Bayesian Nonparametric Methods for Causal Inference: Optimizing Sequential Decisions in Continuous Time

Jason Roy

Department of Biostatistics and Epidemiology
Rutgers School of Public Health
jason.roy@rutgers.edu

October 2022

Overview

Dynamic treatment strategies use past data to guide current treatment decisions.

One challenge in optimizing dynamic treatment strategies is time-dependent confounding.

- e.g. previous treatment decision affects variables L; L affects next treatment decision and outcome
- added complication: time between treatment decisions might itself be informative (e.g., patient needs time to recover)

Overview

We have to make strong causal identifying assumptions (e.g., sequential ignorability.

We do not have to make strong modeling assumptions when it comes to the joint or conditional distributions of the data.

Our overall approach is therefore:

- define causal estimands of interest
  - talk to subject matter experts to decide on treatment strategies to compare
- specify causal identifying assumptions
- use Bayesian nonparametric methods to model observed data distributions
- use g-computation to obtain posterior distribution of causal effects for each treatment strategy under consideration

Motivating Example

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

Patients move through a sequence of chemotherapy courses

- Anthracyclines (ACT) therapy is effective at suppressing cancer.
- ACT is cardiotoxic; known to lower ejection fraction (EF), which can decrease survival.
- Might need to withhold ACT if signs of poor or worsening heart function.

Motivating example

- Common practice involves withholding ACT if EF is below some threshold (e.g. a 10% decline from baseline (Neuendor et al., 2020)).
- Utility of these thresholds still debated
- In a previous trial, about 12% of patients experienced cardiotixicity within a year of treatment initiation.
- Patients who experienced cardiotoxicity were found to have lower overall survival

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

Motivating example

Standard protocol is 4 courses of treatment. However, timing of each course depends on patient data.

Notation:

- Y_k: time of kth treatment decision
- W_k: time between treatment decisions $W_k = Y_k - Y_{k-1}$
- T: survival time
- L_k: time-dependent covariates at time k
- A_k: treatment decision at time k (binary)

Data Structure

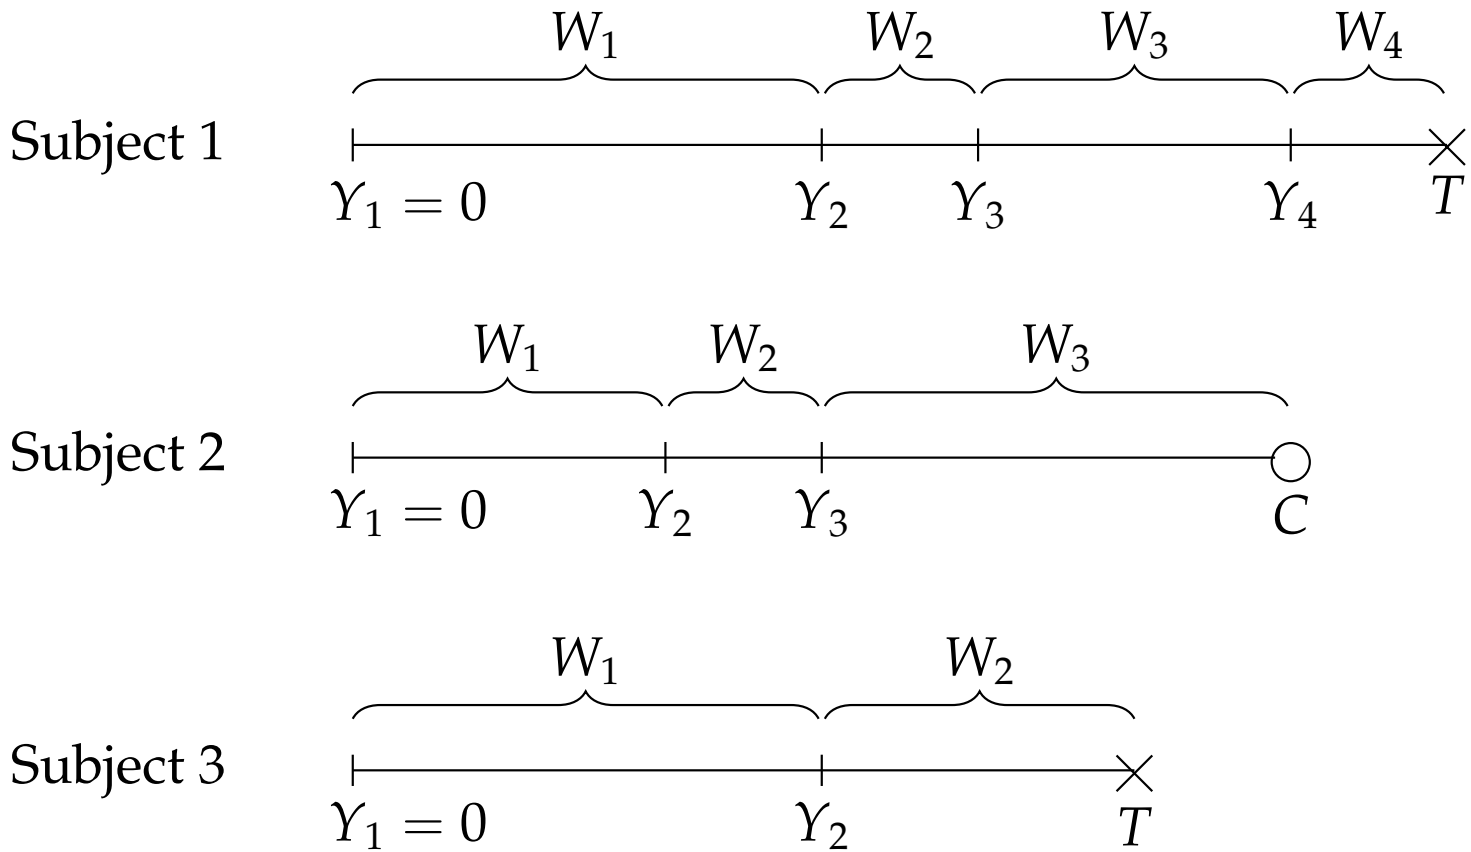

- $\overline{H}_k = (\overline{L}_k, \overline{A}_{k-1}, \overline{W}_{k-1})$ are used to make treatment decision $A_k \in \{1, 0\}$.

Waiting Times

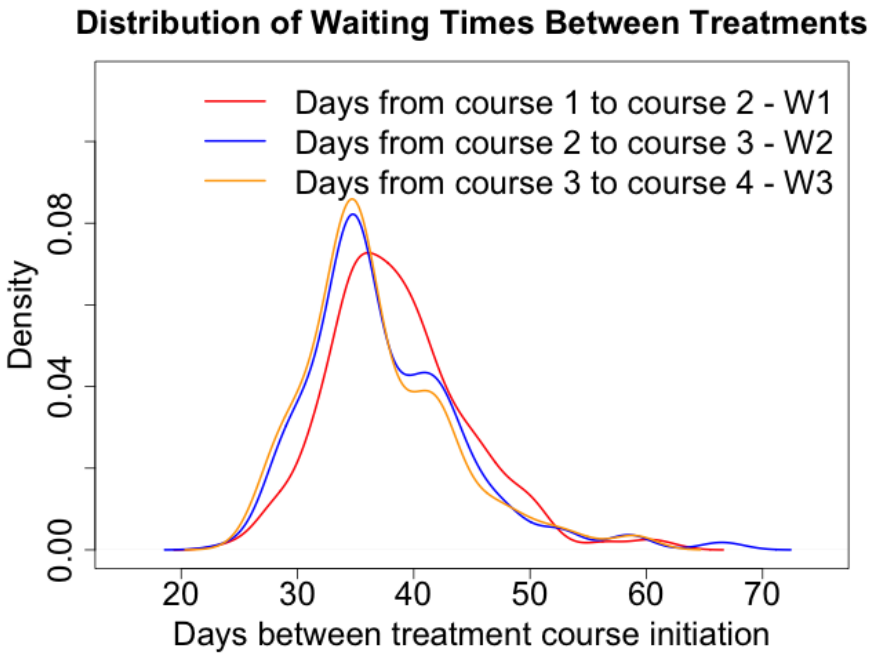

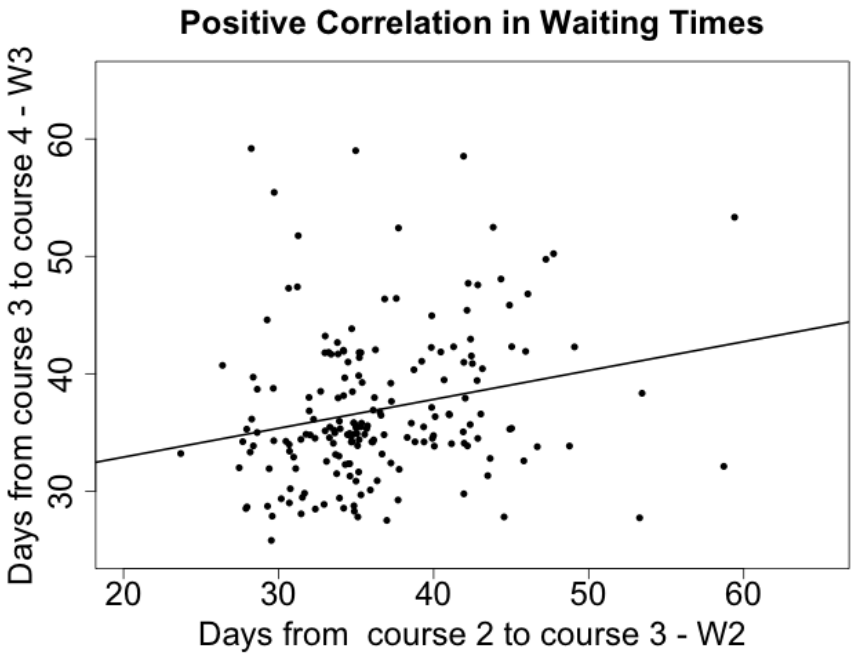

Potential Outcomes and Target Estimand

Define rule $r : \mathcal{H}_k \to \{0, 1\}$

- $T(r)$: potential survival time.
- $W_k(r)$: potential waiting times.
- $\kappa(r)$ potential number of treatments.

Target estimand:

$$\Psi^r(t) = P(T(r) > t)$$

- E.g. if T is measured in months, potential 2-year survival $\Psi^r(24)$.
- Contrast effects of rules: $\Psi^r(t)/\Psi^{r'}(t)$.
- For some t, find optimal rule

$$r^* = \arg\max_{r \in \mathcal{R}} \Psi^r(t)$$

Identification Assumptions

- Sequential Ignorability:

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

- Treatment Positivity:

$$P(A_k = 1 \mid \bar{H}_k = h_k, \kappa \geq k) > 0$$

for feasible $\bar{h}_k$

- Non-informative Censoring: $\lambda_C(t \mid H(t), \bar{W}_K(r)) = \lambda_C(t \mid H(t))$

$$\lambda_C(t \mid H(t)) = \lim_{dt \to 0} \frac{P(t \leq T \leq t + dt, \delta = 0 \mid T \geq t, H(t))}{dt}$$

Conditional on history, hazard does not depend on potential outcomes.

- Censoring Positivity:

$$P(C \geq t \mid H_k) > 0$$

A G-formula for $\Psi^r(t)$

Consider case $K = 2$. Then $\Psi^r(t) = P(T(r) > t)$

$$\begin{aligned}\Psi^r(t) = \int_t^\infty \int_{\bar{L}_2, W_1} & f_1(l_1) S_k(w_1 \mid a_{r,1}, l_1) \lambda_{Y1}(w_1 \mid a_{r,1}, l_1) \\ \times \; & f_2(l_2 \mid a_{r,1}, l_1, w_1, \kappa > 2) S_{T2}(u - w_1 \mid \bar{a}_{r,2}, \bar{l}_2, w_1, \kappa > 2) \\ \times \; & \lambda_{T2}(u - w_1 \mid \bar{a}_{r,K}, \bar{l}_2, w_1, \kappa > 2) \, d\bar{l}_2 \, dw_1 \, du\end{aligned}$$

Dynamic AML Treatment Rules

Ejection-fraction based treatment rules

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

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

Bayesian Proportional Hazard Models

For hazard of subsequent treatment after treatment k,

$$\lambda_{Yk}(w \mid H_k, \lambda_{0Yk}, \beta_{Yk}) = \lambda_{0Yk}(w) \exp\left(H_k' \beta_{Yk}\right)$$

Priors

$$\Lambda_{0Yk} \sim GP(\alpha_{Yk} \Lambda^*_{Yk})$$
$$\beta_{Yk} \sim f_{\beta_{Yk}}(\beta_{Yk})$$

Identical model for $\lambda_{Tk}(w \mid H_k, \lambda_{0Tk}, \beta_{Tk})$

Gamma Process

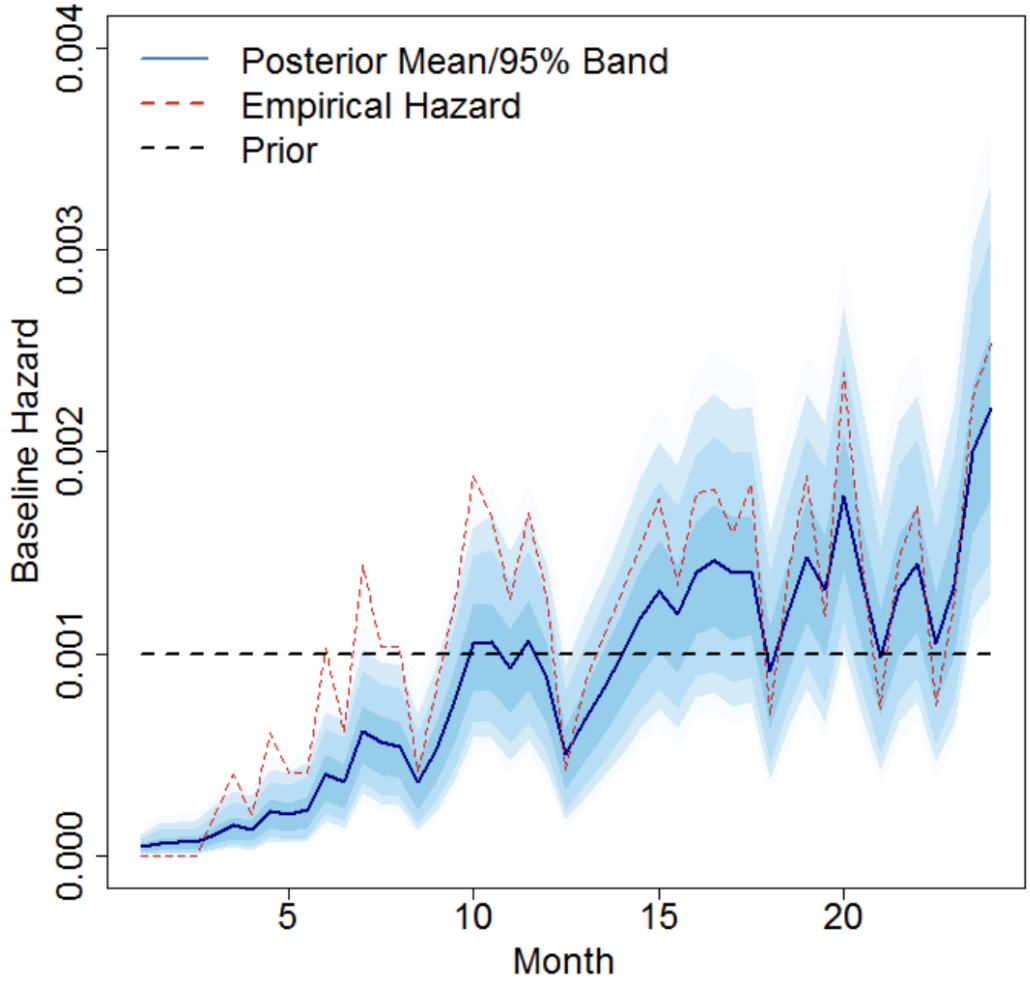

Confounder Trajectory Models

$$f_k(l_k \mid \bar{h}_k, \kappa \geq k) = \prod_{p=1}^{P} f_{kp}(l_{kp} \mid \bar{h}_k, \kappa > k)$$

- Could use your favorite BNP models for each of these
- Can be done in parallel
- e.g., for a binary covariate l_{kp} specify a probit BART model

Posterior Sampling

Want to sample from

$$f(\omega_1, \omega_2, \ldots, \omega_K \mid \mathcal{D})$$

where $\omega_k = \left\{ \lambda_{0Tk}, \beta_{Tk}, \lambda_{0Yk}, \beta_{Yk}, \eta_k \right\}$.

Can do this sequentially - Metropolis-in-Gibbs.

G-Computation for $K = 2$

At end of iteration m, we have for $k = 1, 2$

$$\omega_k^{(m)} = \left\{ \lambda_{0Tk}^{(m)}, \beta_{Tk}^{(m)}, \lambda_{0Yk}^{(m)}, \beta_{Yk}^{(m)}, \eta_k^{(m)} \right\}$$

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

1. Simulate $L_1^{(b)} \sim f_1(L_1 \mid \eta_1^{(m)})$
2. Simulate $W_{Y1}^{(b)} \sim \lambda_{Y1}(w \mid L_1^{(b)}, A_1 = r(L_1^{(b)}), \lambda_{0Y1}^{(m)}, \beta_{Y1}^{(m)})$
3. Simulate $W_{T1}^{(b)} \sim \lambda_{T1}(w \mid L_1^{(b)}, A_1 = r(L_1^{(b)}), \lambda_{0T1}^{(m)}, \beta_{T1}^{(m)})$
4. if $W_{T1}^{(b)} < W_{Y1}^{(b)}$, then set $T^{(b)} = W_{T1}^{(b)}$. Else
5. Simulate $L_2^{(b)} \sim f_2(L_2 \mid L_1^{(b)}, A_1 = r(L_1^{(b)}), W_{Y1}^{(b)}, \eta_2^{(m)})$
6. Simulate $W_{T2}^{(b)} \sim \lambda_{T2}(w \mid \bar{L}_2^{(b)}, \bar{A}_2 = r(\bar{L}_2^{(b)}), W_{Y1}^{(b)}, \lambda_{0T2}^{(m)}, \beta_{T2}^{(m)})$
   Set $T^{(b)} = W_{Y1}^{(b)} + W_{T2}^{(b)}$]]

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

Simulation

Weibull model is true model. The proportion of patients still at risk for subsequent treatment or death at time points 5, 10, 15, and 20 are about 95%, 85%, 60%, and 2%, respectively

95% credible interval.

		Gamma Process			Weibull		
Setting	t	Bias (%)	Coverage	Int. Width	Bias (%)	Coverage	Int. Width
N=1000	5	0.07	95.4	0.04	0.14	94.5	0.04
	10	0.03	95.2	0.07	0.22	94.8	0.07
	15	2.29	90.4	0.08	0.24	95.8	0.08
	20	185.00	18.8	0.13	6.58	94.5	0.06
N=300	5	0.34	94.9	0.07	0.53	95.8	0.07
	10	0.13	94.7	0.12	0.98	95.2	0.12
	15	4.52	89.3	0.14	0.51	94.9	0.14
	20	227.00	14.7	0.18	25.00	90.3	0.11

Posterior Rule Optimization

Consider rules parameterized by

$$\mathcal{R} = \{r_\tau : \tau \in \mathcal{T}\}$$

E.g. $A_k = r(L_k) = I(L_k > \tau)$. Then a posterior draw of the optimal rule is,

$$\tau^*(\omega) = \arg\max_{\tau \in \mathcal{T}} \Psi^\tau(t; \omega)$$

Can be done for each posterior draw $\omega^{(m)}$. Grid search is feasible for low dimensional, finite $\mathcal{T}$.

Capturing Risk-Reward via Utility Functions

Could consider estimands that capture risk-benefit tradeoff. For fixed s, t,

$$U^r(s,t) = P(T^r > t) - P(L_2^r < s)$$

A rule with same survival probability as another, but lower adverse event risk would be preferred.

Analysis of Dynamic AML Treatment Rules

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 ahead of each treatment.
- Presence of bloodstream infection.
- Waiting time since previous treatment.

Dynamic AML Treatment Rules

Ejection-fraction based treatment rules

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

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

Posterior Survival Curves

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

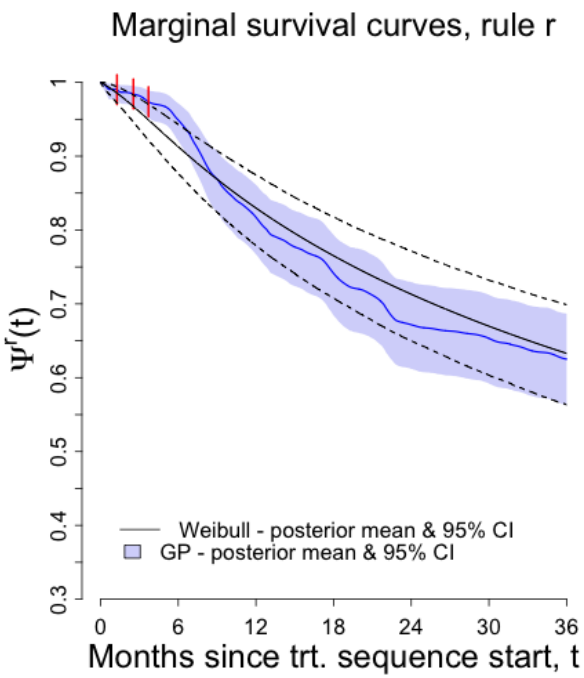

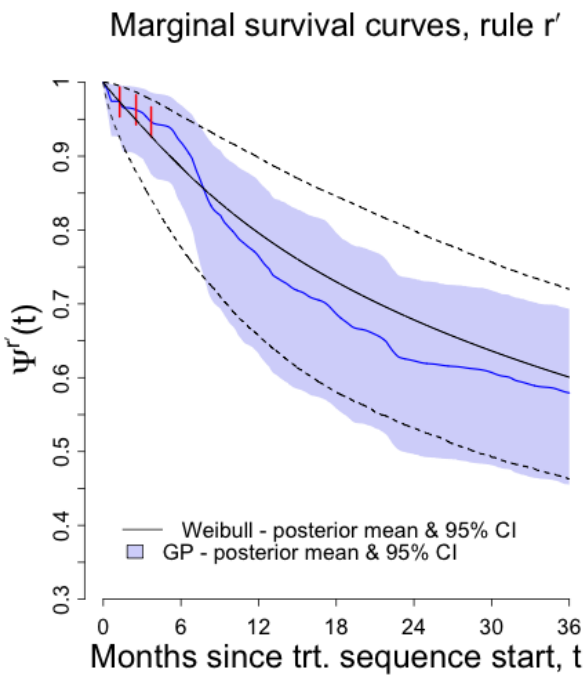

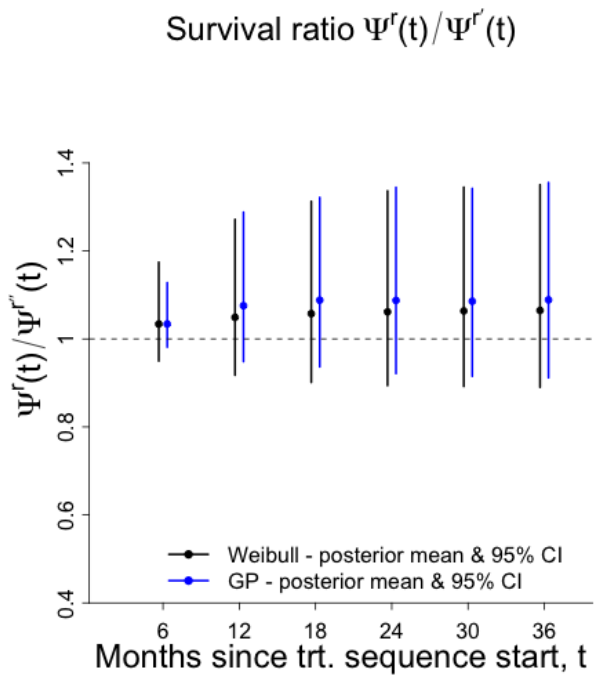

Posterior Over Optimal Rule Parameters

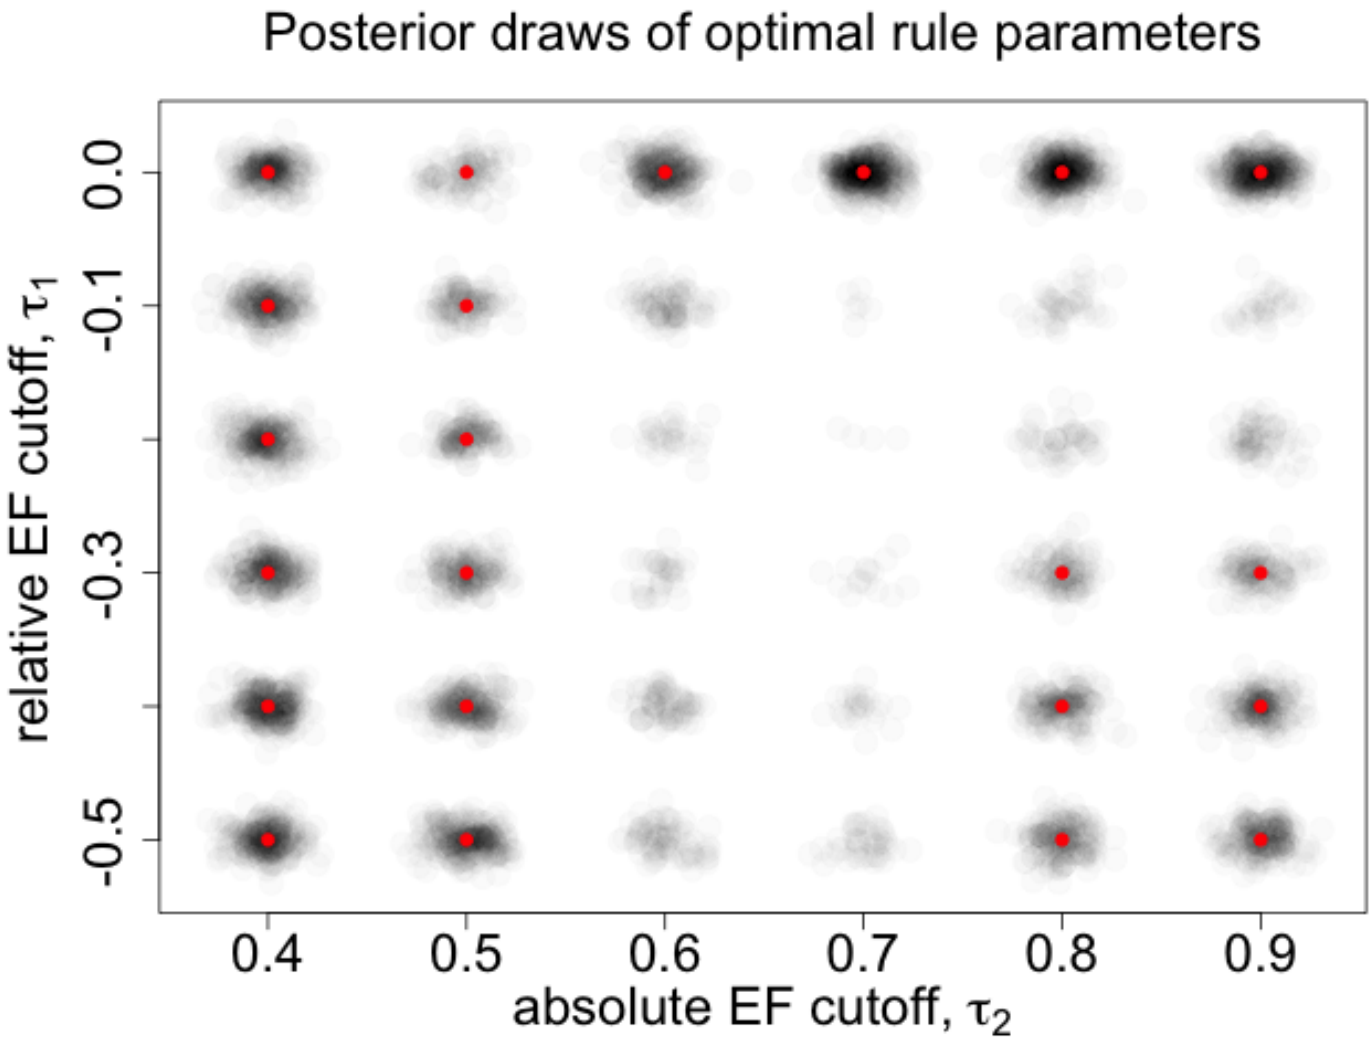

Capturing trade-offs

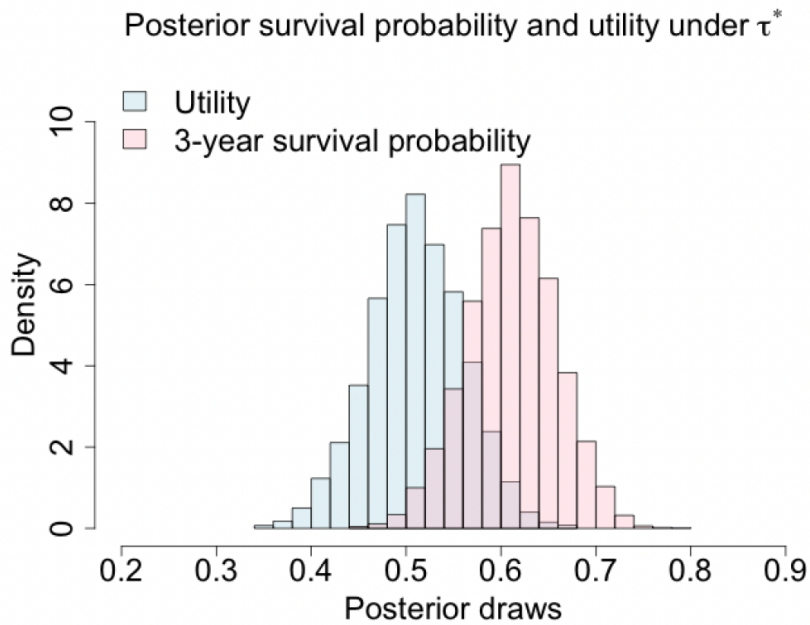

Under this rule, posterior mean 3-year survival is about 60% with about 10% risk of EF dipping below 0.5 during the treatment course.

Discussion

- Developed a BNP approach to optimal dynamic treatment strategies that accounts for waiting times
- An advantage of BNP approach is full posterior for the optimal strategy
- Important to incorporate subject matter knowledge when determining range of strategies
- Future work: develop sensitivity analysis methods for the causal assumptions
- Future work: determine optimal measurement time strategy (how often should a biomarker be measured) and deal with missing data

Acknowledgements

- Co-authors: Arman Oganisian (Brown U), Kelly Getz (UPenn), Todd Alonzo (USC), Richard Aplenc (CHOP)
- Data: AAML1031 phase III clinical trial run by the Children's Oncology Group (COG)
- Funding: PCORI

Questions?