

Cubic Logistic Model Results Analysis

We fitted the extended monotone cubic logistic model for the conditional probability

$$\text{Prob}(AE \mid X = x) \quad (1)$$

in both the FULL and TRIM datasets, using raw- and log-scale biomarker representations. The resulting fitted curves are shown in Figure 1. The figure displays fitted probabilities of the adverse event as a function of the predictor x for the FULL and TRIM cohorts, shown on both the log-transformed and raw scales. In both cohorts, the probability increases from near zero to near one, but the shapes differ substantially. The TRIM cohort exhibits a steep, almost step-like transition over a narrow range of x , consistent with a well-defined threshold. In contrast, the FULL cohort shows a more gradual increase with a clear shoulder in the mid-range of x , where the curve temporarily flattens before rising again. This feature is particularly visible on the log scale. Notably, the location of this flattening aligns with isolated observations in the NC subset, suggesting that these points influence the fitted curve and contribute to the observed deviation from a simple sigmoidal shape.

From a mathematical perspective, these features arise from the cubic specification in the linear predictor, $\beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 x^3$. The derivative of this expression is quadratic, allowing changes in curvature and thus permitting an inflection (saddle-like) point in the resulting probability curve after applying the logistic (expit) transformation. This flexibility enables the model to locally reduce its slope before increasing again, producing the observed shoulder. The position of this saddle point is data-driven and, in this case, coincides with the region where outlying NC observations occur. These points exert leverage on the likelihood, effectively pulling the fit to accommodate them and inducing the change in curvature. Such behavior cannot be captured by lower-order models, highlighting how the cubic ansatz introduces sufficient flexibility to reflect both genuine structure and sensitivity to influential observations.




Figure 1: Maximum likelihood fits of the cubic logistic regression model for the conditional probability of an adverse event, $P(AE \mid X = x)$, for the FULL and TRIM cohorts. Panel (A) shows the fitted curves on the raw scale of X , while panel (B) shows the fits on the log-transformed scale. Observed data for both cohorts are overlaid. The legend applies to both panels.

To further assess parameter identifiability and local optimization behavior, profiles of the negative log-likelihood around the optimum are shown in Figure 2. The richest local structure is observed for the TRIM model fitted on the log-transformed biomarker scale. The profiles reveal that the curvature of the cost function varies substantially across parameters: some directions ex-

hibit very sharp, well-defined minima, while others are comparatively flat or display irregular local features.

From a statistical perspective, this pronounced anisotropy in curvature implies large differences in the precision of parameter estimates, with steep directions corresponding to well-identified parameters (small standard errors) and flat directions indicating weak identifiability (large standard errors). From a computational standpoint, such heterogeneity in curvature leads to an ill-conditioned Hessian matrix with a large condition number. This, in turn, results in slow convergence of optimization algorithms, as progress along flat directions is limited while steep directions require small step sizes for stability.

Moreover, the presence of nearby local minima and irregularities in some profiles suggests the existence of competing solutions in the parameter space. This further complicates inference, as optimization routines may become sensitive to initialization and may struggle to reliably locate the global optimum. Together, these features highlight both statistical and numerical challenges associated with fitting the cubic logistic regression model in this setting.




Figure 2: Profiles of the negative log-likelihood (cost function) around the maximum likelihood estimate for the cubic logistic regression model. Each curve shows the variation in the cost function when perturbing one parameter at a time, as indicated in the legend. Results are shown for the FULL cohort (panel A) and the TRIM cohort (panel B) using log-transformed biomarkers. The vertical axis is displayed on a logarithmic scale. The legend applies to both panels.

Figure 3 compares 95% confidence bands for the fitted cubic logistic regression model obtained using several approaches, which can be grouped into two main categories. The first group relies on asymptotic normality of the parameter estimates around the optimum. In this approach, the parameter vector is assumed to follow a multivariate normal distribution with covariance matrix given by the inverse Hessian of the negative log-likelihood, requiring no additional optimization beyond the original fit. Two variants are shown: the *normal* approach, where confidence bands are constructed from empirical quantiles of model predictions obtained by sampling parameters from this distribution, and the *delta* method, where the nonlinear model is linearized around the optimum, uncertainty is propagated analytically, and the resulting bands are clipped to the admissible range $[0, 1]$.

The second group consists of bootstrap-based methods [1], which explicitly account for sampling variability by generating bootstrap samples and refitting the model. In the *nonparametric bootstrap*, samples $\{(x_i^*, y_i^*)\}_{i=1}^n$ are obtained by resampling the original observations with replacement. In the *parametric bootstrap*, responses y_i^* are generated at the observed x_i by sampling from a Bernoulli distribution with success probability given by the fitted model $p(x_i)$. Confidence bands are then constructed from the empirical quantiles of the resulting fitted curves.

While asymptotic methods are typically expected to produce narrower and smoother confidence bands, reflecting the local curvature of the likelihood captured by the Hessian, this behavior is not observed in the present setting. The estimated parameter uncertainties are highly heterogeneous and, in some directions, extremely large due to weak identifiability and poor conditioning of the Hessian. As a consequence, the *normal* approach yields degenerate or uninformative upper confidence limits, effectively saturating at the boundary value 1 over a wide range of x . In contrast, the bootstrap-based bands are more stable and better reflect the actual variability of the fitted model, although they may appear irregular, particularly in regions with sparse data or strong nonlinearity. These differences are most pronounced in the FULL cohort, where the fitted curve exhibits a shoulder (inflection) and the likelihood surface is less well-conditioned, leading to substantial discrepancies between methods. In the TRIM cohort, where the transition is sharper and the model is better identified, the various confidence bands are more consistent.




Figure 3: Fitted cubic logistic regression model with 95% confidence bands constructed using different methods (see legend). Results are shown for the FULL cohort (panels A–B) and the TRIM cohort (panels C–D). Panels A and C display the model on the log scale of the biomarker (predictor), whereas panels B and D display it on the raw scale. Observed data points are overlaid. The legend applies to all panels.

References

- [1] Bradley Efron. Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, 7 (1):1–26, 1979. doi: 10.1214/aos/1176344552.