Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# A Comparison of Frequentist RCS, Bayesian RCS, and SHAP for Estimating Risk Inflection Points and Crossing Points of Biomarkers

**Yue-peng Liu**

liuyp2080@163.com

Saideng Biotech Co., Ltd.

**Method Article**

**Keywords:** restricted cubic splines, SHAP, dose-response, inflection point, crossing point, Bayesian inference, Monte Carlo simulation, biomarkers

**Posted Date:** November 1st, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-11522242/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** The authors declare no competing interests.

## Abstract

# Objective

Continuous biomarkers often have a non-linear dose-response with the outcome so the need is to identify two key points on the fitted curve. The inflection point is the point at which the risk curve changes direction and the crossing point is the point at which the effect changes from protective to harmful. This is a very common requirement in clinical and epidemiological work but there is still no validated and practical method for estimating these points.

# Methods

Monte Carlo simulation was used to compare frequentist RCS, Bayesian RCS and SHAP in terms of their ability to estimate inflection and crossing points. A range of dose response shapes, levels of collinearity and outcome types were considered and the accuracy and robustness of the three methods were assessed.

# Results

Bayesian MCMC is best for crossing points. For the nonlinear monotone and symmetric U-shaped cases the bias is around ± 0.03 (a lot better than the − 0.23 to + 0.39 seen with the frequentist RCS). If the crossing point is very far from the exposure mean or on a sparse branch however the prior shrinkage in the Bayesian approach can make it less accurate in which case the frequentist method is superior. For inflection points the frequentist RCS is better (bias of -0.06 versus − 0.32 for the Bayesian in the asymmetric U case). INLA is not reliable for inflection points or the right root due to coefficient distortion so Bayesian estimation should use MCMC. SHAP has a built in bias of around − 0.18 to -0.24 which doesn't go away with tuning so it is only of use for exploratory screening. RCS is more precise when the data do a good job of covering the shape but no method can get a crossing point on a sparse branch. The approach is illustrated in two real examples: serum phosphate in CRRT patients and age in ectopic pregnancy patients.

# Conclusion

use Bayesian MCMC for crossing points and frequentist RCS for inflection points. If the crossing point is far from the exposure mean or on a sparse branch then use frequentist RCS. SHAP should be used purely for exploratory screening.

## 1 Introduction

Dose-response analysis is a core tool in both clinical and epidemiological research. A common approach is to use restricted cubic splines (RCS) or machine learning techniques (e.g. SHAP) to model the non-linear relationship between a continuous marker and an outcome, with the aim of then identifying clinically useful cut points. There has been a lot of methodological work on the modelling of non-linear relationships with a number of practical guidelines available for the handling of continuous variables [1] and systematic comparisons have been made of different basis functions (e.g. splines vs fractional polynomials) [2]. The bulk of this work however is concerned with the fitting of the curve itself and provides very little in the way of systematic methodological guidance on the subsequent question of how to reliably identify inflection points or crossing points from the fitted curve. The term 'cut point' is also used in the literature in a very inconsistent manner and can refer to at least two different things. The first is the inflection point (also known as the turning point) which is the point at which the slope of the risk curve changes direction (e.g. the bottom of a U or the top of an inverted U). The second is the crossing point which is the point at which the effect of the marker changes sign (from negative to positive). With a U-shaped relationship there are in fact two crossing points rather than one. The two points can coincide in the case of a monotonic relationship but with a U-shaped relationship they will be distinct and can be more than one standard deviation apart. Mixing the two up will therefore lead to systematic bias.

Whilst identifying inflection points and crossing points is a very common requirement, there is no systematically validated and directly usable method for estimating them. People will typically just go with their experience in choosing their tools, using the default RCS knots, looking at the SHAP dependence plots by eye or fitting a linear term. The reliability of these approaches has not been assessed and the extent of the bias in terms of both magnitude and source is not known. This is well recognised in the applied literature. Researchers will routinely report 'inflection points', 'thresholds' or 'cut points' for biomarkers on the basis of RCS or two segment linear models (e.g. the curvilinear relationship between C-reactive protein and death from COVID-19 and its cut point [3], the nonlinear relationships between a number of physiological markers and all cause mortality and the cut points derived from those [4], the J shaped relationship between mercury exposure and cardiovascular outcomes and its inflection point [5]). These papers will typically report a single key point. They rarely state whether the point they are referring to is an inflection point or a crossing point and they don't go on to assess estimation method bias. The tools they use (RCS, segmented regression, visual inspection etc) are also inconsistent.

As for the methods used to estimate the points in question, the various approaches currently in the literature for identifying inflection or crossing points fall into three broad categories, each with its own drawbacks. The first category comprises parametric regression methods, most notably restricted cubic splines (RCS) [6][7] and fractional polynomials [8][9]. These methods identify the key points by finding the extrema or zero crossings of the fitted curve. The problem with this approach is that the location of the key point is very much dependent on the choice of basis function and knots and they generally don't provide any quantification of the uncertainty in the location of the key point. The second category are the segmented or change-point models (e.g. two segment linear regression, segmented splines [10][11])

which estimate the inflection or break point directly as a model parameter. The problem with this approach is that the break point is estimated very slowly (theoretical rate of convergence is $n^{-1/3}$ [12]) and there is a systematic bias introduced by approximating a smooth curve with piecewise linear segments. The third category are the machine learning interpretation methods, most notably SHAP [13] (which is generally based on tree models such as XGBoost [14]). These methods locate the risk crossing or inflection point by looking at the SHAP dependence plot. SHAP is designed to estimate the marginal contribution of each variable to the model prediction, not to estimate inflection or crossing points per se. It has no analytic definition of a key point and no uncertainty measure so using it to locate key points is a case of borrowing without any examination of the reliability of that borrowing. These three categories are from fundamentally different modelling paradigms (statistical modelling and machine learning) and the tension between them in terms of their aims and their interpretability has been well documented [15]. Bayesian methods such as INLA [16] and MCMC can provide posterior distributions of the key points but the performance of these methods under conditions of collinearity and small sample sizes is not known. In summary, there is no current "gold standard" approach for the reliable estimation of inflection and crossing points, either because the definition of the point is not consistent, because there is no quantification of the uncertainty in the point's location or because the methods were not designed for that purpose. This paper addresses that gap.

To fill this gap the paper employs systematic Monte Carlo simulation to compare the performance of frequentist RCS, Bayesian RCS and SHAP in the estimation of inflection and crossing points and to quantify the estimation error and conditions of applicability of each method. The paper is structured as follows: it first evaluates the three methods, then identifies the conditions under which RCS estimation is valid and finally concludes that Bayesian MCMC is the preferred method for crossing points, frequentist RCS for inflection points and SHAP for screening.

# 2 Methods

# 2.1 Formal definitions of key points

Let f(x) be the log hazard or conditional effect function linking the marker x to the outcome, and let $\bar{f}$ be the population mean of f. The inflection point is the value of x at which f'(x) changes sign (i.e. an extremum of f). The crossing point is the value of x at which the centred effect $f(x) - \bar{f}$ crosses zero (i.e. where the effect changes sign). The number of crossing points will depend on the shape of the relationship. There will be one under a monotonic relationship and two under a U-shaped relationship (the left and right roots) which will need to be estimated as a pair.

# 2.2 Method exploration and determination

We didn't settle on the methods at the outset. We went through a process of trying various approaches, comparing them and working through the trade-offs. We've described that process so that the reader can see how we arrived at the final approach and can assess its robustness.

**Estimation of key points.** We estimate both inflection and crossing points using RCS. We begin by subjecting the RCS to a non-linearity test against a linear term to establish whether there is genuine non-linearity. The inflection point is taken as the extremum of the fitted RCS curve and the crossing point as the zero crossing of the centred curve (we check that each of these lies within the observed range). RCS is robust across all shapes and does not break down.

**Note the two crossing points in particular.** If the curve is U-shaped then the centred effect will have two zero crossings (a left one and a right one) and the crossing points have to be reported as a pair (the left root and the right root) each with its own confidence interval. Reporting a single crossing point would effectively discard the risk boundary on one side. Unlike a quadratic basis function whose roots can be found in closed form the RCS is a piecewise spline and the crossing points are found numerically by identifying the zero crossings of the centred curve. As there is no closed form solution the left and right roots have to be found separately and each one checked against the observed range. A root on the side with sparse data will also require an additional check of data coverage (as described in Section 3.2.3).

**Choosing an estimation paradigm.** We compare frequentist maximum likelihood estimation (MLE) with Bayesian estimation under the RCS basis functions. Table 1 details the implementation of each of the candidate methods, including the software packages, functions and key parameters.

Table 1
Specific implementations of the candidate estimation methods

| Method | Software and functions | Model settings and key parameters | Solving for the inflection and crossing points |
|---|---|---|---|
| Frequentist RCS, default | R rms::cph for Cox outcomes, rms::Glm for logistic outcomes, riskRegression::FGR for competing risks | rcs(X, 4), that is, 4 knots, plus covariate Z | Inflection point as the extremum of the fitted curve; crossing point by numerical zero crossing after centering |
| Bayesian INLA | R INLA::inla, family = "coxph" or "binomial" | rcs basis functions; fixed-effect prior control.fixed = list(mean = 0, prec = 0.01), that is, N(0, 10) | Posterior samples from inla. posterior. Sample rebuild the curve; inflection point as the mean of the extrema across sample curves, crossing point as the mean of the centered zero crossings. |
| Bayesian MCMC, supplement | R brms::brm with cmdstanr backend, family = cox() or bernoulli() | rcs basis functions; prior prior(normal(0, 10), class = "b"); chains = 2, iter = 400, warmup = 200 | Rebuild the curve from posterior samples to find inflection and crossing points; diagnose convergence with Rhat and ESS |
| SHAP | Run xgboost::xgb.train with objective "survival: cox" or "binary: logistic", plus shapviz::shapviz | eta = 0.1, max_depth = 4 for U shapes or 3, nrounds = 50–80; SHAP values from the S matrix of shapviz | Smooth the SHAP dependence-plot trend with loess (span = 0.6); crossing point as the trend zero crossing, inflection point as the trend minimum |

SHAP is designed to estimate the marginal contribution of each variable to the model prediction (i.e. its individual contribution to the final prediction) rather than specifically to estimate inflection or crossing points. It is included in this paper as a comparison because in practice it is very widely used for the visual inspection of risk crossings and inflection points. This paper uses simulation to look at the reliability of that borrowed use and to assess whether SHAP can in fact be used as a key point localisation technique.

Both the Bayesian implementations use the same prior for all of the regression coefficients (a normal with mean 0 and standard deviation 10 i.e. $N(0,10^2)$) which is a weakly informative prior. INLA specifies this prior in an equivalent way via prec = 0.01. The Bayesian prior shrinkage discussed later is a result of the shrinkage effect of this prior when the data information is low.

**Implementation pitfalls.** The frequentist rms::Predict(ref.zero=TRUE) forces the reference point to the median so the 'zero crossing' is always going to be the median and therefore trivial. If you want to estimate a zero crossing you have to centre it yourself (subtract the mean effect) and then find the zero crossing which is of course the sum-to-zero convention used by SHAP.

# 2.3 Simulation design

**Data generation**: Cox proportional hazards [17], logistic and competing risk models are used with log hazard = f(X) + $\beta_z$ Z where X ~ N(1) is the exposure and Z = ρX + √sqrt (1-ρ²) ε is the covariate (ε ~ N(0)) with ρ controlling the degree of collinearity. **Dose response shapes**: linear monotone f = 0.6X, nonlinear monotone f = 0.6X − 0.15X², symmetric U f = 0.3X + 0.25X², asymmetric U f = -0.3X + 0.25X² + 0.1X³. Further non-quadratic shapes are used to test RCS robustness (e.g. sigmoid monotone f = $3/(1 + e^{-2(x-0.3)})$ and plateau f = $1.6/(1 + e^{-4(x+0.5))})$. **Collinearity levels**: ρ = 0, 0.3, 0.5, 0.8. **Outcome types**: Cox survival, binary logistic, competing risk. **Monte Carlo**: most scenarios use n = 200 replications, some use n = 100. **Evaluation**: let θ be the true value, $\theta_{i\ be}$ the estimate from the i-th simulation and N be the number of replications. Bias is (1/N) $\Sigma_i$ ($\theta_i$ - θ) which is the average displacement of the estimate (it is directional and can be positive or negative) and measures systematic bias. RMSE is √sqrt ((1/N) $\Sigma_i$ ($\theta_i$ - θ)²) which is the typical magnitude of the deviation of the estimate from the true value (it is always non-negative) and satisfies RMSE² = bias² + variance (so it contains both the average displacement and the variation across the replications). In addition the median absolute error (median(|$\theta_i$ - θ|)) is used which is the median distance of the estimate from the true value and is robust to outliers.

The bias and RMSE are both dimensional absolute quantities (i.e. they have the same units as the exposure variable X) rather than being percentages. Since X is standardised to a standard normal distribution (with SD = 1) in the simulation, the bias and RMSE in Tables 2 to 8 are all in units of the standard deviation of X. A bias of + 0.20 for example would mean that the estimated inflection point is on average 0.20 standard deviations higher than the true value (not 20%). The clinical interpretation of such a bias will of course depend on the actual variability of the exposure variable. The point estimates and 95% CIs in the real data cases are given in the original units of the exposure variable (e.g. mg/dL, years etc).

# 2.4 Real data validation design

To assess the method's operability on real data and the clinical interpretability of the results we apply it to two real datasets.

**Case 1**

**Phosphate and death.** A cohort of patients with AKI on CRRT (from a public dataset [18]) uses 24 hour serum phosphate ($P_{24h)}$ as the exposure and 90 day death as the outcome under a Cox proportional hazards model. The relationship is monotonically increasing with no inflection point so the level of phosphate to be targeted can be determined solely by the crossing point. This is a case of estimating a crossing point in a setting with no inflection point.

**Case 2**

**Age and ectopic pregnancy.** A dataset of 2060 patients with ectopic pregnancy from a public dataset [19], aged 16 to 45 years, uses age as the exposure and ectopic pregnancy as the outcome, adjusting for 12 clinical covariates. The relationship is expected to be U-shaped. This case looks at the full application of the method in a scenario with an inflection point and two crossing points.

Both cases follow the same 4 step workflow: SHAP screening, frequentist RCS estimation, Bayesian RCS (MCMC) validation and consistency assessment. The frequentist RCS key points are derived via bootstrap resampling (1000 resamples with replacement), each of which involves refitting the model and re-solving the key points. The 95% confidence intervals are then derived as the 2.5th and 97.5th percentiles. The Bayesian RCS uses the MCMC posterior distribution to derive the 95% credible intervals.

# 3 Results

This chapter has two halves. The first (sections 3.1 and 3.2) looks at the performance of frequentist RCS, Bayesian RCS and SHAP in the estimation of inflection and crossing points, in terms of their error and their applicable range. The second half (section 3.3) deals with when RCS estimation is in fact the more accurate approach.

# 3.1 Comparison of methods for estimating inflection and crossing points

Table 2
Comparison of methods for inflection-point estimation

| Method | Bias | RMSE |
|---|---|---|
| Frequentist RCS extremum | + 0.200 | 0.223 |
| Bayesian RCS extremum, MCMC | + 0.248 | 0.258 |
| SHAP minimum | + 0.020 | 0.241 |

Note. Symmetric U-shape, n = 2000, ρ = 0.5, true inflection point at -0.6, 200 replications for all three methods. Inflection points are defined as: the frequentist RCS extremum, the mean of the Bayesian MCMC posterior extrema and the minimum of the SHAP dependence plot.

For inflection point estimation (as per Table 2) none of the three methods fail. The frequentist RCS extremum has a bias of around + 0.20 and the Bayesian MCMC extremum around + 0.25, both of which are similar and stem from the approximation of the true curve by the RCS spline. The SHAP minimum has a bias of approximately + 0.02 and is essentially unbiased. It does however have the highest RMSE of 0.24 so a single estimate is often off the true value by around 0.24 SDs. If INLA is used to estimate the inflection point the bias appears to be very close to zero. That bias is due to the distortion of the highest order spline coefficient under the Poisson approximation and so the estimate is not trustworthy. The Bayesian estimate of the inflection point is therefore MCMC.

Table 3
Bias and RMSE of crossing-point estimation under different dose-response shapes for RCS, Bayesian, and SHAP

| True shape | True crossing point | RCS | | Bayesian | | SHAP | |
|---|---|---|---|---|---|---|---|
| | | Bias | RMSE | Bias | RMSE | Bias | RMSE |
| Nonlinear monotone | -0.236 | -0.230 | 0.249 | -0.029 | 0.065 | -0.120 | 0.141 |
| Symmetric U, right root | 0.566 | + 0.389 | 0.393 | + 0.034 | 0.063 | -0.182 | 0.206 |
| Sigmoid monotone | 0.118 | + 0.025 | 0.066 | -0.367 | 0.378 | -0.183 | 0.186 |
| Plateau | -0.316 | -0.025 | 0.068 | -0.159 | 0.179 | -0.119 | 0.124 |

Note. ρ = 0.5, n = 2000 and 200 replications for RCS, Bayesian and SHAP. The Bayesian column is the mean of the MCMC posterior zero crossings. The nonlinear monotone and symmetric U shapes are quadratic whereas the sigmoid and plateau shapes are non-quadratic. The symmetric U shape has two zero crossings and only the right one is reported here (the left one is in Table 5).

Table 3 demonstrates that method performance varies considerably depending on the dose-response shape. For the quadratic-type shapes, RCS exhibits a clear spline approximation bias (biases of -0.23 and + 0.39 for the nonlinear monotone and the symmetric U right root respectively). Bayesian MCMC is considerably better in this case (biases of -0.03 and + 0.03). For the non-quadratic shapes RCS is essentially unbiased (bias ~ ± 0.03). The Bayesian method is worse in this case as the prior shrinkage introduces bias (biases of -0.16 to -0.37) and has a higher RMSE than RCS. SHAP does not break down under any of the four shapes but its crossing point is systematically biased in the negative direction by approximately − 0.12 to -0.18. Bayesian MCMC is therefore the better method for crossing points in most of the shapes, though frequentist RCS is more stable when the crossing point is far from the mean or on a sparse branch.

# 3.2 Detailed comparison of estimation methods

At the level of the estimation methods we compare frequentist RCS, Bayesian RCS and SHAP systematically. The key point is that Bayesian RCS naturally provides posterior intervals and its merits are a function of the net offset between the RCS spline bias and the prior shrinkage. For quadratic type shapes the spline bias is offset and the Bayesian method can be superior. For non-quadratic shapes the prior shrinkage is a pure cost and the Bayesian method is inferior (see Table 3). At the inflection point for the symmetric U shape Bayesian MCMC and frequentist RCS are comparable with both suffering from an RCS spline approximation bias of approximately + 0.2. INLA is not reliable at the inflection point or at the right root as the highest order spline coefficient is distorted. SHAP is robust to shape and collinearity and does not break down. It does however have an inherent negative bias of around − 0.18 to -0.24 which is constant with respect to max_depth, nrounds and the loess smoothing parameter and it is unreliable on sparse branches and for identifying two roots (lacks an analytic basis function

justification). Bayesian MCMC is therefore preferred for crossing points, frequentist RCS for inflection points and frequentist RCS is used for crossing points that are far from the mean or on a sparse branch. SHAP is only suitable for initial visual screening and not for precise localisation.

# 3.2.1 Crossing-point estimation

Table 4
Bias and RMSE of crossing-point estimation under the nonlinear monotone shape at ρ = 0.5 and 0.8

| Collinearity ρ | SHAP | | Bayesian MCMC | | Frequentist RCS | |
|---|---|---|---|---|---|---|
| | Bias | RMSE | Bias | RMSE | Bias | RMSE |
| 0.5 | -0.120 | 0.141 | -0.029 | 0.065 | -0.230 | 0.249 |
| 0.8 | -0.128 | 0.157 | -0.032 | 0.069 | -0.227 | 0.249 |

Table 4 (nonlinear monotone shape with one crossing point) compares SHAP, Bayesian MCMC and frequentist RCS at the crossing point for ρ = 0.5 and 0.8. The frequentist column is the frequentist RCS spline (rms::cph with rcs(X,4)). The Bayesian column is the mean of the MCMC posterior zero crossings from brms with cmdstanr. The SHAP column is the dependence plot zero crossing. All three use the zero crossing closest to 0. The linear shape is not included in this table as the crossing point is always just the exposure mean (trivial). The U-shaped crossing points are a pair (the left and right roots) so the single root convention doesn't apply; their estimation is in Section 3.2.3 and Table 5.

**Note**

the actual crossing point under the nonlinear monotone curve is -0.236, with n = 2000 and 200 replications for all three methods.

**Nonlinear monotone shape.** SHAP has a bias of -0.120 to -0.128 and an RMSE of 0.141 to 0.157, stable across all levels of collinearity. Frequentist RCS has a bias of -0.227 to -0.230 and an RMSE of around 0.25, again constant across all levels of collinearity. This is the approximation bias of the RCS spline basis functions for the quadratic polynomial truth (i.e. the bias at the level of the basis functions rather than of the frequentist estimation paradigm per se). When a matching quadratic basis function is used the frequentist crossing point becomes nearly unbiased (bias of about + 0.002). Bayesian MCMC has a bias of only − 0.029 to -0.032, again stable across all levels of collinearity. The conclusion that the Bayesian crossing point is superior to the frequentist one is therefore credible.

**U-shaped relationships.** In contrast to the monotonic non-linear shape, the roots of a U-shaped relationship occur in pairs (the left root and the right root) rather than at a single point. To report just a single root would be a definitional error and so the U-shape is not present in Table 4. When the two roots are properly identified as in Table 5, the frequentist RCS has biases of + 0.335 and + 0.389 for the left and right roots of the symmetric U-shape respectively, whilst the Bayesian MCMC has biases of + 0.348

and + 0.034. For the asymmetric U-shape the right root on the dense branch has a bias of + 0.117, the left root crossing point on the sparse branch has a bias of -0.630 and an RMSE of 0.672 (a very poor performance) and this is due to the data not adequately covering that branch. This is irrespective of the estimation method as discussed in Section 3.2.3.

**Best option is** Bayesian MCMC for crossing point estimation. With the nonlinear monotonic shape the bias is around − 0.03 which is considerably less than the − 0.23 of the frequentist RCS (see Table 4) and it does provide credible intervals. That advantage is due to the prior shrinkage cancelling out the RCS spline bias. When the crossing point is far from the exposure mean or on a very sparse branch the prior shrinkage pulls it off target and the frequentist approach is superior (see Tables 3 and 5) and in that situation the crossing point should be taken from the frequentist RCS. SHAP has the smallest bias but it has an inherent negative bias and is dependent on the smoothing parameters so is not recommended for crossing point estimation.

Table 5
Estimation of key points under U-shaped relationships for frequentist RCS and Bayesian MCMC

| Shape | Key point | True value | Collinearity ρ | Frequentist RCS | | Bayesian MCMC | |
|---|---|---|---|---|---|---|---|
| | | | | Bias | RMSE | Bias | RMSE |
| Symmetric U | Inflection point | -0.600 | 0.5 | + 0.202 | 0.214 | + 0.248 | 0.258 |
| | | | 0.8 | + 0.205 | 0.226 | + 0.247 | 0.264 |
| | Left-root crossing point | -1.766 | 0.5 | + 0.335 | 1.426 | + 0.348 | 0.361 |
| | | | 0.8 | + 0.380 | 1.435 | + 0.364 | 0.379 |
| | Right-root crossing point | 0.566 | 0.5 | + 0.389 | 0.393 | + 0.034 | 0.063 |
| | | | 0.8 | + 0.402 | 0.409 | + 0.045 | 0.076 |
| Asymmetric U | Inflection point | 0.468 | 0.5 | -0.055 | 0.062 | -0.316 | 0.320 |
| | | | 0.8 | -0.060 | 0.069 | -0.315 | 0.319 |
| | Sparse-branch left-root crossing point | -0.604 | 0.5 | -0.630 | 0.672 | -0.802 | 0.814 |
| | | | 0.8 | -0.615 | 0.740 | -0.759 | 0.773 |
| | Dense-branch right-root crossing point | 1.298 | 0.5 | + 0.117 | 0.125 | -0.466 | 0.468 |
| | | | 0.8 | + 0.102 | 0.284 | -0.461 | 0.463 |

Note: the inflection point is taken as the extremum of the RCS fitted curve and the crossing point as the zero crossing of the centred curve. Both the Frequentist RCS and the Bayesian MCMC use 200 replications. The true roots of the symmetric U are − 1.766 and 0.566 and the true roots of the asymmetric U are − 0.604 and 1.298.

# 3.2.2 Inflection-point estimation

Table 6
Estimation of key points under three outcome types for frequentist RCS

| Outcome | Collinearity ρ | Inflection point | | Right-root crossing point | |
|---|---|---|---|---|---|
| | | Bias | RMSE | Bias | RMSE |
| Cox | 0.5 | + 0.202 | 0.214 | + 0.388 | 0.393 |
| Cox | 0.8 | + 0.205 | 0.226 | + 0.402 | 0.409 |
| Logistic | 0.5 | + 0.198 | 0.247 | + 0.378 | 0.392 |
| Logistic | 0.8 | + 0.147 | 0.311 | + 0.367 | 0.387 |
| Competing risk | 0.5 | + 0.115 | 0.149 | + 0.199 | 0.205 |
| Competing risk | 0.8 | + 0.116 | 0.172 | + 0.205 | 0.215 |

Table 6 compares the inflection point and crossing point estimates of frequentist RCS under the three outcome types. The bias of the inflection point under Cox, logistic and competing risk outcomes is approximately + 0.20, + 0.15 to + 0.20 and + 0.12 respectively and the bias of the right root crossing point is approximately + 0.20 to + 0.40. None of them break down in a systematic way and all are robust to collinearity (rho). This confirms that frequentist RCS is a robust baseline across the different outcome types. For the crossing point estimation, Bayesian MCMC has lower bias in most of the cases as shown in Table 4 so the crossing point favours the Bayesian and the inflection point favours the frequentist.

### 3.2.3 Asymmetric U shape

The asymmetric U shape presents a problem for both types of RCS as can be seen in Table 5. Firstly, the left root crossing point on the sparse branch fails miserably regardless of whether you use frequentist RCS or Bayesian MCMC, with biases ranging from − 0.62 to -0.80 and RMSEs from 0.67 to 0.81. That establishes the fundamental point that no method can fix the problem that the data don't cover the shape. Secondly, at the inflection point and at the right root on the dense branch frequentist RCS is acceptable (biases of -0.06 and + 0.12 respectively). Bayesian MCMC is worse in this case as the weakly informative prior shrinkage is pulling that key point towards the mean (biases of -0.32 and − 0.47 respectively). This is the opposite of the situation we saw at the right root of the symmetric U shape where the Bayesian approach was superior. The lesson is that the Bayesian prior shrinkage is either beneficial or detrimental depending on the net offset between the RCS spline bias and the prior shrinkage, not on the location of any one particular key point.

## 3.3 Conditions for RCS estimation precision

RCS estimation accuracy is a function of sample size, number of events, collinearity and shape coverage. Frequentist RCS is collinearity robust (as demonstrated in Table 4) and a crossing point on a sparse branch cannot be 'rescued' by the method alone as discussed in Section 3.2.3. This section looks at the impact of sample size and number of events.

## 3.3.1 Sample size and number of events

Table 7
Effect of sample size on RCS estimation of inflection and crossing points

| Sample size n | Inflection point | | | Crossing point, right root | | |
|---|---|---|---|---|---|---|
| | Bias | RMSE | Missing | Bias | RMSE | Missing |
| 100 | + 0.135 | 0.477 | 0 | + 0.043 | 0.217 | 47 |
| 200 | + 0.213 | 0.429 | 0 | + 0.055 | 0.181 | 47 |
| 500 | + 0.174 | 0.353 | 0 | + 0.181 | 0.214 | 42 |
| 1000 | + 0.226 | 0.256 | 0 | + 0.284 | 0.294 | 35 |
| 2000 | + 0.202 | 0.214 | 0 | + 0.388 | 0.393 | 25 |
| 5000 | + 0.192 | 0.197 | 0 | + 0.503 | 0.505 | 14 |

Table 8
Effect of event rate on RCS estimation of inflection and crossing points

| Event rate | Inflection point | | Crossing point, right root | |
|---|---|---|---|---|
| | Bias | RMSE | Bias | RMSE |
| 0.1 | -0.003 | 0.763 | + 0.433 | 0.456 |
| 0.3 | + 0.167 | 0.319 | + 0.371 | 0.387 |
| 0.5 | + 0.198 | 0.247 | + 0.378 | 0.392 |
| 0.7 | + 0.186 | 0.264 | + 0.344 | 0.374 |

The accuracy of RCS estimation improves with increasing information, though a structural bias is not eliminated by increasing sample size. The RMSE of the inflection point drops from 0.477 to 0.197 as sample size increases and from 0.763 to 0.247 as the event rate increases, but the bias remains at approximately + 0.20. The bias of the right root crossing point remains at around + 0.4 to + 0.5 and is improved with sample size (the number of missing estimates drops from 47 at n = 100 to 14 at n = 5000). This bias is due to the use of the four knot RCS spline to approximate the true curve rather than due to finite sample variance and is in agreement with the theoretical result that break point estimation converges slowly at $n^{-1/3}$ [12].

# 3.4 Real data case analysis

Applying the method to the two real datasets described in Section 2.4 gives results which are in line with the simulation and are clinically interpretable.

# 3.4.1 Phosphate and death

The data are from a public dataset (n = 1144 AKI patients receiving CRRT, 90 day death as the outcome) [18]. First we do a SHAP screen of the variables. The SHAP values for P24h generally increase with phosphate level and the dependence plot crosses zero at around 4.7 mg/dL which suggests a directional effect on the outcome and potential non-linearity; this variable would be of interest in the key point analysis. The dependence plot for P24h does, however, show some local wiggles so the point at which it crosses zero should not be reported as exact.

In step 2 we use frequentist RCS to get the results. The P24h dose response curve is monotonic over the observed range with no inflection point (so we don't estimate one). The centred curve crosses zero at 6.51 mg/dL (95% bootstrap CI 6.07 to 6.96). The assumption checks are as follows: the nonlinearity test of RCS against the linear term gives a $\chi^2$ of 8.48 and a P of 0.014 so nonlinearity is significant but only modest in size. With n = 1144 and 821 events the amount of information is way beyond what is required to estimate a crossing point. The crossing point is within the observed range so the estimation assumptions are satisfied.

In step 3 the Bayesian RCS (MCMC) is used to check for stability and to produce posterior intervals. The Bayesian MCMC converges following standardisation of the covariates and the posterior mean of the crossing point is 4.69 mg/dL (95% credible interval 4.32 to 4.92).

In step 4 we look at consistency. The crossing point is 0.8 SDs from the phosphate mean so the fallback rule is not invoked. The Bayesian MCMC crossing point is 4.69 mg/dL (95% CI 4.32 to 4.92) which is quite clearly at odds with the frequentist RCS value of 6.51 mg/dL. The two values are different in both their definition of the crossing point and in the amount of prior shrinkage applied. In accordance with the principle that the crossing point is Bayesian, this case takes the Bayesian crossing point of 4.69 mg/dL which is in line with the SHAP dependence plot zero crossing of 4.7 and the clinical upper normal limit of 4.5 mg/dL for phosphate (making it the more clinically interpretable value). The frequentist value of 6.51 is reported in parallel as a point of contrast with a larger spline approximation bias.

**Summary**

A full application of the 4 step workflow in a monotone shape with no inflection point. SHAP confirms the variable is worth looking at, the frequentist RCS gives a crossing point of 6.51 mg/dL (95% CI 6.07 to 6.96), the Bayesian MCMC gives a crossing point of 4.69 mg/dL (95% credible interval 4.32 to 4.92) and in accordance with the principle that the crossing point is Bayesian we take 4.69 mg/dL. This is the phosphate level at which the risk goes from negative to positive relative to the population mean and is consistent with the clinical upper normal limit of 4.5 mg/dL.

# 3.4.2 Age and ectopic pregnancy

The data are from a publicly available dataset of patients with ectopic pregnancy (age range 16–45 years) [19]. SHAP is used to screen the variables and the SHAP importance ranks age as second. The dependence plot for age is clearly U-shaped with a trough at around 32.9 years, indicating a strong non-linearity between age and risk of ectopic pregnancy and thus justifies its inclusion in the key point analysis.

In the second step we report the frequentist RCS results. The dose response curve is U-shaped after fitting, and after adjustment for 12 clinical covariates it has one inflection point and two zero crossings. We take the inflection point as the curve extremum (32.6 years, 95% CI: 29.8 to 33.9). The crossing points are found by numerically determining the zero crossings of the centred curve, which gives two roots. The left root is 23.1 years (95% CI: 20.0 to 25.2) and the right root is 38.6 years (95% CI: 37.5 to 39.7). These are the key points of interest. The non-linearity test of age gives chi-square of 34.49 and $P < 0.0001$ so the U shape is very clear. With n = 2060 and 1030 events we have sufficient information. All of the key points are within the observed range and there are no sparse branches so the estimation assumptions are satisfied.

Step 3 is Bayesian RCS (MCMC) which checks for stability and provides posterior intervals. The Bayesian MCMC converges with a posterior mean of 30.9 years for the inflection point and a 95% credible interval of 29.1 to 31.8. The two crossing points are 25.7 and 34.4 years with 95% credible intervals of 22.9 to 27.5 and 33.3 to 35.4 respectively.

For the inflection point the two methods are in agreement so the conclusion is sound (the frequentist value of 32.6, the Bayesian value of 30.9 or something near the midpoint would all be acceptable). The two crossing points are both a little over 1.1 to 1.3 SDs from the age mean which invokes the fallback rule so the crossing points take the frequentist values of 23.1 and 38.6. The Bayesian values of 25.7 and 34.4 are reported in parallel as an example of prior shrinkage. The event rates are all consistent with the U-shape (66.3% in the 16–20 age group, a minimum of 45.9% in the 28–32 age group and then a rise back to 68.0% in the 40–45 age group).

This is a full application of the 4 step workflow in the case of a U-shaped relationship with an inflection point and two crossing points. SHAP confirms the non-linearity, the frequentist RCS gives an inflection point of 32.6 years (the frequentist estimate) and the Bayesian MCMC gives 30.9 years (again confirming it). The two crossing points are both more than one SD from the mean and so revert to the frequentist values of 23.1 and 38.6 years. Ectopic pregnancy risk is high on both the very young and the old sides and is at its lowest at the intermediate ages, in line with the observed event rates.

The two real world cases also serve to ground the simulation bias (expressed in units of standard deviation) on a clinically meaningful scale. In the ectopic pregnancy case the SD of age is 6.28 years so the frequentist RCS inflection point bias of approximately + 0.20 SD in the simulation is equivalent to about 1.3 years and the RMSE of 0.22 SD is about 1.4 years. In the phosphate case the SD of $P_{24h}$ is 2.37 mg/dL and the 95% credible interval for the Bayesian crossing point of 4.69 mg/dL is 4.32 to 4.92 (a width of 0.6 mg/dL or about 0.25 SD). In other words the RCS bias is not some abstract percentage but

an absolute displacement of some fraction of a standard deviation of the exposure variable and its clinical interpretation must be in terms of the variability of the exposure variable.

## 4 Discussion

# 4.1 Rationale and applicable conditions of the RCS method

The methodological conclusions are as follows: use Bayesian MCMC for crossing points and frequentist RCS for inflection points. If the crossing point is a long way from the exposure mean or on a sparse branch then switch to frequentist RCS. Use SHAP purely for exploratory screening of direction, importance, non-linearity and interaction. The two real examples go through the full four step workflow and produce clinically interpretable results, confirming the method is both workable and sensible. In the phosphate example SHAP confirms that $P_{24h}$ is worth looking at. Frequentist RCS finds no inflection point and a crossing point of 6.51 mg/dL. Bayesian MCMC gives a crossing point of 4.69 mg/dL. The crossing point is 0.8 SDs from the mean so the fallback is not triggered and the Bayesian value of 4.69 mg/dL is used. In the ectopic pregnancy example SHAP confirms the strong non-linearity of age. Frequentist RCS gives an inflection point of 32.6 years and Bayesian MCMC gives 30.9 years, which is consistent with this. The two crossing points are both more than 1 SD from the age mean so the fallback is triggered and the crossing points are therefore the frequentist values of 23.1 and 38.6 years. The two cases illustrate the two paths (Bayesian for crossing points and fallback to frequentist) and confirm that the rule for choosing by the location of the key point is both workable and sensible.

For the merits of the various methods, state the conclusions separately for each key point. For crossing point estimation, Bayesian MCMC is the preferred method. It is more accurate in most situations (bias of approximately − 0.03 for the nonlinear monotone crossing point and + 0.03 for the symmetric U right root as per Tables 4 and 5) and it naturally provides credible intervals. The frequentist RCS crossing point has greater bias which arises from the approximation of the true polynomial with the RCS spline. The Bayesian advantage here is that the weakly informative prior shrinkage offsets the spline bias. When the crossing point is far from the exposure mean or on a sparse branch the prior shrinkage pulls it off target and the frequentist method is in fact superior (see Tables 3 and 5). For inflection point estimation, the frequentist RCS is the preferred method with a bias of approximately − 0.06 to + 0.20 (vs -0.32 for the Bayesian MCMC under the asymmetric U shape). INLA is unreliable at the inflection point and the right root as the highest order spline coefficient is distorted. Bayesian estimation at these points therefore has to be via MCMC. SHAP has an inherent negative bias and is unreliable on sparse branches so it is only useful for exploratory screening of direction, importance, nonlinearity and interaction.

Whether the bias is clinically acceptable is a function of how much the risk estimate changes as a result of the key point being shifted, i.e. the product of the bias and the risk gradient near the key point, rather than the bias in itself. The real cases provide a clear answer. At the ectopic pregnancy trough of 32.6 years the risk only changes by around 0.3 percentage points per year so the 0.2 SD inflection point bias has essentially no effect on the risk assessment and is acceptable. At the phosphate crossing point of

4.69 mg/dL the relative risk increases by around 39% for each 1 mg/dL increase so 0.2 SDs (0.47 mg/dL) corresponds to an 18% difference in risk which is of the same order of magnitude as the sampling uncertainty of around ± 23% as reflected in the 95% credible interval of 4.32 to 4.92. The precision bottleneck at the crossing point is therefore the width of the confidence interval (determined by the sampling error) rather than the systematic bias. Report the key points with their confidence intervals and when a key point is being used as a hard intervention threshold then assess the clinical consequences of the bias in conjunction with the risk gradient in the vicinity of the key point.

The Bayesian RCS is implemented throughout using MCMC via brms and cmdstanr with 200 replications per scenario and 400 iterations per chain. The inflection point and crossing point scenarios all converge quite nicely (though there are a few very rare instances of divergence). That said the scale is more than sufficient to support the qualitative conclusion we've already drawn above that the benefit or detriment of prior shrinkage will depend on the dose response shape. As the final values prior to submission, we'd recommend rechecking the key scenarios with longer chains and a higher adapt_delta.

# 4.2 Decomposition and validation of bias sources

The bias of the three methods in estimating key points can be broken down into several sources. **Source 1 is spline approximation.** RCS has a spline approximation bias of approximately + 0.2 to + 0.5 (as shown in Table 7), due to the use of a 4-knot spline to approximate the true curve and this bias does not go to zero with increasing sample size. This bias is present at the inflection point for both the frequentist and Bayesian MCMC approaches. **Source 2 is the inherent bias** of the SHAP dependence plot. The SHAP dependence plot has an inherent bias of approximately − 0.18 to -0.24 at the crossing point and this bias does not improve with tuning. **Source 3 is the prior shrinkage and the choice of implementation in the Bayesian approach.** The weakly informative prior will shrink the Bayesian posterior mean towards the prior centre and the benefit or detriment of this will be a function of the net offset between the RCS spline bias and the prior shrinkage. For quadratic-type shapes the spline bias is offset and the Bayesian approach is superior. For non-quadratic shapes the prior shrinkage is a pure burden and the Bayesian approach is inferior (as shown in Table 3). The choice of Bayesian implementation is also critical. INLA breaks down at the inflection point as the highest order spline coefficient is distorted under the Poisson approximation and it drifts with collinearity at the crossing point, whereas MCMC is robust at both. Bayesian estimation therefore must use MCMC. Ratio-type and root finding functionals are extremely sensitive to coefficient errors. **Source 4 is insufficient information.** The RMSE of the RCS inflection point drops from 0.477 to 0.197 as sample size increases and from 0.763 to 0.247 as the event rate increases (as shown in Tables 7 and 8) and when information is insufficient the bias is primarily due to high variance in keeping with the theory of slow convergence in break point estimation [12]. **Source 5 is a shape not adequately covered by the data.** For the asymmetric U shape the sparse branch crossing point has a bias of -0.80 and an RMSE of 0.81 whereas the dense branch is almost unbiased. This source is independent of the estimation method. In summary the spline approximation, the intrinsic SHAP bias and the choice of Bayesian implementation are all inherent limitations of the three methods. At the same time insufficient information and insufficient coverage are hard constraints at the data level.

## 4.3 Necessary conditions for the dataset and simulation evidence

The simulation results taken together demonstrate that in order to get accurate estimates of the RCS key points the dataset has to meet a number of necessary conditions, each of which is supported by the corresponding simulation values. **The data must cover the shape.** The asymmetric U simulation demonstrates that the sparse branch left root crossing point performs very poorly in both the frequentist RCS and the Bayesian MCMC cases (biases of- 0. 62 to- 0. 80 and RMSEs of 0. 67 to 0. 81) whereas the frequentist RCS bias for the dense branch right root is only + 0. 117. If one branch of the true shape is not observed then no matter how good the method is the key point on that branch cannot be estimated at all. **Nonlinearity must be testable.** The nonlinearity test of RCS against the linear term will confirm whether or not there is genuine nonlinearity and hence whether or not an inflection point exists (which is a prerequisite to avoiding a false inflection point). Visual inspection of the SHAP dependence plot in the case of the nonlinear monotone shape is only 53% correct and therefore cannot be the basis for any judgement. **There must be sufficient events and a sufficiently large sample size.** The sample size sensitivity simulation shows that the inflection point RMSE is approximately 0. 477 at n = 100 and falls to 0. 197 at n = 5000. The event rate simulation shows that the inflection point RMSE is approximately 0. 763 at an event rate of 0. 0.1 and drops below 0. 0.3 when the event rate is at least 0.3. 3 (see Tables 7 and 8). This is in line with the slow convergence theory of break point estimation [12] and with the general requirement of a sufficiently large sample size in continuous variable modelling [1]. **The target must be unambiguously an inflection point or a crossing point.** These are two very different things: the inflection point is an extremum and the crossing point is a zero crossing centred at the point of interest. In the case of a U shape the crossing point is also a pair (the left and right roots) and these must be estimated together. The most fundamental point is that whether or not the data cover the true dose response shape is determined at the data collection stage and statistical modelling cannot fix that.

## 4.4 Operational workflow and prerequisites for practical application

Combine the simulation and the real case evidence and you can use frequentist RCS in real data analysis to estimate inflection and crossing points. That said you do have to meet a number of prerequisites and carry out the corresponding operations. The first step is to fit the frequentist RCS and test for non-linearity. From the curve you then read the extremum and zero crossing and decide which of these key points you report. A monotone shape will only have one crossing point and a U-shaped relationship will have an inflection point and a pair of crossing points. The second step is to check that the key point lies within the observed range. An inflection or crossing point that lies outside the data support doesn't exist or can't be estimated and should not be reported. The third step is to check shape coverage. If there is sufficient data in the vicinity of the key point then the estimate is reliable. If the key point lies on a sparsely populated branch then the crossing point estimate is unreliable and you must state that the coverage is insufficient. The fourth step is to check the amount of information. The crossing point is

fairly stable when n is about 200 or the event rate is at least 10%. The inflection point requires n to be at least 500 or an event rate of at least 30% to be stable. The inflection point precision clearly deteriorates when the data are insufficient. The fifth step is to solve correctly. The crossing point is the zero crossing of the centred curve (i.e. after subtracting the mean effect) and the inflection point is the extremum of the fitted curve. Don't confuse the reference point zero crossing with a crossing point. The sixth step is to report uncertainty. Report 95% confidence intervals for the inflection and crossing points via bootstrap resampling rather than single point estimates. The seventh step is to interpret the key points. After you've read the key points you must answer two questions. Firstly whether the prerequisites have been met (i.e. is the non-linearity test significant, is the key point within the observed range, is there sufficient data in the vicinity and is the sample size and number of events sufficient). Secondly how large the bias is and what the clinical consequences of that bias are. Convert the structural bias to the original units by multiplying the exposure variable's standard deviation by the risk gradient at the key point and compare that with the width of the 95% confidence interval to assess whether the bias is acceptable. The eighth step is to choose the estimation paradigm. The inflection point will always be estimated using frequentist RCS. The crossing point will be estimated using Bayesian MCMC by default but will fall back to frequentist RCS under two rules after estimation. The first rule is that if the crossing point is more than about one standard deviation from the exposure mean then the prior shrinkage will clearly pull it off target and the frequentist approach is better. The second rule is that if the crossing point lies on a sparsely populated branch then both methods are unreliable, the frequentist approach is relatively more stable in that situation and the insufficient coverage must be reported honestly in that case. The judgement of how far from the mean the crossing point is is based on the difference between the two point estimates of the crossing point from the Bayesian and the frequentist approaches. If the difference between the Bayesian and frequentist crossing points is more than about 0.3 standard deviations then the crossing point is usually quite far from the mean and the frequentist value should be used. If you are using the Bayesian approach you should prefer MCMC as INLA is unreliable at the inflection point and the right root due to the distortion of the spline coefficients.

The above can be boiled down to a recommended analysis workflow. Step 1: use SHAP to screen the variables. Use the SHAP dependence plots and the importance ranking to make your initial exploratory judgements as to direction, importance, non-linearity and interactions and decide which variables are worth going on to the key point analysis (without reporting any key point values from SHAP). Step 2: use frequentist RCS to estimate the inflection point. Fit the RCS, test for non-linearity and take the extremum of the fitted curve as the inflection point. Provide a 95% CI via bootstrapping. Step 3: use Bayesian MCMC to estimate the crossing point by fitting Bayesian RCS and taking the zero crossing of the centred posterior curve as the crossing point. The credible interval is provided directly from the posterior. Step 4: fall back by rule and calculate the difference between the Bayesian and the frequentist crossing points. If the difference is greater than about 0.3 SD or the crossing point is on a sparse branch then the crossing point is switched to the frequentist RCS and the disagreement between the two methods is reported as model uncertainty. In this case the frequentist interval is a confidence interval and the Bayesian interval is a credible interval. They are of different origin and meaning and must be labelled as such.

## 5 Conclusion

This paper looks at frequentist RCS, Bayesian RCS and SHAP in turn and comes to the following conclusions. Use Bayesian MCMC for crossing points and frequentist RCS for inflection points. If a crossing point is more than approximately one SD away from the exposure mean or is on a sparse branch then switch to frequentist RCS. Bayesian estimation has to be by MCMC as INLA is unreliable at the inflection point and at the right root due to the distorted coefficients. SHAP is not used to estimate the key points. It is used as a screening tool prior to key point estimation. It uses importance ranking to select variables to be considered and dependence plots to look at the direction of the effect and to check for non-linearity and interactions. It is then used to decide whether and how to enter the RCS analysis, with no reported inflection or crossing point values from it. The precision of both types of RCS improves with sample size and number of events and both are validated under Cox, logistic and competing risk outcomes. This gives us a mechanism which is clear, validated and directly applicable for the identification of key points in dose response analysis.

## Declarations

**Ethics approval and consent to participate:**

Not applicable.

**Consent for publication:**

Not applicable.

# Competing interests:

The author declares no competing interests.

# Funding:

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

# Authors' contributions:

YL designed the study, performed the simulations and analyses, interpreted the results, and wrote the manuscript.

# Acknowledgements:

Not applicable.

# Availability of data and materials:

The datasets analysed during the current study are publicly available in the Dryad digital repository (references [18] and [19]). Simulation code is available from the corresponding author on reasonable request.

## References


1. Lopez-Ayala P, Riley RD, Collins GS, Zimmermann T (2025) Dealing with continuous variables and modelling non-linear associations in healthcare data: a practical guide. BMJ 390:e082440. 10.1136/bmj-2024-082440
2. Kahan BC, Rushton H, Morris TP, Daniel RM (2016) A comparison of methods to adjust for continuous covariates in the analysis of randomised trials. BMC Med Res Methodol 16:70. 10.1186/s12874-016-0141-3
3. Li F, He M, Zhou M et al (2023) Association of C-reactive protein with mortality in Covid-19 patients: a secondary analysis of a cohort study. Sci Rep 13:20361. 10.1038/s41598-023-47680-x
4. Nguyen VK, Colacino J, Chung MK et al (2021) Characterising the relationships between physiological indicators and all-cause mortality (NHANES): a population-based cohort study. Lancet Healthy Longev 2(10):e651–e662. 10.1016/S2666-7568(21)00212-9
5. Hu XF, Lowe M, Chan HM (2021) Mercury exposure, cardiovascular disease, and mortality: a systematic review and dose-response meta-analysis. Environ Res 193:110538. 10.1016/j.envres.2020.110538
6. Durrleman S, Simon R (1989) Flexible regression models with cubic splines. Stat Med 8(5):551–561. 10.1002/sim.4780080504
7. Harrell FE Jr (2015) Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis, 2nd edn. Springer, New York. 10.1007/978-3-319-19425-7
8. Royston P, Altman DG (1994) Regression using fractional polynomials of continuous covariates: parsimonious parametric modelling. J Roy Stat Soc: Ser C (Appl Stat) 43(3):429–467. 10.2307/2986270
9. Royston P, Sauerbrei W (2008) Multivariable Model-Building: A Pragmatic Approach to Regression Analysis Based on Fractional Polynomials for Modelling Continuous Variables. Wiley, Chichester. 10.1002/9780470770771
10. Muggeo VMR (2003) Estimating regression models with unknown break-points. Stat Med 22(19):3055–3071. 10.1002/sim.1545

11. Pastor-Barriuso R, Guallar E, Coresh J (2003) Transition models for change-point estimation in logistic regression. Stat Med 22(7):1141–1162. 10.1002/sim.1045
12. Chen Y (2021) Jump or kink: on super-efficiency in segmented linear regression breakpoint estimation. Biometrika 108(1):215–222. 10.1093/biomet/asaa049
13. Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. Adv Neural Inform Process Syst (NeurIPS) 30:4765–4774
14. Chen T, Guestrin C (2016) XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. :785–794. 10.1145/2939672.2939785
15. Breiman L (2001) Statistical modeling: the two cultures. Stat Sci 16(3):199–231. 10.1214/ss/1009213726
16. Rue H, Martino S, Chopin N (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. J Roy Stat Soc B 71(2):319–392. 10.1111/j.1467-9868.2008.00700.x
17. Cox DR (1972) Regression models and life-tables. J Roy Stat Soc B 34(2):187–220. 10.1111/j.2517-6161.1972.tb00899.x
18. Jung SY, Kwon J, Park S et al (2019) Data from: phosphate is a potential biomarker of disease severity and predicts adverse outcomes in acute kidney injury patients undergoing continuous renal replacement therapy [Dataset]. 10.5061/dryad.6v0j9. Dryad
19. Jacob L, Kalder M, Kostev K (2017) Data from: risk factors for ectopic pregnancy in Germany: a retrospective study of 100,197 patients [Dataset]. 10.5061/dryad.pv805. Dryad

# Figures

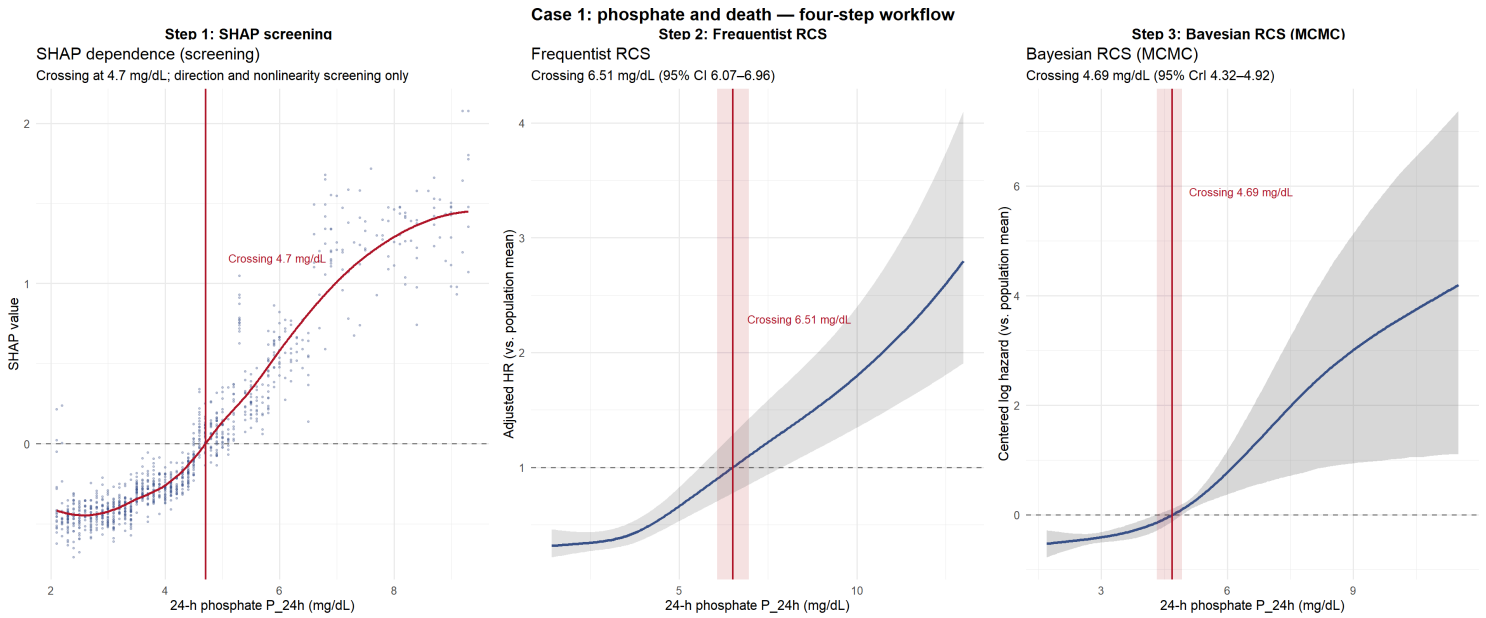



**Figure 1**

**Case 1, phosphate and death. The four-step workflow of SHAP screening, frequentist RCS, and Bayesian RCS.**

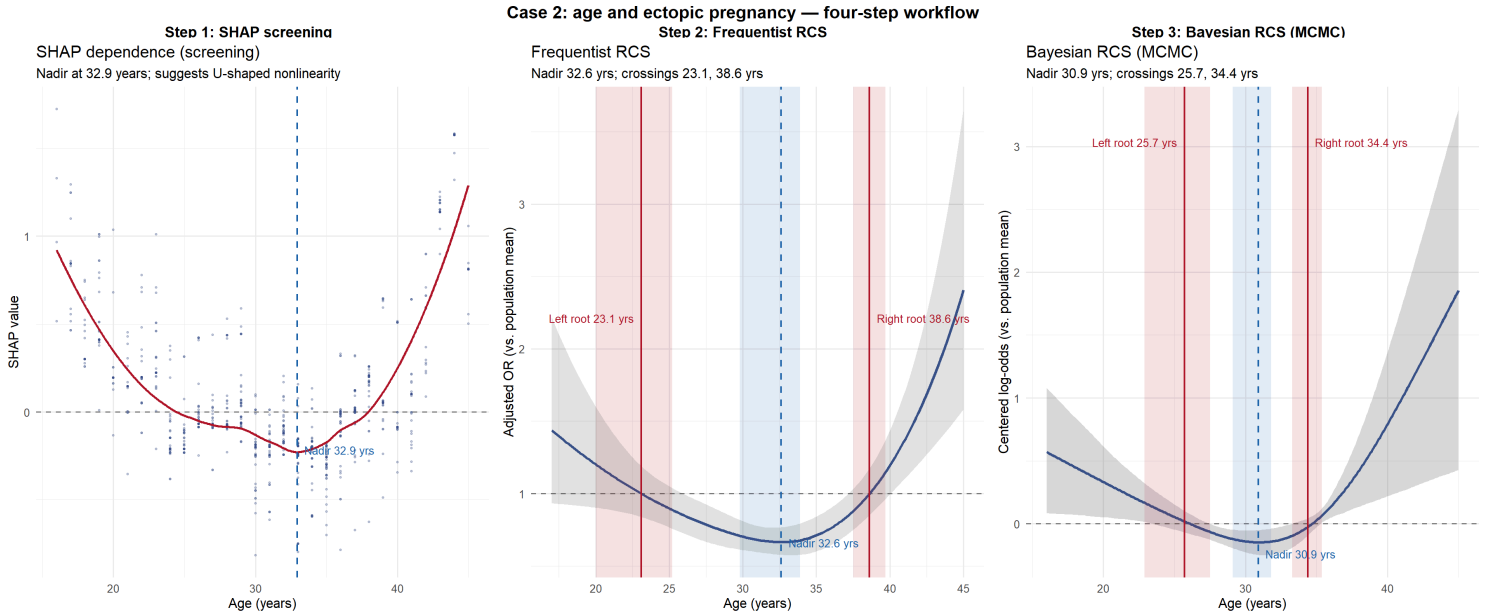



**Figure 2**

**Case 2, age and ectopic pregnancy. The four-step workflow of SHAP screening, frequentist RCS, and Bayesian RCS.**