparisons, fixed-margin sensitivity calculations, likelihood re-evaluation, seeded formula-validation Monte Carlo, the oracle fixed-margin assessment, and the complete-estimator validation. The complete-estimator folder contains the analytic-gradient implementation, parameter limits, seeds, optimizer settings, likelihood-ridge diagnostic, 300 replicate-level records, doubled-start retries, and scenario

summaries. The numerical clinical profile remains grid-based and is not globally certified; profile-grid envelopes are not confidence regions.

References

1. Tsiatis, A. A nonidentifiability aspect of the problem of competing risks. *Proc. Natl. Acad. Sci. USA* **1975**, *72*, 20–22. [CrossRef] [PubMed]
2. Prentice, R.L.; Kalbfleisch, J.D.; Peterson, A.V.; Flournoy, N.; Farewell, V.T.; Breslow, N.E. The analysis of failure times in the presence of competing risks. *Biometrics* **1978**, *34*, 541–554. [CrossRef] [PubMed]
3. Tian, Y.; Gui, W. Statistical inference of dependent competing risks from Marshall–Olkin bivariate Burr-XII distribution under complex censoring. *Commun. Stat. Simul. Comput.* **2024**, *53*, 2988–3012. [CrossRef]
4. Sarhan, A.M.; Tolba, A.H.; Ramadan, D.A.; Manshi, T. A dependent bivariate Burr XII inverse Weibull model: Application to diabetic retinopathy and dependent competing risks data. *Mathematics* **2026**, *14*, 120. [CrossRef]
5. Gorfine, M.; Hsu, L. Frailty-based competing risks model for multivariate survival data. *Biometrics* **2011**, *67*, 415–426. [CrossRef] [PubMed]
6. Eriksson, F.; Scheike, T.H. Additive gamma frailty models with applications to competing risks in related individuals. *Biometrics* **2015**, *71*, 677–686. [CrossRef] [PubMed]
7. Martins, A.; Aerts, M.; Hens, N.; Wienke, A.; Abrams, S. Correlated gamma frailty models for bivariate survival time data. *Stat. Methods Med. Res.* **2019**, *28*, 3437–3450. [CrossRef] [PubMed]
8. Gorfine, M.; Zucker, D.M. Shared frailty methods for complex survival data: A review of recent advances. *Annu. Rev. Stat. Its Appl.* **2023**, *10*, 51–73. [CrossRef]
9. 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]
10. Clayton, D.G. A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometrika* **1978**, *65*, 141–151. [CrossRef]
11. Hougaard, P. *Analysis of Multivariate Survival Data*; Springer: New York, NY, USA, 2000. [CrossRef]
12. Duchateau, L.; Janssen, P. *The Frailty Model*; Springer: New York, NY, USA, 2008. [CrossRef]
13. Scrucca, L.; Santucci, A.; Aversa, F. Competing risk analysis using R: An easy guide for clinicians. *Bone Marrow Transplant.* **2007**, *40*, 381–387. [CrossRef] [PubMed]

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.