Bayesian Semiparametric Inference for Dynamic Treatment Strategies with Informative Timing

Arman Oganisian

Kelly Getz, Todd Alonzo, Richard Aplenc, Jason Roy

Department of Biostatistics
Brown University

Study Motivation

Pediatric acute myeloid leukemia (AML) is a cancer of the blood and bone marrow.

Patients move through sequence of chemotherapy courses

- Anthracyclines (ACT) therapy is effective at suppressing cancer.
- ACT also lowers ejection fraction (EF), which can worsen survival.

Goal: Estimate (and optimize) effect of ACT treatment rules on survival.

Data: COG AAML1031 Phase III Trial

Reference:

Arman Oganisian, Kelly D Getz, Todd A Alonzo, Richard Aplenc, Jason A Roy, Bayesian semiparametric model for sequential treatment decisions with informative timing, Biostatistics, 2024; kxad035, https://doi.org/10.1093/biostatistics/kxad035

Data Generating Process - AAML1031

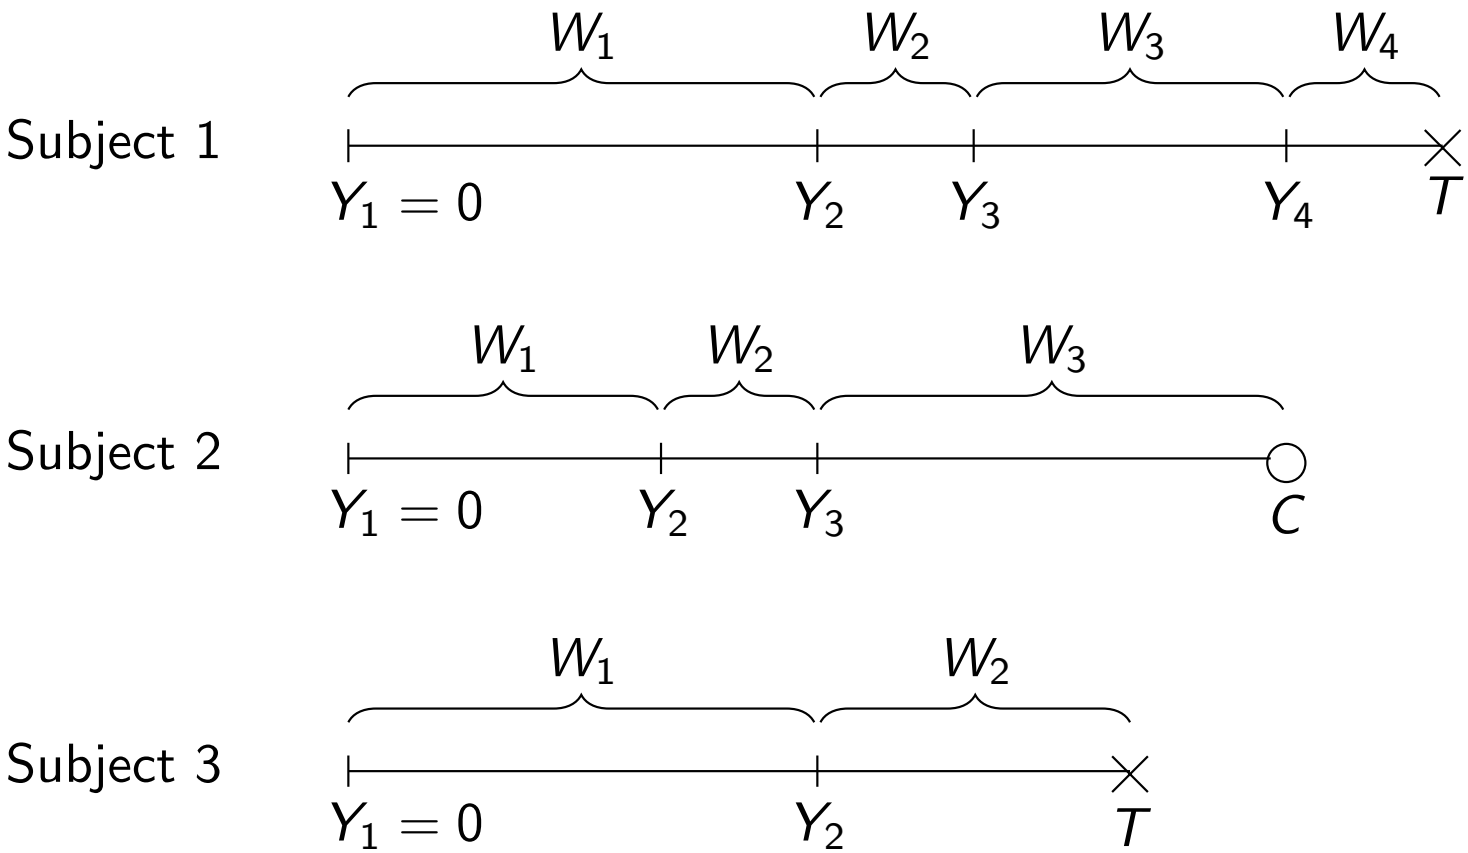

- Time of k^{th} decision, Y_k.
- At time Y_k, confounders L_k are measured...
- ... combined with previous history to decide treatment $A_k \in \{0, 1\}$.
- Timing varies across patients.

Our approach

- Bayesian semiparametric hazard models robust to misspecification.
- Respects continuous-time nature of transitions.
- Allows for covariate-dependent censoring and death before treatment course completion.
- Causal estimation: posterior survival probabilities adjusted for time-varying confounding and informative timing.
- Optimization: Posterior distribution over optimal rule parameters.
- Full posterior inference for other functionals of the joint.

Notation

History: $\bar{X}_k = (X_1, X_2, \ldots, X_k)$; Future: $\underline{X}_k = (X_k, X_{k+1}, \ldots, X_K)$

- κ: number of treatment courses.
- For $k = 1, 2, \ldots, \kappa$, define W_k waiting time from treatment k to next event

$$W_k = min(T, Y_{k+1}, C) - Y_k$$

  with $\delta_k \in \{1, 0, -1\}$,
- Confounder history $\bar{L}_k$.
- Available information ahead of A_k, $H_k = (\bar{L}_k, \bar{W}_{k-1}, \bar{A}_{k-1})$.

The observed data for subject i consists of $\mathcal{D}_i = (\bar{L}_{\kappa_i}, \bar{A}_{\kappa_i}, \bar{W}_{\kappa_i}, \bar{\delta}_{\kappa_i})$. Full data denoted by $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^n$.

Dynamic ACT Assignment Rules

For $k = 1, 2, \ldots, K$, define rule

$$r_k : \mathcal{H}_k \to \mathcal{S}_k$$

- Available history space $\mathcal{H}_k$.
- Feasible set of treatment options $\mathcal{S}_k \subset \mathcal{A} = \{0, 1\}$
- $r = \{r_k\}_{k=1}^{K}$

Distinct from static treatments:

- Assignment is determined dynamically: $A_k = r_k(h_k)$.
- Example: $r_k(L_k; \tau) = I(L_k > \tau) I(A_{k-1} \neq 1)$.

Potential Outcomes and Target Estimand

- $K(r)$: potential number of treatment courses.
- $\bar{W}_{K(r)}(r) = \{W_1(r), W_2(r), \ldots, W_{K(r)}(r)\}$: potential waiting times.
- $T(r) = \sum_{k=1}^{K(r)} W_k(r)$: potential survival time.

Target estimands:

- Population-level survival rate: $\Psi^r(t) = P(T(r) > t)$.
- Contrast effects of rules: $\Psi^r(t)/\Psi^{r'}(t)$.
- For some t, find optimal rule $r^* = \text{argmax}_{r \in \mathcal{R}} \Psi^r(t)$

Must identify joint distribution of potential outcomes: $f^*(\bar{w}_k(r), K(r) = k)$

Identification Assumptions

- Sequential Ignorability:

$$\underline{W}_k(r), \underline{L}_k(r) \perp A_k \mid \bar{A}_{k-1}, \bar{L}_k, \bar{W}_{k-1}, \kappa \geq k$$

- Treatment Positivity:

$$P(A_k = a_k \mid \bar{H}_k = h_k, \kappa \geq k) > 0$$

for each $\bar{h}_k \in \mathcal{H}_k$ in support and each feasible $a_k \in \mathcal{S}_k$.

Other assumptions: ignorable censoring; censoring positivity; SUTVA.

A G-formula for $\Psi^r(t)$: case $K(r) = 2$

Identification via a version of the g-formula (Robins, 1986; Tsiatis et al., 2020).

$$f^*(\bar{w}_2(r), K(r) = 2) = \int_{\bar{\mathcal{L}}_2} f_{21}(w_2 \mid \bar{w}_1, \bar{a}_2^r, \bar{l}_2) g_2(l_2 \mid \bar{w}_1, \bar{a}_1^r, \bar{l}_1) f_{10}(w_1 \mid a_1^r, l_1) g_1(l_1) d\bar{l}_2$$

- $f_{ks}(w_k \mid \bar{w}_{k-1}, \bar{a}_k^r, \bar{l}_k)$: sub-density function for waiting time, w_k, until event of type $s \in \{0, 1\}$.
- $g_k(l_k \mid \bar{w}_{k-1}, \bar{a}_{k-1}^r, \bar{l}_{k-1})$: model for distribution of confounders measured at course k.

A G-formula for $\Psi^r(t)$: case $K(r) = k$

Identification via a g-formula (Robins, 1986; Tsiatis et al., 2020).

$$f^*(\bar{w}_k(r), K(r) = k) = \int_{\bar{\mathcal{L}}_k} f_{k1}(w_k \mid \bar{w}_{k-1}, \bar{a}^r_k, \bar{l}_k) g_k(l_k \mid \bar{w}_{k-1}, \bar{a}^r_{k-1}, \bar{l}_{k-1}) \times \prod_{j=1}^{k-1} f_{j0}(w_j \mid \bar{w}_{j-1}, \bar{a}^r_j, \bar{l}_j) g_j(l_j \mid \bar{w}_{j-1}, \bar{a}^r_{j-1}, \bar{l}_{j-1}) \, d\bar{l}_k$$

- $f_{ks}(w_k \mid \bar{w}_{k-1}, \bar{a}^r_k, \bar{l}_k)$: sub-density function for waiting time, w_k, until event of type $s \in \{0, 1\}$.
- $g_k(l_k \mid \bar{w}_{k-1}, \bar{a}^r_{k-1}, \bar{l}_{k-1})$: model for distribution of confounders measured at course k.

The Cause-Specific Hazard

Model sub-density indirectly by modeling the cause-specific hazard:

$$f_{ks}(w_k \mid x_k) = \lambda_s(w_k \mid x_k) \exp\Big\{ - \int_0^{w_k} \sum_{v \in \{0,1\}} \lambda_v(u \mid x_k) du \Big\}$$

with feature vector $x_k = (\bar{w}_{k-1}, \bar{a}_k^r, \bar{l}_k)$.

Bayesian Proportional Hazard Models

Specify a proportional hazard model for waiting times:

$$\lambda_s(w_k \mid x_k) = \lambda_{ks0}(w_k) \exp\left(x_k' \beta_{ks}\right)$$

Priors

$$\lambda_{ks0} \sim GP(\alpha_{ks} \lambda^*_{ks0})$$
$$\beta_{ks} \sim f_{\beta_{ks}}(\beta_{ks})$$

- Empirically calibrated hyperparameters
- Good frequentist properties in simulations.

Illustration of Posterior Inference

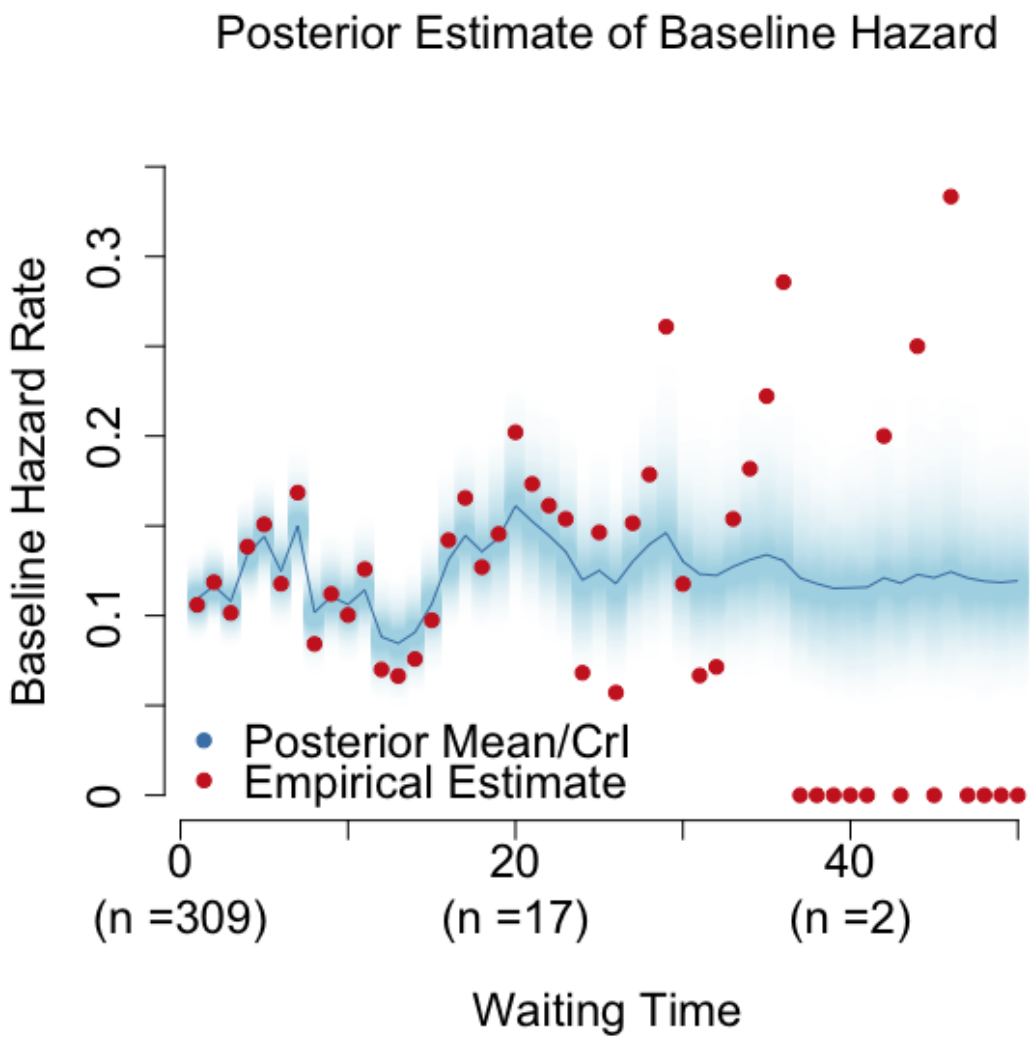

Confounder Trajectory Models

Abuse of notation, but let $x_k = (\bar{w}_{k-1}, \bar{a}^r_{k-1}, \bar{l}_{k-1})$

$$g_k(l_k \mid x_k; \eta_k)$$

- E..g. Gaussian with $E[L_k \mid x_k; \eta_k] = q_k(x_k; \eta_k)$,
- E.g. Bernoulli with $E[L_k \mid \bar{x}_k; \eta_{kp}] = \Phi(q_k(x_k; \eta_k))$

Can be as non-parametric or parametric as desired:

- E.g. $q_k \sim BART \rightarrow \ \eta_k$ is collection of tree structures.
- E.g. $q_k(x_k) = x_k' \eta_k \rightarrow \ \eta_k$ is collection of regression coefficients.

Posterior G-Computation - for $K = 2$

We develop adaptive MCMC algorithms to obtain $m = 1, 2, \ldots, M$ posterior draws for course $k = 1, 2$

$$\omega_k^{(m)} = \left\{ \lambda_{k00}^{(m)}, \beta_{k0}^{(m)}, \lambda_{k10}^{(m)}, \beta_{k1}^{(m)}, \eta_k^{(m)} \right\}$$

For $b = 1, 2, \ldots, B$

1. Simulate $L_1^{(b)} \sim g_1(l_1 \mid \eta_1^{(m)})$ and make decision $a_1^r = r(L_1^{(b)})$
2. Simulate waiting time until next treatment $W_{01}^{(b)} \sim \lambda_0(w_1 \mid L_1^{(b)}, A_1 = a_1^r; \lambda_{100}^{(m)}, \beta_{10}^{(m)})$
3. Simulate waiting time until death $W_{11}^{(b)} \sim \lambda_1(w_1 \mid L_1^{(b)}, A_1 = a_1^r; \lambda_{110}^{(m)}, \beta_{11}^{(m)})$

Posterior G-Computation - for $K = 2$

We develop adaptive MCMC algorithms to obtain $m = 1, 2, \ldots, M$ posterior draws for course $k = 1, 2$

$$\omega_k^{(m)} = \left\{ \lambda_{k00}^{(m)}, \beta_{k0}^{(m)}, \lambda_{k10}^{(m)}, \beta_{k1}^{(m)}, \eta_k^{(m)} \right\}$$

For $b = 1, 2, \ldots, B$

1. Simulate $L_1^{(b)} \sim g_1(l_1 \mid \eta_1^{(m)})$ and make decision $a_1^r = r(L_1^{(b)})$
2. Simulate waiting time until next treatment $W_{01}^{(b)} \sim \lambda_0(w_1 \mid L_1^{(b)}, A_1 = a_1^r; \lambda_{100}^{(m)}, \beta_{10}^{(m)})$
3. Simulate waiting time until death $W_{11}^{(b)} \sim \lambda_1(w_1 \mid L_1^{(b)}, A_1 = a_1^r; \lambda_{110}^{(m)}, \beta_{11}^{(m)})$
4. if $W_{11}^{(b)} < W_{01}^{(b)}$, then set $T^{(b)} = W_{11}^{(b)}$ and STOP. Else

Posterior G-Computation - for $K = 2$

We develop adaptive MCMC algorithms to obtain $m = 1, 2, \ldots, M$ posterior draws for course $k = 1, 2$

$$\omega_k^{(m)} = \left\{ \lambda_{k00}^{(m)}, \beta_{k0}^{(m)}, \lambda_{k10}^{(m)}, \beta_{k1}^{(m)}, \eta_k^{(m)} \right\}$$

For $b = 1, 2, \ldots, B$

1. Simulate $L_1^{(b)} \sim g_1(l_1 \mid \eta_1^{(m)})$ and make decision $a_1^r = r(L_1^{(b)})$
2. Simulate waiting time until next treatment $W_{01}^{(b)} \sim \lambda_0(w_1 \mid L_1^{(b)}, A_1 = a_1^r; \lambda_{100}^{(m)}, \beta_{10}^{(m)})$
3. Simulate waiting time until death $W_{11}^{(b)} \sim \lambda_1(w_1 \mid L_1^{(b)}, A_1 = a_1^r; \lambda_{110}^{(m)}, \beta_{11}^{(m)})$
4. if $W_{11}^{(b)} < W_{01}^{(b)}$, then set $T^{(b)} = W_{11}^{(b)}$ and STOP. Else
5. Simulate $L_2^{(b)} \sim g_2(l_2 \mid L_1^{(b)}, A_1 = a_1^r, W_{01}^{(b)}, \eta_2^{(m)})$ and make decision $a_2^r = r(L_2^{(b)})$.
6. Simulate waiting time to death $W_{12}^{(b)} \sim \lambda_1(w_2 \mid \bar{L}_2^{(b)}, \bar{A}_2 = \bar{a}_2^r, W_{01} = W_{01}^{(b)}; \lambda_{210}^{(m)}, \beta_{21}^{(m)})$.
7. Set $T^{(b)} = W_{01}^{(b)} + W_{12}^{(b)}$

Posterior Inference for Estimands

Could compute survival rate as

$$\psi^r(t)^{(m)} = \frac{1}{B}\sum_{b=1}^{B} I(T^{(b)} > t)$$

- Using $m = 1, 2, \ldots, M$ draws, form point and interval estimates.
- Doing this for a range of t traces out entire survival curve.
- Could compute $\psi^r(t)^{(m)}$ for various rules and find the maximizing rule.
- Could also consider more general utility functions.

Analysis of AML Treatments in AAML1031

Features	$N = 292$
Follow-up (years)	2.6 (1.3-3.7)
Death	114 (40%)
Num. Treatment Courses, (κ)	
1	22 (8%)
2	36 (12%)
3	46 (16%)
4	188 (63%)
AML Risk Classification (high)	64 (23%)
WBC (cells/uL)	20 (6.8-65)
Male	152 (52%)

Time-varying confounders:

- Ejection fraction (EF_k) measured before each treatment decision.
- Presence of bloodstream infection.
- Waiting time since previous treatment.

A look at the data

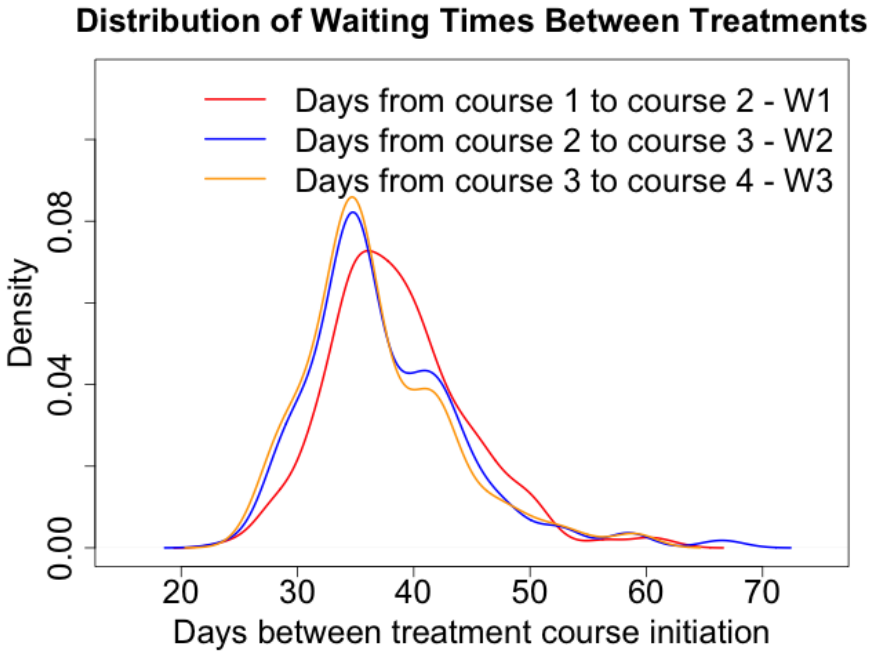

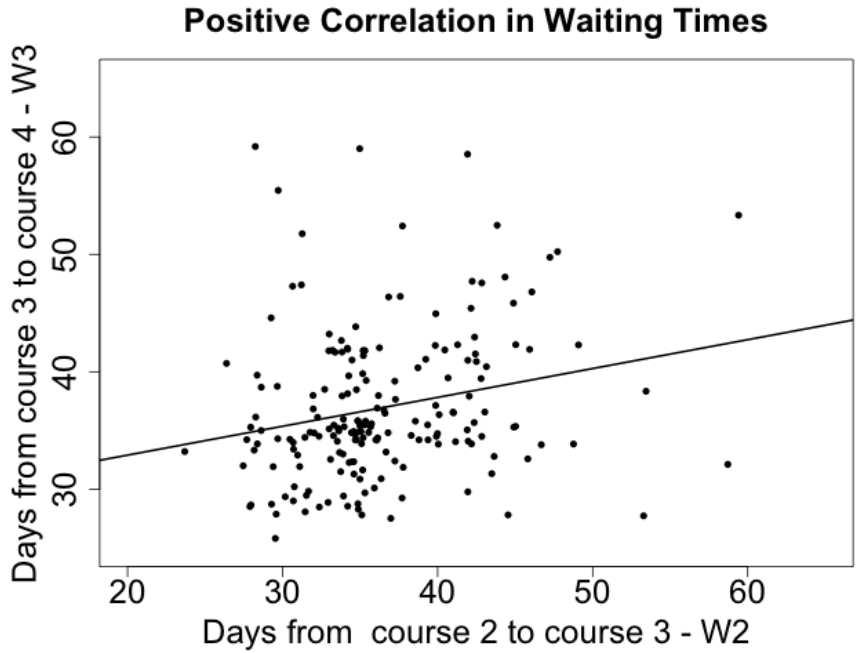

Ejection-Fraction Based Treatment Rules for AML

Withhold ACT at $k = 1$ if $EF_k < \tau_2$. Withhold ACT at $k = 2, 4$ if ejection fraction declined by $\tau_1\%$ from baseline to a value $< \tau_2$. Always withhold at course $k = 3$.

$$A_k = r(EF_k, EF_1; \tau) = \begin{cases} 1 - I(EF_1 < \tau_2) & k = 1 \\ 1 - I(EF_k/EF_1 - 1 < \tau_1) I(EF_k < \tau_2) & k \in 2, 4 \\ 0 & k = 3 \end{cases}$$

$$\tau = (\tau_1, \tau_2) \in \mathcal{T} = \{0, -.1, -.2, -.3, -.4, -.5\} \times \{.4, .5, .6, .7, .8, .9\}.$$

Note: $\mathcal{S}_3 = \{0\}$.

Posterior Survival Curves

Compare rule r with $(\tau_1 = -.1, \tau_2 = .5)$ and rule r' with $(\tau_1' = -.1, \tau_2' = 1)$.

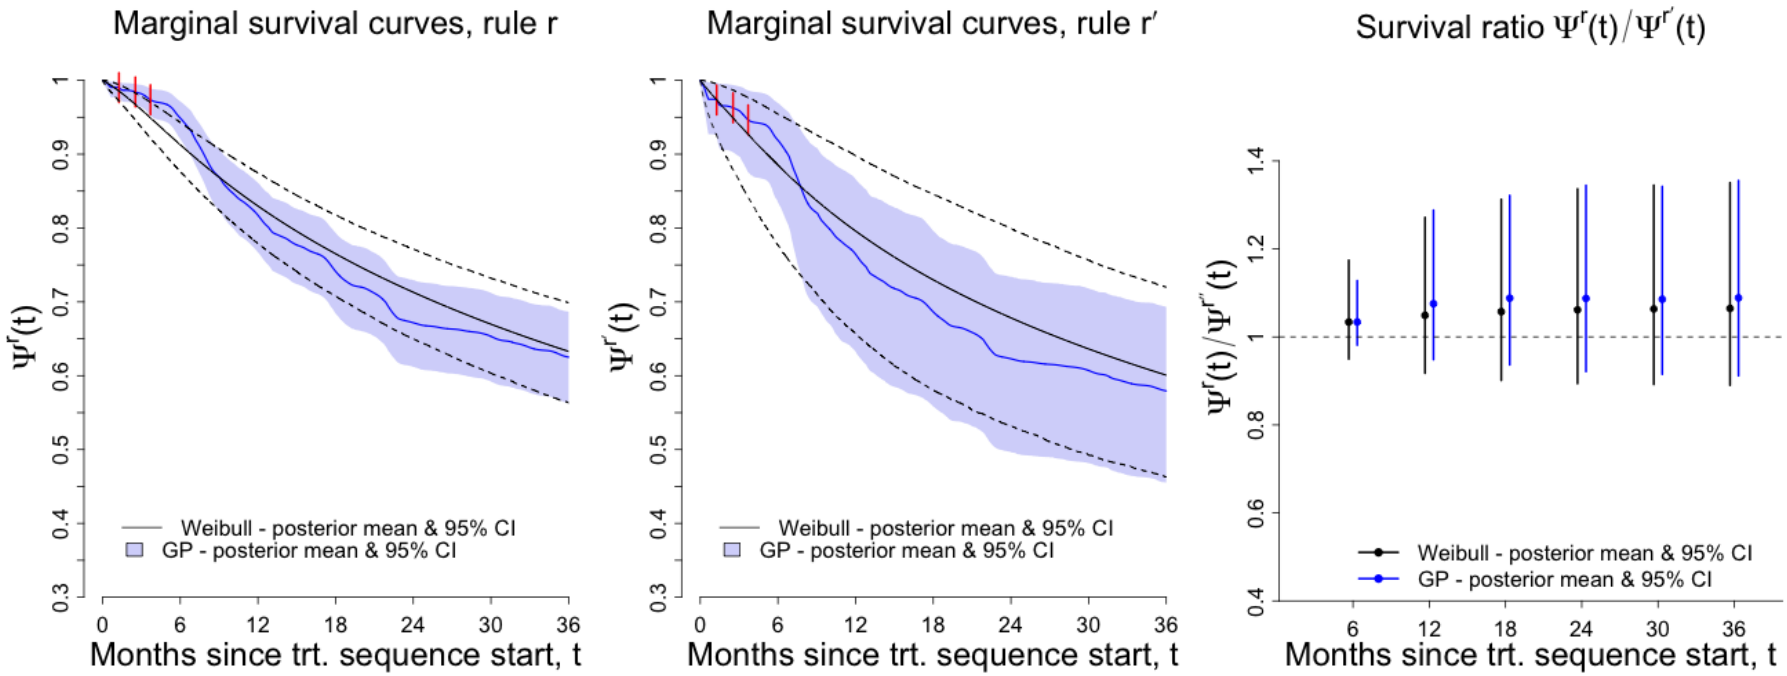

Thank you!

- Oganisian, A., Getz, K. D., Alonzo, T. A., Aplenc, R., & Roy, J. A. (2024). Bayesian semiparametric model for sequential treatment decisions with informative timing. *Biostatistics*. `doi.org/10.1093/biostatistics/kxad035`
- `causalBETA R` package: `https://arxiv.org/abs/2310.12358`
- Oganisian and Roy (2020). A practical introduction to Bayesian estimation of causal effects: parametric and nonparametric approaches. *Statistics in Medicine*.
- We're hiring a post-doc!
- Funding: PCORI ME2021C324942; PCORI MNJ3PA7MXFN6;
- Email: arman_oganisian@brown.edu

Scan for Paper

Appx: High Posterior Uncertainty for Infeasible Treatments

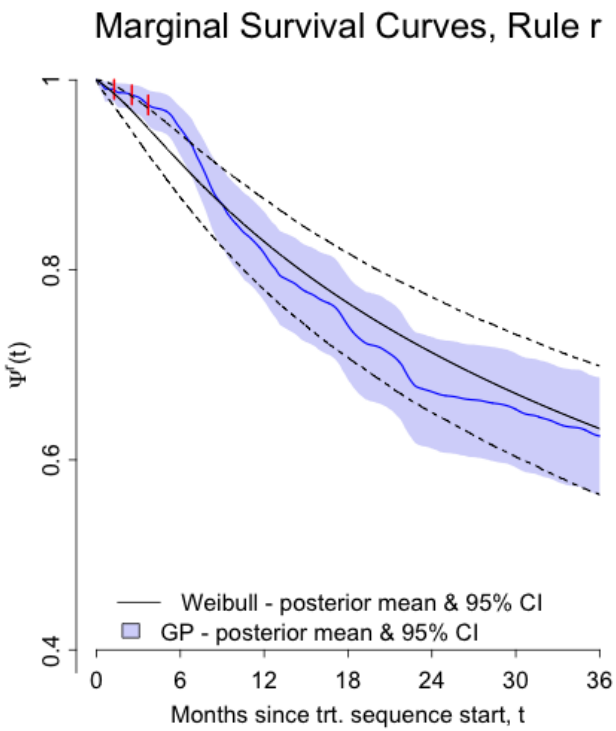

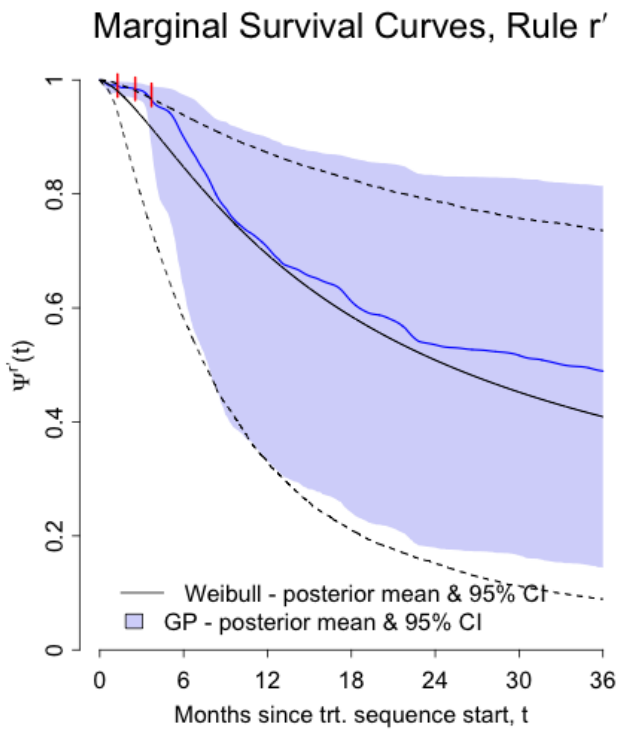

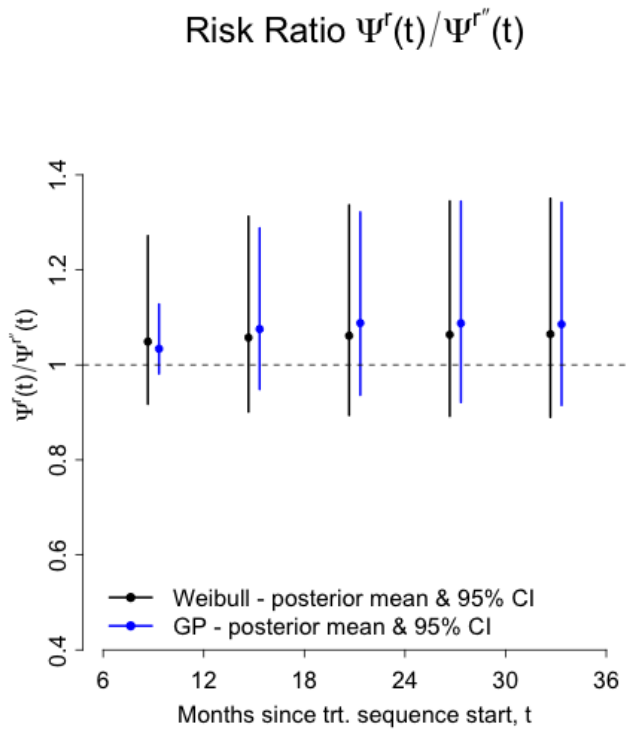

Appx: Posterior Over Optimal Rule Parameters

$$\tau^*(\omega^{(m)}) = \underset{\tau \in \mathcal{T}}{\text{argmax}}\, \Psi^{\tau}(t; \omega^{(m)})$$