Utility Based Bayesian Personalized Treatment Selection for Advanced Breast Cancer

Juhee Lee *

Department of Statistics, University of California Santa Cruz, Santa Cruz, CA

Peter F. Thall

Department of Biostatistics, UT M.D. Anderson Cancer Center, Houston, TX

Bora Lim

Department of Hematology & Oncology, Breast Oncology,

Dan L Duncan Comprehensive Cancer Center, Baylor College of Medicine, Houston, TX

Pavlos Msaouel

Departments of Genitourinary Medical Oncology and Translational Molecular Pathology,

UT M.D. Anderson Cancer Center, Houston, TX

Abstract

A Bayesian method is proposed for personalized treatment selection in settings where data are available from a randomized clinical trial with two or more outcomes. The motivating application is a randomized trial that compared letrozole plus bevacizimab to

*Address for Correspondence: Department of Statistics, Baskin School of Engineering, University of California 1156 High Street Mail Stop SOE2 Santa Cruz, CA 95064 USA. E-mail: juheelee@soe.ucsc.edu.

letrozole alone as first-line therapy for hormone receptor positive advanced breast cancer. The combination treatment arm had larger median progression-free survival time, but also a higher rate of severe toxicities. This suggests that the risk-benefit trade-off between these two outcomes should play a central role in selecting each patient's treatment, particularly since older patients are less likely to tolerate severe toxicities. To quantify the desirability of each possible outcome combination for an individual patient, we elicited from breast cancer oncologists a utility function that varied with age. The utility was used as an explicit criterion for quantifying risk-benefit trade-offs when making personalized treatment selections. A Bayesian nonparametric multivariate regression model with a dependent Dirichlet process prior was fit to the trial data. Under the fitted model, a new patient's treatment can be selected based on the posterior predictive utility distribution. For the breast cancer trial dataset, the optimal treatment depends on the patient's age, with the combination preferable for patients 70 years or younger and the single agent preferable for patients older than 70.

Keywords: Bayesian nonparametrics; Dependent Dirichlet process; Multivariate probit regression; Precision medicine; Statistical decision making; Utility function.

1 Introduction

Nearly all published clinical trial results focus on statistical inferences about effects of treatments and patient prognostic variables on clinical outcomes. This may fall short of what is needed by practicing physicians to make informed treatment decisions for individual patients. In many settings, estimated effects on efficacy and toxicity lead to conflicting treatment choices, and the relative desirability of two treatments also may vary with patient prognostic variables. Our motivating dataset, which illustrates this class of problems, arose from a phase III study of targeted agents for treating hormone receptor-positive advanced breast cancer, reported by Dickler et al. (2016). Patients were randomized between letrozole

plus bevacizumab (L+B) and letrozole plus placebo (L). The primary efficacy endpoint was progression-free survival (PFS) time, defined as the time from the treatment to disease progression or death from any cause. Due to safety concerns, 21 different types of toxicity were monitored, including the type and grade (0=none to 5=fatal) of each occurrence. A statistically significant PFS improvement was seen with L+B compared to L (one-sided p-value = 0.016), with estimated median PFS 20.2 months (95% confidence interval, CI, 17.0 – 24.1) with L+B compared to 15.6 months (95% CI 12.9 – 19.7) with L. Consideration of toxicities led to the opposite conclusion, with 46.8% of patients treated with L+B experiencing severe (grade ≥ 3) toxicities compared to 14.2% with L.

Considering each outcome alone, selecting an optimal treatment is straightforward, since longer PFS and less toxicity each is more desirable. This leads to the problematic conclusions that L+B is preferable in terms of PFS but L is preferable in terms of toxicity. Thus, when considering these two outcomes together, as must be done in practice by a physician when choosing between the treatments for an individual patient, decision making is not straightforward. In our analyses, rather than dichotomizing toxicity severity, we will use total toxicity burden (TTB) (Bekele and Thall, 2004; Le-Rademacher et al., 2020) to summarize each patient's adverse events. In general, for K toxicities $\boldsymbol{z} = (z_1, \cdots, z_K)$, where each z_k is a grade in $\{0, \ldots, J-1\}$, we define the scaled TTB to be $q = \sum_{k=1}^{K} z_k / \{K \times (J-1)\}$, which takes on values between 0 and 1. Fig 1 illustrates the TTB distributions and Kaplan-Meier estimates of the PFS survival function for each treatment arm in the breast cancer dataset, with L+B represented by blue and L by red. The questions that we will address in this paper are how one may use the available data to choose between L+B and L for a new breast cancer patient, and how this may be done in other, similar settings.

We assume that each patient's data can be summarized as a vector, $\boldsymbol{y}$, of clinical outcomes, a vector, $\boldsymbol{x}$, of prognostic covariates, and a treatment indicator variable, τ. A Bayesian regression model, $f(\boldsymbol{y} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta})$ is assumed and fit to the data, where $\boldsymbol{\theta}$ denotes the model's

parameter vector. For the breast cancer trial data, $\boldsymbol{y} = (y_1, y_2)$ with y_1 = progression-free survival (PFS) time and y_2 = TTB. Using the breast cancer trial data for illustration, we extend the usual statistical process of data analysis by connecting it with medical decision making by practicing physicians. To do this, we first construct a family of utility functions, with each utility assigning numerical desirability scores $U(\boldsymbol{y}, \boldsymbol{x})$ to all $\boldsymbol{y}$ = (PFS time, TTB) pairs for a patient with prognostic variables $\boldsymbol{x}$. In a given setting, a physician and patient may choose a particular utility function from the family that best represents the patient's subjective trade-offs. Because the trade-off between PFS and TTB may vary with $\boldsymbol{x}$, the desirability of a particular pair of treatment options may not be the same for all patients, and this may lead to different treatment preferences for two patients having different $\boldsymbol{x}$.

Decision analysis based on utility functions certainly is not new. This has been studied and applied in many areas, including business (e.g., Pennings and Smidts (2003); Loewenstein et al. (1989)), engineering (e.g., Chen et al. (1998); Bagočius et al. (2014)), and operations research (e.g., Walsh et al. (2004); Roy et al. (2017)). Two papers of a five-part primer on medical decision analysis are given by Detsky et al. (1997) and Naglie et al. (1997). However, formal utility-based decision procedures are seldom included in statistical data analysis reports. Our application of the methodology to the breast cancer dataset illustrates how a utility function and Bayesian statistical model can be used to choose between two treatments for a patient with given prognostic variables. To establish the idea that outcome utilities can be used as practical tools for decision making, and illustrate the range of potential applications, Supplementary § 2 provides examples of utility functions for different types of outcome vectors. These include a one-dimensional ordinal outcome, binary (response, toxicity) indicators, and the two ordinal categorical outcomes (disease status, toxicity severity).

For the breast cancer data analysis, we develop a robust Bayesian regression model, $f(\boldsymbol{y} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta})$, that assumes latent patient frailties to account for association among the elements of $\boldsymbol{y}$ and $\boldsymbol{z}$, and describes how each outcome varies as a function of treatment, τ,

and baseline covariates, $\boldsymbol{x}$. We formulate a joint Bayesian nonparametric (BNP) multivariate regression model that includes a vector of continuous latent variables defined to represent ordinal toxicity outcomes, using the dependent Dirichlet process (DDP) (MacEachern, 1999). We use a linear DDP developed by De Iorio et al. (2004, 2009). DDP models are highly flexible and provide a robust basis for inferences about regression relationships. BNP models have been applied to a broad range of statistical problems, including density estimation, regression, clustering, and survival analysis. See, for example, Müller et al. (2015) for general applications, Mitra and Müller (2015) for applications in biostatistics, or Müller and Mitra (2013); Thall et al. (2017) for overviews and illustrations. Given a Bayesian model $f(\boldsymbol{y} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta})$ and utility function $U(\boldsymbol{y}, \boldsymbol{x})$, we use posterior predictive utility distributions as a basis for deciding between treatments for a new patient with prognostic variables $\boldsymbol{x}$. We choose the treatment that yields the greatest posterior mean utility.

The remainder of the paper is organized as follows. § 2.1 formulates a bivariate regression model for clinical outcomes PFS time and TTB, and provides a predictive distribution of (PFS, TTB) for a future patient as a function of the patient's covariates and each potential treatment that may be given to the patient. § 2.2 provides computational details for implementation. In § 3, we describe a utility function that varies with PFS, TTB, and covariates in order to represent a personalized risk-benefit trade-off between PFS and TTB. The utility function is constructed using separate contributions from PFS time and TTB, with each contribution constrained so that it is logically consistent and reflects elicited expert opinion. The predictive distribution allows one to estimate the utility of each treatment for a future patient with given covariates, and provides a way to compute the probability that each treatment is preferred for the future patient. In § 4, we illustrate the methodology by applying it to make personalized treatment selections based on the breast cancer trial data. In § 5, a simulation study is presented to illustrate general properties of the proposed decision making approach. We close with a brief discussion in § 6.

2 A Bayesian Nonparametric Regression Model

2.1 Sampling Distribution and Prior Specification

Let $t_i \in \mathbb{R}^+$ denote PFS time during the follow-up period $(0,\ c_i]$, for patient $i = 1, \ldots, n$. The observed time of failure (progression or death) or independent administrative censoring at c_i is $t_i^o = \min(t_i, c_i)$, with $\delta_i = 1$ if PFS time was observed, $t_i \leq c_i$, and $\delta_i = 0$ if censored, $t_i > c_i$. Denote the ordinal variable $z_{i,k} \in \{0, 1, \ldots, J-1\}$ for the maximum grade that patient i experienced of toxicity type $k = 1, \ldots, K$. Censoring is assumed to be independent of t_i, toxicity occurrences, and covariates. Denote the i^{th} patient's vector of baseline covariates by $\boldsymbol{x}_i = (x_{i,1}, \ldots, x_{i,P})$, and the observed data by $\mathcal{D} = \{(t_i^o, \delta_i, \boldsymbol{z}_i, \tau_i, \boldsymbol{x}_i), i = 1, \ldots, n\}$, where $\boldsymbol{z}_i = (z_{i,1}, \ldots, z_{i,K})$. In the breast cancer dataset, there are $n = 340$ patients after removing three patients having $t^o = 0$, $K = 21$ toxicity categories, and $J = 6$ severity grades, where grade 0 = no occurrence of that toxicity type and 5 = death. If a subject died due to a toxicity type k occurrence, the corresponding $z_{i,k} = 5$ was recorded with observed survival time $t_i^o = t_i$ and $\delta_i = 1$. We include $P = 3$ prognostic covariates, $x_1 =$ age, an indicator x_2 of measurable disease at baseline, and an indicator x_3 of whether the patient's disease free interval prior to trial entry was greater than 24 months, in addition to the indicator τ of treatment $L + B$.

To construct a model that accounts for heterogeneity between patients not explained by the covariates, we introduce real-valued $(K+1)-$dimensional multivariate normal latent frailty vectors, $\boldsymbol{s}_i = (s_{i,0}, s_{i,1}, \ldots, s_{i,K})'$ for $i = 1, \cdots, n$. We assume $\boldsymbol{s}_i \mid \Omega \overset{iid}{\sim} \mathrm{N}_{K+1}(\boldsymbol{0}, \Omega)$ and $\Omega \sim$ Inv-Wishart(a_Ω, Ω_0). Following Chib and Greenberg (1998), we construct a multinomial probit model for the ordinal toxicity outcomes $\boldsymbol{z}_i$ by introducing the unobserved real-valued latent variables $\tilde{\boldsymbol{z}}_i = (\tilde{z}_{i,1}, \cdots, \tilde{z}_{i,K})$, where $\tilde{z}_{i,k} \in \mathbb{R}$, and define $z_{i,k} = j$ if and only if $u_{k,j} < \tilde{z}_{i,k} \leq u_{k,j+1}$, where $u_{k,0} < u_{k,1} < \cdots < u_{k,J}$ denote toxicity type-specific cutoffs for each k. This is a common modeling strategy that uses real-valued latent variables to induce

a tractable multivariate distribution on a vector of observed ordinal categorical variables, and greatly facilitates computation. For logarithm transformed PFS time, $\tilde{t} = \log(t)$, latent variables $\tilde{\boldsymbol{z}}_i$, treatment τ_i, and covariates $\boldsymbol{x}_i$, we assume

$$(\tilde{t}_i, \tilde{\boldsymbol{z}}_i) \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i \overset{indep}{\sim} h(\tilde{t}_i, \tilde{\boldsymbol{z}}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i), \text{ where } (\tilde{t}_i, \tilde{\boldsymbol{z}}_i) \in \mathbb{R}^{K+1}. \tag{1}$$

We take a BNP approach for modeling h in (1) that allows flexible regression structures by assuming the dependent Dirichlet process (DDP) (MacEachern, 1999), which is a family of random probability distributions indexed by $(\tau, \boldsymbol{x})$. Specifically, we use a linear-DDP that induces covariate dependence through linear regression structures (De Iorio et al., 2004, 2009). Denoting $\widetilde{\boldsymbol{x}}_i' = (1, \tau_i, \boldsymbol{x}_i')$, $\boldsymbol{\beta} = (\beta_0, \beta_\tau, \beta_1, \ldots, \beta_P)'$ and $\boldsymbol{\alpha}_k = (\alpha_{k,0}, \alpha_{k,\tau}, \alpha_{k,1}, \ldots, \alpha_{k,P})'$, we assume the simple parametric linear combinations $\eta_0(\widetilde{\boldsymbol{x}}_i) = \boldsymbol{\beta}'\widetilde{\boldsymbol{x}}_i$ and $\eta_k(\widetilde{\boldsymbol{x}}_i) = \boldsymbol{\alpha}_k'\widetilde{\boldsymbol{x}}_i$, for each $k = 1, \ldots, K$. We denote $\boldsymbol{\eta}(\widetilde{\boldsymbol{x}}_i) = (\eta_0(\widetilde{\boldsymbol{x}}_i), \ldots, \eta_K(\widetilde{\boldsymbol{x}}_i))$ and construct a model for h through a convolution with a normal kernel

$$h(\tilde{t}_i, \tilde{\boldsymbol{z}}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i) = \int \phi_{K+1}(\tilde{t}_i, \tilde{\boldsymbol{z}}_i \mid \boldsymbol{\eta}(\widetilde{\boldsymbol{x}}_i) + \boldsymbol{s}_i, \Sigma) dG(\boldsymbol{\beta}, \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_K), \tag{2}$$

where $\phi_d(\cdot \mid \boldsymbol{a}, \boldsymbol{B})$ is the density function of the $d-$variate normal distribution with mean vector $\boldsymbol{a}$ and $d \times d$ covariance matrix $\boldsymbol{B} > 0$. We use the Dirichlet process (DP) as a prior for the random mixing distribution G in (2). This gives a DP mixture of multivariate normal linear models,

$$h(\tilde{t}_i, \tilde{\boldsymbol{z}}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i) = \sum_{m=1}^{\infty} w_m \phi_{K+1}(\tilde{t}_i, \tilde{\boldsymbol{z}}_i \mid \boldsymbol{\eta}_m(\widetilde{\boldsymbol{x}}_i) + \boldsymbol{s}_i, \Sigma), \tag{3}$$

where $\boldsymbol{\eta}_m(\widetilde{\boldsymbol{x}}_i) = (\eta_{m,0}(\widetilde{\boldsymbol{x}}_i), \ldots, \eta_{m,K}(\widetilde{\boldsymbol{x}}_i))$ with $\eta_{m,0}(\widetilde{\boldsymbol{x}}_i) = \boldsymbol{\beta}_m'\tilde{\boldsymbol{x}}_i$ and $\eta_{m,k}(\widetilde{\boldsymbol{x}}_i) = \boldsymbol{\alpha}_{m,k}'\tilde{\boldsymbol{x}}_i$. The weights $\{w_m\}$ are constructed via Sethuraman's (1994) so-called "stick-breaking" process by assuming $w_m / \prod_{m'=1}^{m-1}(1 - w_{m'}) \overset{iid}{\sim} \text{Be}(1, \xi)$, with fixed $\xi > 0$. For the covariate and treatment

effect parameters in (3), we assume

$$\boldsymbol{\beta}_m \overset{iid}{\sim} \mathrm{MVN}_{P+2}(\bar{\boldsymbol{\beta}}, \kappa^2 \mathrm{I}_{P+2}) \quad \text{and} \quad \boldsymbol{\alpha}_{m,k} \mid \bar{\boldsymbol{\alpha}}_k, \boldsymbol{V} \overset{iid}{\sim} \mathrm{MVN}_{P+2}(\bar{\boldsymbol{\alpha}}_k, \boldsymbol{V}) \tag{4}$$

with $\bar{\boldsymbol{\beta}}$ and κ^2 fixed, where MVN_{P+2} represents a $(P+2)-$dimensional multivariate normal distribution. In (4), $\bar{\boldsymbol{\beta}} = (\bar{\beta}_0, \bar{\beta}_\tau, \bar{\beta}_1, \ldots, \bar{\beta}_P)$ and $\bar{\boldsymbol{\alpha}}_k = (\bar{\alpha}_{k,0}, \bar{\alpha}_{k,\tau}, \bar{\alpha}_{k,1}, \ldots, \bar{\alpha}_{k,P})$ are $(P+2)-$dimensional mean vectors and $\boldsymbol{V} = \mathrm{diag}[v_p^2]$ is a $(P+2) \times (P+2)$ matrix. We assume $\bar{\alpha}_{k,p} \overset{indep}{\sim} \mathrm{N}(\bar{\bar{\alpha}}_p, v_\alpha^2)$ with fixed $\bar{\bar{\alpha}}_p$ and v_α^2, and $v_p^2 \overset{iid}{\sim} \mathrm{IG}(a_v, b_v)$ with fixed a_v and b_v for $p = 0, \ldots, P$. The hierarchical structure for $\boldsymbol{\alpha}_{m,k}$ enables the model to borrow information across the toxicity categories. The model in (3) incorporates τ and $\boldsymbol{x}$ linearly in the mean of each normal summand. Due to the fact that the distribution of $(\tilde{t}, \tilde{\boldsymbol{z}})$ is a weighted average of multivariate normal distributions, each with its own linear term, the model accounts for possible effects of τ and $\boldsymbol{x}$ on $(\tilde{t}, \tilde{\boldsymbol{z}})$ that can be nonlinear and quite complex, including interactions between two or more variables in $(\tau, \boldsymbol{x})$. This construction thus provides a flexible modeling framework for inference and prediction, avoiding restrictive assumptions of linearity or additivity in the covariate effects. This facilitates accurate decision making by avoiding restrictions imposed by conventional parametric models, such as the proportional hazards model. We let $\Sigma = \mathrm{diag}(\sigma_t^2, \sigma_z^2, \ldots, \sigma_z^2)$, which implies conditional independence between $\tilde{t}_i$ and $\tilde{z}_{i,k}$ given $\boldsymbol{s}_i$. Due to the conditional independence, the marginal distributions from (3) can be expressed as the weighted averages

$$\begin{aligned} \tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0} &\overset{indep}{\sim} f(\tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0}) = \sum_{m=1}^{\infty} w_m \phi_1(\tilde{t}_i \mid \eta_{m,0}(\widetilde{\boldsymbol{x}}_i) + s_{i,0}, \sigma_t^2), \\ \tilde{z}_{i,k} \mid \tau_i, \boldsymbol{x}_i, s_{i,k} &\overset{indep}{\sim} g_k(\tilde{z}_{i,k} \mid \tau_i, \boldsymbol{x}_i, s_{i,k}) = \sum_{m=1}^{\infty} w_m \phi_1(\tilde{z}_{i,k} \mid \eta_{m,k}(\widetilde{\boldsymbol{x}}_i) + s_{i,k}, \sigma_z^2) \end{aligned} \tag{5}$$

That is, f and each g_k also is a linear DDP mixture distribution. The marginal distribution

of each $z_{i,k}$ is obtained by integrating over the latent variables,

$$\mathrm{P}(z_{i,k} = j \mid \tau_i, \boldsymbol{x}_i, s_{i,k}) = \sum_{m=1}^{\infty} w_m \int_{u_{k,j}}^{u_{k,j+1}} \phi_1(\tilde{z} \mid \eta_{m,k}(\widetilde{\boldsymbol{x}}_i) + s_{i,k}, \sigma_z^2) d\tilde{z}. \tag{6}$$

Marginalizing by averaging over $\boldsymbol{s}$ in (3), the resulting DP mixture model, $h(\tilde{t}, \tilde{\boldsymbol{z}} \mid \tau, \boldsymbol{x})$, has covariance matrix $\Sigma + \Omega$. Thus, Ω induces dependence between PFS time and the toxicities within each patient, in addition to explaining additional variability between patients not explained by τ and $\boldsymbol{x}$.

To ensure identifiability in the multivariate ordinal regression model, we fix σ_z^2 and set $u_{k,1} = 0$ for all k. We also set $u_{k,0} = -\infty$, and $u_{k,J} = \infty$, and $\mathrm{P}(z_{i,k} < 0) = 0$ and $\mathrm{P}(z_{i,k} \leq J-1) = 1$. We let the cut-offs $u_{k,j}$, $j = 2, \ldots, J-1$ be random for flexibility, by defining $u_{k,j} = u_{k,j-1} + e_{k,j-1}$ $k = 1, \ldots, K$ and with error terms $e_{k,j} \overset{iid}{\sim} \mathrm{Ga}(a_e, b_e)$, for $j = 2, \ldots, J-1$. Lastly, we assume $\sigma_t^2 \sim \mathrm{IG}(a_t, b_t)$.

2.2 Posterior Inference

Collecting terms, $\boldsymbol{\theta} = (w_m, \boldsymbol{\beta}_m, \sigma_t^2, \boldsymbol{\alpha}_{m,k}, \boldsymbol{e}, \bar{\alpha}_{k,p}, v_p^2, \Omega)$ is the vector of all model parameters, and $\tilde{\boldsymbol{\theta}} = (M, \bar{\boldsymbol{\beta}}, a_t, b_t, \kappa^2, \bar{\bar{\alpha}}_p, v_\alpha^2, a_v, b_v, a_\Omega, \Omega_0)$ is the vector of all fixed hyper-parameters. Given $\tilde{\boldsymbol{\theta}}$ and data $\mathcal{D}$, the joint posterior of $\boldsymbol{\theta}$ and the patient specific random effects $\boldsymbol{s} = \{\boldsymbol{s}_i, i = 1, \ldots, n\}$ is

$$p(\boldsymbol{\theta}, \boldsymbol{s} \mid \mathcal{D}, \tilde{\boldsymbol{\theta}}) \;\propto\; \left\{ \prod_{i=1}^{n} p(\tilde{t}_i^o, \delta_i, \boldsymbol{z}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i, \boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \times p(\boldsymbol{s}_i \mid \boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \right\} p(\boldsymbol{\theta} \mid \tilde{\boldsymbol{\theta}}), \tag{7}$$

where the joint likelihood of the observed data for the i^{th} patient is the product

$$\begin{aligned} p(\tilde{t}_i^o, \delta_i, \boldsymbol{z}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i, \boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \;=\; & \{f(\tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0}, \boldsymbol{\beta}, \sigma_t^2)\}^{\delta_i} \{1 - F(\tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0}, \boldsymbol{\beta}, \sigma_t^2)\}^{1-\delta_i} \\ & \times \prod_{k=1}^{K} p(z_{ik} \mid \tau_i, \boldsymbol{x}_i, s_{i,k}, \boldsymbol{\alpha}_k, \boldsymbol{u}_k). \end{aligned}$$

We use Markov chain Monte Carlo (MCMC) simulation to generate posterior samples of the parameter and latent variable vectors, $(\boldsymbol{\theta}, \boldsymbol{s})$. For computational convenience, we approximate the DDP in (5) by truncating the infinite number of mixture components of F and G_k to the finite value M. The final weight is set to $w_M = 1 - \sum_{m=1}^{M-1} w_m$ to ensure that F and G_k are proper distributions. For sufficiently large M, the truncated sum produces inferences virtually identical to those with the infinite sum (Ishwaran and James, 2001; Rodriguez and Dunson, 2011). As discussed in Rodriguez and Dunson (2011), if there is a discrepancy between the posterior distributions under the truncated and infinite sums, then the model is sensitive to the choice of M. Any value of M that has a small value for w_M is sufficiently large to produce a negligible discrepancy. We examined the posterior distribution of w_M, and assessed sensitivity of the model to several different M values for the breast cancer dataset. We found that the truncated process is robust to the choice of M, if M is sufficiently large. This led us to use $M = 15$ for the data analysis and simulation studies. Computational details are given in Supplementary §1.1. A computer program "utility-analysis" for fitting the proposed model is available from `https://users.soe.ucsc.edu/~juheelee/`.

3 Utility Functions for PFS and TTB

In this section, we describe how a utility function was constructed for the breast cancer data analysis. While we focus on the case where $\boldsymbol{y}$ consists of PFS time and TTB, the methodology may be applied generally in settings where $\boldsymbol{y}$ is a single variable, a bivariate binary or ordinal variable, or some combination of two or more discrete and continuous outcomes. Examples are given in Supplementary § 2.

Given the reduction of the K-dimensional toxicity vector $\boldsymbol{z}$ to the scaled TTB q, $0 \leq q \leq 1$, we will construct a utility function for the pair (t, q). A departure of our utility formulation from previous published outcome utilities is that we construct U so that it

varies with covariates $\boldsymbol{x}$ as well as the outcomes (t, q). We let $\tau = 1$ for $L{+}B$ and $\tau = 0$ for L, so choosing τ for a future patient with a covariates $\boldsymbol{x}^{\text{new}}$ is the target of our decision analysis. To ensure a consistent utility function that quantifies trade-offs between t and q for each $\boldsymbol{x}$, we require

$$U(t, q, \boldsymbol{x}) > U(t, q', \boldsymbol{x}) \text{ if } q < q' \text{ for any } t, \quad \text{and } U(t, q, \boldsymbol{x}) < U(t', q, \boldsymbol{x}) \text{ if } t < t' \text{ for any } q.$$

That is, considered individually with the other outcome variable fixed, smaller TTB and longer PFS each must be more desirable.

The form of the utility function given here, and the numerical values that it takes on, were obtained based on the consensus of two medical oncologists who are co-authors of this paper, PM and BL, one of whom is a breast cancer subspecialist. The first step of our construction was to specify a parametric total utility function, which we defined generally as the product

$$U_{\text{tot}}(t, q, \boldsymbol{x}) = U_{\text{PFS}}(t, \boldsymbol{x}) \times U_{\text{TTB}}(q, \boldsymbol{x}), \tag{8}$$

subject to the constraints $0 \leq U_{\text{PFS}}(t, \boldsymbol{x}) \leq U_{\max}$ and $0 < U_{\text{TTB}}(q, \boldsymbol{x}) \leq 1$. For a given $\boldsymbol{x}$, the utility component $U_{\text{TTB}}(q, \boldsymbol{x})$ acts multiplicatively to decrease the utility component $U_{\text{PFS}}(t, \boldsymbol{x})$, and one may regard multiplying by $U_{\text{TTB}}(q, \boldsymbol{x})$ as penalizing the PFS utility, where the magnitude of the penalty is determined by q.

To apply this to the breast cancer dataset, we constructed a functional form for (8) to reflect this particular treatment setting. We denote the prognostic covariates by $x_1 =$ Age, an indicator x_2 of measurable disease, and an indicator x_3 of whether the patient's disease free interval prior to trial entry was > 24 months, with all three included in the regression models for outcomes t and $\boldsymbol{z}$. While the randomization for the breast cancer trial was stratified by x_2 and x_3 to improve precision, based on clinical experience PM and

BL decided that neither x_2 nor x_3 should have any effect on the utility function, whereas x_1 = Age is included in U_{TTB}. This is because, in clinical practice, the utility gained by greater PFS is similar regardless of age group, while older patients tend to care more about maintaining a good quality of life, quantified by a lower TTB. Thus, older patients are less likely than younger patients to accept a higher level of toxicity for the same PFS benefit. Accordingly, we only used the prognostic covariate x_1 = Age in the utility function, and constructed $U_{\text{tot}}(t, q, \boldsymbol{x}) = U_{\text{tot}}(t, q, Age)$ so that, for any PFS time t, a given value $q > 0$ for TTB decreases the utility for an older patient more than for a young patient. We also assumed that $U_{\text{PFS}}(t, Age) = U_{\text{PFS}}(t)$. We set $U_{\max} = 100$ since the domain (0, 100) is easy to interpret, and set $U_{\text{TTB}}(0, Age) = 1$ for any Age. Thus, if a patient has no toxicity ($q = 0$) then $U_{\text{tot}}(t, 0, Age) = U_{\text{PFS}}(t)$. For example, if $q = 0.50$, $U_{\text{PFS}}(36) = 80$, and $U_{\text{TTB}}(0.50, 60) = 0.50$, then the utility of 60-month PFS time and TTB = 0.5 is $U_{\text{tot}}(36, 0.50, 60) = 80\times 0.5 = 40$ for a 60-year-old patient, whereas $U_{\text{TTB}}(0.50, 40) = 0.70$ for a 40-year old patient gives the much higher total utility $U_{\text{tot}}(36, 0.50, 40) = 80\times\ 0.7 = 56$.

For the PFS component, the clinicians PM and BL specified the particular values $U_{\text{PFS}}(24) = 50$, $U_{\text{PFS}}(48) = 95$, and required $\lim_{t\to\infty} U_{\text{PFS}}(t) = 100$ for patients with any Age. This was based on the clinical experience that, in the hormone receptor-positive metastatic breast cancer setting, the increase in utility with increasing PFS is linear for up to four years, after which there is generally not much increase in utility with increasing PFS. This is due, in part, to the availability of newer regimens that can be given as salvage therapy to patients whose disease progresses after four years. To reflect this, we constructed the following parametric function for $U_{\text{PFS}}(t)$:

$$U_{\text{PFS}}(t) = \begin{cases} U_0 \left(\frac{t}{t_0}\right)^a & \text{if } t < t_0 \\ \frac{U_{\max}}{1+\exp(-b_1 t)} & \text{if } t \geq t_0, \end{cases} \tag{9}$$

and we set t_0=48 months, U_0=95, and $U_{\max}$=100. This function increases in t up to 48 months, with a small additional increase for $t > 48$ months. Since $U_{\text{PFS}}(24) = 50$, the equation $95 \times (24/48)^a = 50$ gives $a = 0.926$. Similarly, since $U_{\text{PFS}}(48) = 100/\{1+\exp(-b_1 48)\} = 95$, this gives $b_1 = 0.061$. The resulting function $U_{\text{PFS}}(t)$ is plotted in Fig 2(a).

We defined $U_{\text{TTB}}(q, Age)$ to vary with Age so that it decreases at a faster rate for older Age. PM and BL established numerical values of U_{TTB} in [0, 1] for each pair of (q, Age) values specified on a grid. These are tabulated in Supplementary Tab 2. We constructed a parametric function that closely approximates these elicited TTB utilities by exploring various functional forms, and chose

$$U_{\text{TTB}}(q, Age) = \exp\{-q^2/(2g^2(Age))\}, \quad \text{for } 0 \le q \le 1, \tag{10}$$

where $g(Age) = \exp(c_0 + c_1 Age)$. As q increases, the total utility given in (8) decreases by a factor of the above exponential function of q^2, and U_{PFS} is penalized in U_{tot} using this exponential value. If no toxicity occurs, i.e., $q = 0$, then $U_{\text{TTB}}(0, Age) = 1$ and $U_{\text{tot}}(t, q, Age) = U_{\text{PFS}}(t)$, as required. The function $U_{\text{TTB}}(q, Age)$ has an inflection point at $q = g(Age)$, and for any $q > g(Age)$, U_{tot} decreases to 0 very quickly. We defined $g(Age)$ to be a decreasing function of Age so that $U_{\text{TTB}}(q, Age)$ decreases in q faster for larger values of Age. To obtain this, we restricted $c_1 < 0$ and calibrated the values of c_0 and c_1 using the elicited numerical utilities. The numerical values $c_0 = 0.823$ and $c_1 = -0.05$ yield a good approximation. Details are given in Supplementary §2. Fig 2(b)-(c) compares $U_{\text{tot}}(t, q, Age)$ (solid line) to the elicited values (dots connected by dotted lines) for $Age = 50$, 65, and 85. In each plot, different colors represent different values of q. The figure illustrates how $U_{\text{tot}}(t, q, Age)$ decreases with q, and how the magnitude of decrease changes with Age. The specific utility function described here may be questioned due to its subjectivity. However, ordering the consequences of decisions is inherently subjective and is necessary for decision

making, and using an explicit utility function that is constructed based on expert knowledge produces meaningful decisions.

To use this structure for individualized treatment selection, we exploit the Bayesian model to compute the posterior predictive (PP) distribution of $(\tilde{t}, \boldsymbol{z})$ for a new patient with prognostic covariates $\boldsymbol{x}^{\text{new}} = (x_1^{\text{new}}, x_2^{\text{new}}, x_3^{\text{new}})$ assuming a particular τ is given to the patient, similar to the examples in Supplementary § 2,

$$p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D}) = \int \int p(\tilde{t}, \boldsymbol{z} \mid \boldsymbol{\theta}, \boldsymbol{s}^{\text{new}}, \boldsymbol{x}^{\text{new}}, \tau) p(\boldsymbol{\theta}, \boldsymbol{s}^{\text{new}} \mid \mathcal{D}) d\boldsymbol{s}^{\text{new}} d\boldsymbol{\theta}. \tag{11}$$

By averaging the likelihood of the new patient's future outcomes $(\tilde{t}, \boldsymbol{z})$ over the joint posterior distribution of $(\boldsymbol{\theta}, \boldsymbol{s}^{\text{new}})$, the PP distribution in (11) provides a fully model-based criterion for making inferences to compare treatments, with appropriate quantification of uncertainty. The PP distribution of the utility, $p\{U_{\text{tot}}(t, q, Age^{\text{new}}) \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D}\}$, for a new patient with prognostic covariates $\boldsymbol{x}^{\text{new}}$ can be derived directly from (11).

We use the predictive mean total utility for each τ as a basis for treatment selection. For $\tau = 0$ corresponding to L and $\tau = 1$ for $L{+}B$, the predictive mean total utility is

$$\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}}) = \sum_{z_1=0}^{J-1} \cdots \sum_{z_K=1}^{J-1} \int_{\mathbb{R}} U_{\text{tot}}(\tilde{t}, q, Age^{\text{new}}) p(\tilde{t}, \boldsymbol{z} \mid \tau, \mathcal{D}, \boldsymbol{x}^{\text{new}}) d\tilde{t}. \tag{12}$$

One may choose the treatment τ having larger $\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ for the new patient. In general, $\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ is a function of the entire $\boldsymbol{x}^{\text{new}}$ vector and τ through the utility function and/or the probability distribution, and this still would be the case if U_{tot} did not depend on Age.

In general, another criterion for comparing treatments τ and τ' studied in a trial is to compare the PP distributions $p\{U_{\text{tot}}(t, q, Age^{\text{new}}) \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D}\}$ and $p\{U_{\text{tot}}(t, q, Age^{\text{new}}) \mid \tau', \boldsymbol{x}^{\text{new}}, \mathcal{D}\}$ of the total utilities. To do this, we define the posterior probability that treatment

τ' has a larger total utility than treatment τ for a new patient with prognostic variables $\boldsymbol{x}^{\text{new}}$,

$$\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau') = \Pr\left\{U_{\text{tot}}(t(\tau), q(\tau), Age^{\text{new}}) < U_{\text{tot}}(t(\tau'), q(\tau'), Age^{\text{new}}) \mid \tau, \tau', \mathcal{D}, \boldsymbol{x}^{\text{new}}\right\}. \quad (13)$$

One may select τ' to treat a new patient with $\boldsymbol{x}^{\text{new}}$ if $\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau') \geq 0.5$, and otherwise select τ. This can be done by first computing the joint PP distribution of $(\tilde{t}(\tau), \boldsymbol{z}(\tau), \tilde{t}(\tau'), \boldsymbol{z}(\tau'))$, i.e., of two sets of outcomes with one for each treatment,

$$\begin{aligned} p(\tilde{t}(\tau), \boldsymbol{z}(\tau), \tilde{t}(\tau'), \boldsymbol{z}(\tau') \mid \tau, \tau', \boldsymbol{x}^{\text{new}}, \mathcal{D}) \;=\; & \int\int p(\tilde{t}(\tau), \boldsymbol{z}(\tau) \mid \boldsymbol{\theta}, \boldsymbol{s}^{\text{new}}, \boldsymbol{x}^{\text{new}}, \tau) \quad (14) \\ & \times p(\tilde{t}(\tau'), \boldsymbol{z}(\tau') \mid \boldsymbol{\theta}, \boldsymbol{s}^{\text{new}}, \boldsymbol{x}^{\text{new}}, \tau') p(\boldsymbol{s}^{\text{new}}, \boldsymbol{\theta} \mid \mathcal{D}) d\boldsymbol{s}^{\text{new}} d\boldsymbol{\theta}, \end{aligned}$$

which can be used in turn to compute $\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau')$. Although U_{tot} depends on $\boldsymbol{x}^{\text{new}}$ only through Age, the joint distribution in (14) depends on the entire $\boldsymbol{x}^{\text{new}}$ and (τ, τ'), and Δ does as well. The decision criteria $\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ and $\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau')$ may be computed numerically using MCMC samples of $\boldsymbol{\theta}$ simulated from $p(\boldsymbol{\theta} \mid \mathcal{D}, \tilde{\boldsymbol{\theta}})$. Computational details are given in Supplementary § 1.2.

4 Decision Making for the Breast Cancer Data

In this section, we illustrate the proposed decision making procedures by application to the breast cancer dataset. We fit the statistical model in § 2 to the data and used the elicited utility in § 3. To fit the Bayesian model, we specified values of the fixed hyperparameters $\tilde{\boldsymbol{\theta}}$ as follows. We let the DP concentration parameter be $\xi = 1$, $\kappa^2 = 5$ for the prior of $\beta_{m,p}$, $a_v = 5$ and $b_v = 1$ for the priors of v_p^2, $a_e = b_e = 3$ for the priors of $e_{k,j}$, and $a_t = 3$ and $b_t = 1$ for the prior of σ_t^2. We fixed $\bar{\beta}_p = 0$, $p \geq 1$ and used the empirical average of the observed $\tilde{t}$ to specify $\bar{\beta}_0$. Similarly, we let $\bar{\bar{\alpha}}_p = 0$, $p = 1, \ldots, P$ and used the empirical probabilities of $z_{i,k}$ being 0 to determine $\bar{\bar{\alpha}}_0$. We fixed $\sigma_z^2 = 9$. For Ω, we let $a_\Omega = 50$,

and $\Omega_0 = 0.01(a_\Omega - K - 1)\mathrm{I}_{K+1}$. We discarded the first 20,000 iterates for burn-in, and kept the next 5,000 iterates for posterior inference. We examined mixing and convergence of the Markov chains using trace plots, and did not find evidence of poor mixing or bad convergence.

Estimates of the posterior predictive distributions of t, q, and U_{tot} are shown in Fig 3 for a future patient with $\boldsymbol{x}^{\text{new}} = (Age^{\text{new}}, 0, 0)$. The top, middle, and bottom rows of the figure correspond to $Age^{\text{new}} = 55$, 65, and 75 years, with treatment L ($\tau = 0$) represented by red and L+B ($\tau = 1$) by blue. PP estimates of the survival functions $S(t \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ with 95% pointwise credible intervals are given in the left column, and the middle column gives estimates of the PP distributions, $p(q \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$, of TTB. Estimated $S(t \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ for L+B are slightly above those for L at all ages, indicating a small overall improvement in PFS with L+B. In contrast, estimated $p(q \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ for L+B have much longer and thicker right tails, indicating an increased risk of toxicity events with L+B compared to L. From the figures, effects of Age on PFS and TTB are small. Additional comparisons of the PP distributions of PFS (t) and TTB (q) for treatments L and L+B for more values of $\boldsymbol{x}^{\text{new}}$ are given in panels (b) and (c) of Supp. Figs 1 and 2. These give PP point estimates of t and q as functions of $\boldsymbol{x}^{\text{new}}$. Panel (a) of the figures shows estimated PP probabilities that L+B yields greater PFS than L, and that L+B yields greater TTB than L. The figures also indicate that L+B tends to yield better PFS than L, while L+B is more likely to be associated with higher TTB. Swanson and Lin (1994) also noted that older patients are more likely to respond to hormone therapies such as L+B or L. Our inferences for PFS and TTB considered individually agree with the findings reported by Dickler et al. (2016).

As noted earlier, because treatment comparisons based on TTB and PFS considered separately lead in opposite directions, these results do not provide a clear basis for choosing one treatment over the other, either overall or for individual patients. Considering the outcomes together by using the utility function of treatment and Age provides a useful

tool for resolving this problem. The plots in the rightmost column of Fig 3 compare PP distributions of U_{tot} for the two treatments, as functions of *Age*, indicating that the utility benefit of *L+B* over *L* diminishes with increasing *Age*. In the plot, we assume that $x_2^{\text{new}} = 0$ and $x_3^{\text{new}} = 0$. Posterior predictive estimates $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L+B)$ of the probabilities that treatment *L+B* has greater utility than treatment *L* are 0.56, 0.54, and 0.46 for 55, 65 and 75 year old patients, respectively. Thus, in terms of overall utilty accounting for both PFS and TTB, decisions based on $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L+B)$ would be to give *L+B* to patients with $Age^{\text{new}} < 70$, but give *L* to patients with $Age^{\text{new}} \geq 70$.

Fig 4(a) illustrates $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L{+}B)$ on a grid of $Age^{new} = x_1^{\text{new}}$ values for the combinations $(x_2^{\text{new}}, x_3^{\text{new}})$ =(0, 0), (0, 1), (1, 0) and (1, 1). Overall, while *L+B* tends to yield a greater utility for younger patients, *L* is expected to have a greater utility for older patients. Although PFS is improved by *L+B* for patients of all ages, increases in TTB with both *L+B* and age make it less desirable for older patients. For example, $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L+B) < 0.5$ for a 70-year-old patient, with $\boldsymbol{x}^{\text{new}} = (70, 0, 1)$ (dashed line), implying that *L* is a better treatment option than *L+B* for this patient. The differences in Δ for varying $(x_2^{\text{new}}, x_3^{\text{new}})$ values are small, and the values of x_2 and x_3 do not change any decisions significantly. Posterior expected utility estimates $\hat{\bar{u}}_{\text{tot}}(\boldsymbol{x}^{\text{new}}, \tau)$ are computed on the age grid for different treatment options, shown in panels (b) and (c) of the figure. The figure in panel (c) shows that the expected utility decreases rapidly with age for *L+B*, while it increases slightly for *L*. Thus, Age-specific recommendations in terms of $\hat{\bar{u}}_{\text{tot}}(\boldsymbol{x}^{\text{new}}, \tau)$ and $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L{+}B)$ are the same for these values of $(x_2^{\text{new}}, x_3^{\text{new}})$.

5 Simulation Study

In this section, we summarize a simulation study to illustrate the performance of the proposed utility-based decision making procedure. To generate data similar to the breast cancer

dataset, we set the number of patients to be $n = 350$, with three covariates and $K = 20$ toxicity types each having $J = 6$ grades. To mimic the covariate distribution in the breast cancer dataset, we randomly drew a sample of size 350 from $(\tau_i, x_{i,1}, x_{i,2})$ with replacement, where τ, x_1 and x_2 are a binary treatment indicator, age, and a binary indicator of disease measurability, respectively, so our $P = 2$. We simulated patient-specific frailty vectors $\boldsymbol{s}_i^{\text{TR}}$ for correlation between $\boldsymbol{y}_i^{\text{TR}} = (\tilde{t}_i^{\text{TR}}, \tilde{\boldsymbol{z}}_i^{\text{TR}})$ within a patient. To illustrate how the BNP model flexibly accommodates complicated relationships between τ_i, $\boldsymbol{x}_i$ and $\boldsymbol{y}_i$, we generated $\boldsymbol{y}_i^{\text{TR}}$ from a mixture of two regression functions, each having main effects for τ and x_p, and their interaction effects. Specifically, we generated $\boldsymbol{s}_i^{\text{TR}} \overset{iid}{\sim} \text{N}_{K+1}(\boldsymbol{0}, \Omega^{\text{TR}})$, where Ω^{TR} assumes that variances of $s_{i,k}^{\text{TR}}$ are 0.05 and correlations between $s_{i,k}^{\text{TR}}$ and $s_{i,k'}^{\text{TR}}$ are 0.5 if $k, k' > 0$ and $k \neq k'$, and -0.5 if k or $k' = 0$. Given $\boldsymbol{s}_i^{\text{TR}}$, we then generated $\boldsymbol{y}_i^{\text{TR}}$ from the following distribution with probability 0.4,

$$\begin{aligned}
\tilde{t}_i \mid \boldsymbol{x}_i, s_{i,0}^{\text{TR}} &\overset{indep}{\sim} \text{N}(4.0 - 0.5\tau_i + 0.5x_{i,1} + 0.5\tau_i x_{i,1} + s_{i,0}^{\text{TR}}, 0.6), \\
\tilde{z}_{i,k} \mid \boldsymbol{x}_i, s_{i,k}^{\text{TR}}, \boldsymbol{\alpha}^{\text{TR}} &\overset{indep}{\sim} \text{N}(\alpha_{1,k,0}^{\text{TR}} + \alpha_{1,k,1}^{\text{TR}}\tau_i + \alpha_{1,k,2}^{\text{TR}}x_{i,1} + \alpha_{1,k,3}^{\text{TR}}\tau_i x_{i,1} + s_{i,k}^{\text{TR}}, 4),
\end{aligned}$$

where $\alpha_{1,k,0}^{\text{TR}} \overset{iid}{\sim} \text{U}(-4.5, -4.0)$, $\alpha_{1,k,1}^{\text{TR}} \overset{iid}{\sim} \text{U}(0.5, 1.0)$, $\alpha_{1,k,2}^{\text{TR}} \overset{iid}{\sim} \text{U}(0.3, 0.8)$, and $\alpha_{1,k,3}^{\text{TR}} \overset{iid}{\sim} \text{U}(0, 0.5)$ were simulated for all k. With the remaining probability 0.6, we generated

$$\begin{aligned}
\tilde{t}_i \mid \boldsymbol{x}_i, s_{i,0}^{\text{TR}} &\overset{indep}{\sim} \text{N}(1.5 + 0.8\tau_i - 0.5x_{2,i} + s_{i,0}^{\text{TR}}, 0.6) \\
\tilde{z}_{i,k} \mid \boldsymbol{x}_i, s_{i,k}^{\text{TR}}, \boldsymbol{\alpha}^{\text{TR}} &\overset{indep}{\sim} \text{N}(\alpha_{2,k,0}^{\text{TR}} + \alpha_{2,k,1}^{\text{TR}}\tau_i + \alpha_{2,k,2}^{\text{TR}}x_{i,2} + s_{i,k}^{\text{TR}}, 4),
\end{aligned}$$

where $\alpha_{2,k,0}^{\text{TR}} \overset{iid}{\sim} \text{U}(-4.5, -4.0)$, $\alpha_{2,k,1}^{\text{TR}} \overset{iid}{\sim} \text{U}(1.0, 1.5)$, and $\alpha_{2,k,2}^{\text{TR}} \overset{iid}{\sim} \text{U}(0.5, 1.0)$ were simulated for all k. We simulated censoring times $c_i \overset{iid}{\sim} \text{N}(4.63, 2)$, and let $t_i^o = \min(t_i, c_i)$ and $\delta_i = 1(c_i < t_i)$. In the simulated data, 22.86% of t_i's was censored. For ordinal outcomes $z_{i,k}$, we simulated toxicity type specific cutoff points $u_{k,j}^{\text{TR}}$, $j = 2, \ldots, J-1$ for each k, and let $z_{i,k} = j$

if $u^{\text{TR}}_{k,j-1} < \tilde{z}^{\text{TR}}_{i,k} \leq u^{\text{TR}}_{k,j}$. Details of the simulation setup are given in Supplementary § 5. The utility function $U_{\text{tot}}(t, q, Age)$ elicited in § 3 is assumed.

Supplementary Fig 3(a), (b) and (e) illustrate the true total utility $p(U_{\text{tot}} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$, TTB distribution $p(q \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$, and PFS survival function $S(t \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$ (solid lines) in the simulation. The covariate vector $\boldsymbol{x} = (65, 0)$ of a 65-year old patient with no measurable disease is used for illustration. In each plot, red represents $\tau = 0$ and blue represents $\tau = 1$. From panels (b) and (e), treatment $\tau = 1$ has a greater TTB and greater expected PFS time than treatment $\tau = 0$, which complicates treatment choice. This is resolved by panel (a), which shows that U_{tot} with $\tau = 1$ (blue) is stochastically greater than U_{tot} with $\tau = 0$ (red), implying that $\tau = 1$ is better for a patient with $\boldsymbol{x} = (65, 0)$. Supplementary Fig 4 gives

$$\Delta^{\text{TR}}(\boldsymbol{x}, 0, 1) = \Pr\left(U_{\text{tot}}(t(0), q(0), Age) < U_{\text{tot}}(t(1), q(1), Age) \mid \tau = 0, \tau' = 1, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}}\right),$$

with the true expected utilities $\overline{U}^{\text{TR}}(\tau, \boldsymbol{x}) = \text{E}(U_{\text{tot}} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$ shown in dark green symbols for varying *Age*. For both treatments, the expected utility decreases with *Age*, shown in panels (a) and (b). The difference in $\overline{U}^{\text{TR}}(\tau, \boldsymbol{x})$ between the treatments and $\Delta^{\text{TR}}(\boldsymbol{x}, 0, 1)$ decrease with *Age*, indicating that the superiority of treatment $\tau = 1$ compared to treatment $\tau' = 0$ diminishes with *Age*.

We specified values of the fixed hyperparameters similar to those in § 4, and ran the MCMC simulation as described in § 2.2. Posterior inferences are summarized in Supplementary Fig 3. The posterior predictive distributions, $p(q \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ and $S(t \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ with $\boldsymbol{x}^{\text{new}} = (65, 0)$, are shown in Supplementary Fig 3(c)-(e), respectively, where $\tau = 0$ and 1 are in red and blue, respectively. In panel (e), the dashed lines represent the posterior mean estimates with 95% pointwise credible intervals in the shaded areas. Comparing the estimates in panels (d) and (e) to the truth under in panels (b), and (e) shows that the flexible BNP regression model captures the simulation truth reasonably well, which provides a good basis

for accurate statistical decision making. For example, the posterior predictive distribution of U_{tot} in panel (c) that provides a comprehensive criterion for treatment comparison, and is close to the truth in panel (a).

Supplementary Fig 4 illustrates posterior estimates $\hat{\Delta}(\boldsymbol{x}, 0, 1)$ of the probabilities that $\tau = 1$ has greater utility than $\tau = 0$, and compares posterior estimates of expected utilities $\hat{\bar{u}}(\tau, \boldsymbol{x}^{\text{new}})$ for each of $\tau = 0$ and 1, the varying $\boldsymbol{x}^{\text{new}} = (x_1^{\text{new}}, x_2^{\text{new}})$. For $x_1^{\text{new}} = Age$, a grid from 40 years to 80 years in 5 year increments was used, and values of $x_2^{\text{new}} \in \{0, 1\}$ are indicated by the symbols, + and x, respectively, with true values given in dark green. The model recovers the simulation truth reasonably well, and the decision-making procedure based on $\Delta(\boldsymbol{x}, 0, 1)$ in § 3, selects the truly superior treatment for all cases of $\boldsymbol{x}^{\text{new}}$. From panels (b) and (c), we observe discrepancies between $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ and $U_{\text{tot}}^{\text{TR}}(\tau, \boldsymbol{x}^{\text{new}})$ for small x_1 (Age). However, the ranks of $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ between the treatments are well estimated and the procedure of choosing τ to maximize $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ reliably selects the treatments with truly greater utility.

We simulated 100 datasets to further examine the performance of the proposed decision making procedures. The results are summarized in Supplementary Fig 5. Panels (a) and (d) of the figure show the distributions of $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, 0, 1)$ over the 100 datasets, for $x_2^{\text{new}} = 0$ and 1, respectively, where the symbols + and x represent $\Delta^{\text{TR}}(\boldsymbol{x}^{\text{new}}, 0, 1)$ with $x_2^{\text{new}} = 0$ and 1, respectively. In most cases, the proposed decision making procedure based on $\Delta(\boldsymbol{x}, 0, 1)$ produces the correct decisions. Panels (b), (c), (e) and (f) illustrate the distributions of predictive mean utility estimates, $\hat{\bar{u}}_{\text{tot}}(\boldsymbol{x}^{\text{new}}, \tau)$, using the 100 datasets for $(x_1, x_3) = (0, 0), (1, 0)$, (0, 1), and (1, 1), respectively. The proposed model produces reasonably good estimates of the expected utilities, and the ordering of $\bar{u}_{\text{tot}}(\boldsymbol{x}^{\text{new}}, \tau)$ are also well estimated overall.

6 Discussion

We have presented a formal decision making framework based on utility functions to address the goal of statistical decision making based on data from a randomized clinical trial. The key elements of our methodology are a multivariate Bayesian regression model and a utility function of outcomes and covariates. We assumed a DDP, which is a general, flexible family of Bayesian regression models.

The methodology was illustrated with a breast cancer dataset from a randomized clinical trial. This required close collaboration with oncologists to elicit a utility function that reflected this clinical setting, and the resulting utility function varied with Age to reflect different risk-benefit trade-offs between PFS and TTB for older versus younger patients. Our application illustrates that, by establishing a utility function that quantifies the risk-benefit trade-off between two competing outcomes, one can derive a rational basis for making treatment choices in settings where simpler comparisons in terms of individual outcomes give different, contradictory choices.

The particular form of our utility function was tailored to the breast cancer setting, and may not be appropriate in other clinical settings. In general, the requirements are that $U(\boldsymbol{y}, \boldsymbol{x})$ must be consistent in the arguments of $\boldsymbol{y}$ so that it makes sense, and that it be tractable enough to facilitate application. Beyond that, a utility should be constructed so that it provides a sensible basis for quantifying trade-offs between different outcomes. It also should be kept in mind that, if U is constructed to be a function of $\boldsymbol{y}$ but not $\boldsymbol{x}$, both the posterior mean utility $\bar{u}(\tau, \boldsymbol{x})$ and $\Delta(\boldsymbol{x}, \tau, \tau')$ still will vary with $\boldsymbol{x}$ due to the regression structure of $f(\boldsymbol{y} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta})$, and thus both criteria still will serve as a basis for making personalized treatment decisions.

Our framework may be generalized to accommodate more complex clinical settings, such as meta-analysis of multiple studies, optimization of a multi-stage treatment strategies, or

non-medical applications. Although the key elements will remain the same, such applications may be complex and will require tailoring to particular settings.

Acknowledgments

Juhee Lee's research was supported by NSF grant DMS-1662427. Pavlos Msaouel is supported by a Career Development Award by the American Society of Clinical Oncology, a Research Award by KCCure, the MD Anderson Khalifa Scholar Award, and the MD Anderson Physician-Scientist Award. Peter Thall's research was supported by NIH/NCI grants P01 2P30CA016672 and R01 CA261978.

This manuscript was prepared using data from Datasets NCT00601900-D1 and NCT00601900-D2 from the NCTN Data Archive of the National Cancer Institute's (NCI's) National Clinical Trials Network (NCTN). Data were originally collected from clinical trial NCT number NCT00601900, Endocrine Therapy With or Without Anti-VEGF Therapy: A Randomized, Phase III Trial of Endocrine Therapy Alone or Endocrine Therapy Plus Bevacizumab (NSC 704865) for Women With Hormone Receptor-Positive Advanced Breast Cancer. All analyses and conclusions in this manuscript are the sole responsibility of the authors and do not necessarily reflect the opinions or views of the clinical trial investigators, the NCTN, or the NCI. The CALGB-40503 dataset was made available through an NCI-Genentech agreement under which Genentech supplied the drug to support the study.

Supplementary Materials

Supplementary materials are available under the Paper Information link at the journal website.

References

Bagočius, V., Zavadskas, E. K., and Turskis, Z. (2014). Multi-person selection of the best wind turbine based on the multi-criteria integrated additive-multiplicative utility function. *Journal of civil engineering and management*, 20(4):590–599.

Bekele, B. and Thall, P. (2004). Dose-finding based on multiple toxicities in a soft tissue sarcoma trial. *Journal of the American Statistical Association*, 99:71–84.

Chen, W., Wiecek, M. M., and Zhang, J. (1998). Quality utility: a compromise programming approach to robust design. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 80326, page V002T02A032. American Society of Mechanical Engineers.

Chib, S. and Greenberg, E. (1998). Analysis of multivariate probit models. *Biometrika*, 85:347–361.

De Iorio, M., Johnson, W. O., Müller, P., and Rosner, G. L. (2009). Bayesian nonparametric nonproportional hazards survival modeling. *Biometrics*, 65(3):762–771.

De Iorio, M., Müller, P., Rosner, G. L., and MacEachern, S. N. (2004). An anova model for dependent random measures. *Journal of the American Statistical Association*, 99(465):205–215.

Detsky, A. S., Naglie, G., Krahn, M. D., Naimark, D., and Redelmeier, D. A. (1997). Primer on medical decision analysis: part 1getting started. *Medical Decision Making*, 17(2):123–125.

Dickler, M. N., Barry, W. T., Cirrincione, C. T., Ellis, M. J., Moynahan, M. E., Innocenti, F., Hurria, A., Rugo, H. S., Lake, D. E., Hahn, O., et al. (2016). Phase iii trial evaluating letrozole as first-line endocrine therapy with or without bevacizumab for the treatment

of postmenopausal women with hormone receptor–positive advanced-stage breast cancer: Calgb 40503 (alliance). *Journal of Clinical Oncology*, 34(22):2602.

Ishwaran, H. and James, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association*, 96(453):161–173.

Le-Rademacher, J., Hillman, S., Storrick, E., Mahoney, M., Thall, P., Jatoi, A., and Mandrekar, S. (2020). Adverse event burden score a versatile summary measure in cancer clinical trials. *Cancers*, 12:3256.

Loewenstein, G. F., Thompson, L., and Bazerman, M. H. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social psychology*, 57(3):426.

MacEachern, S. N. (1999). Dependent nonparametric processes. In *ASA proceedings of the section on Bayesian statistical science*, volume 1, pages 50–55. Alexandria, Virginia. Virginia: American Statistical Association; 1999.

Mitra, R. and Müller, P. (2015). *Nonparametric Bayesian Inference in Biostatistics*. Springer-Verlag.

Müller, P. and Mitra, R. (2013). Bayesian nonparametric inference why and how. *Bayesian Analysis*, 8(2):269–302.

Müller, P., Quintana, A., Jara, A., and Hanson, T. (2015). *Bayesian Nonparametric Data Analysis*. Springer-Verlag.

Naglie, G., Krahn, M. D., Naimark, D., Redelmeier, D. A., and Detsky, A. S. (1997). Primer on medical decision analysis: part 3estimating probabilities and utilities. *Medical Decision Making*, 17(2):136–141.

Pennings, J. M. and Smidts, A. (2003). The shape of utility functions and organizational behavior. *Management Science*, 49(9):1251–1263.

Rodriguez, A. and Dunson, D. B. (2011). Nonparametric bayesian models through probit stick-breaking processes. *Bayesian analysis (Online)*, 6(1).

Roy, S. K., Maity, G., and Weber, G.-W. (2017). Multi-objective two-stage grey transportation problem using utility function with goals. *Central European Journal of Operations Research*, 25(2):417–439.

Sethuraman, J. (1994). A constructive definition of dirichlet priors. *Statistica sinica*, pages 639–650.

Swanson, G. M. and Lin, C.-S. (1994). Survival patterns among younger women with breast cancer: the effects of age, race, stage, and treatment. *Journal of the National Cancer Institute. Monographs*, (16):69–77.

Thall, P., Müller, P., Xu, Y., and Guidani, M. (2017). Bayesian nonparametric statistics: A new toolkit for discovery in cancer research. *Pharmaceutical Statistics*, 16:414–423.

Walsh, W. E., Tesauro, G., Kephart, J. O., and Das, R. (2004). Utility functions in autonomic systems. In *International Conference on Autonomic Computing, 2004. Proceedings.*, pages 70–77. IEEE.

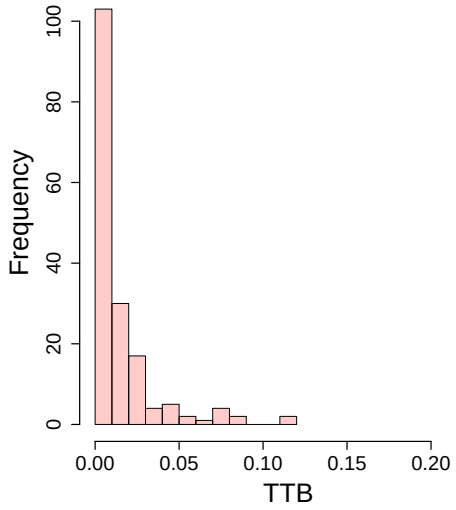

(a) scaled TTB for L

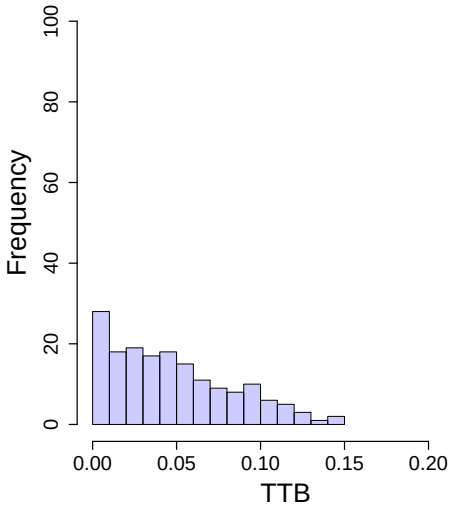

(b) scaled TTB for $L + B$

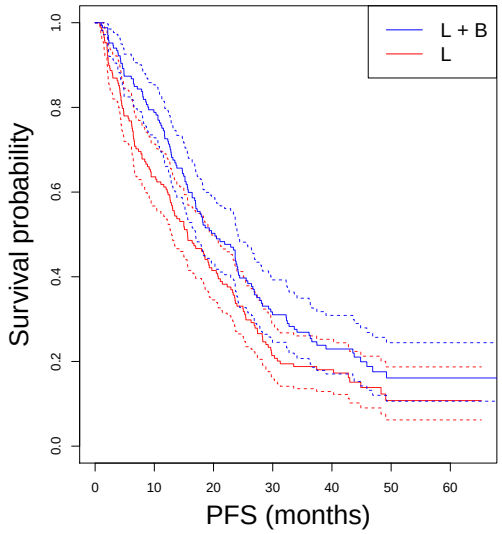

(c) Kaplan-Meier estimates of $S(t)$

Figure 1: For the breast cancer data, histograms of scaled total toxicity burden (TTB) are given in panels (a) and (b). Panel (c) illustrates Kaplan-Meier estimates of survival functions, $S(t)$. Blue and red represent treatments, letrozole plus bevacizumab (L+B) and letrozole plus placebo (L), respectively.

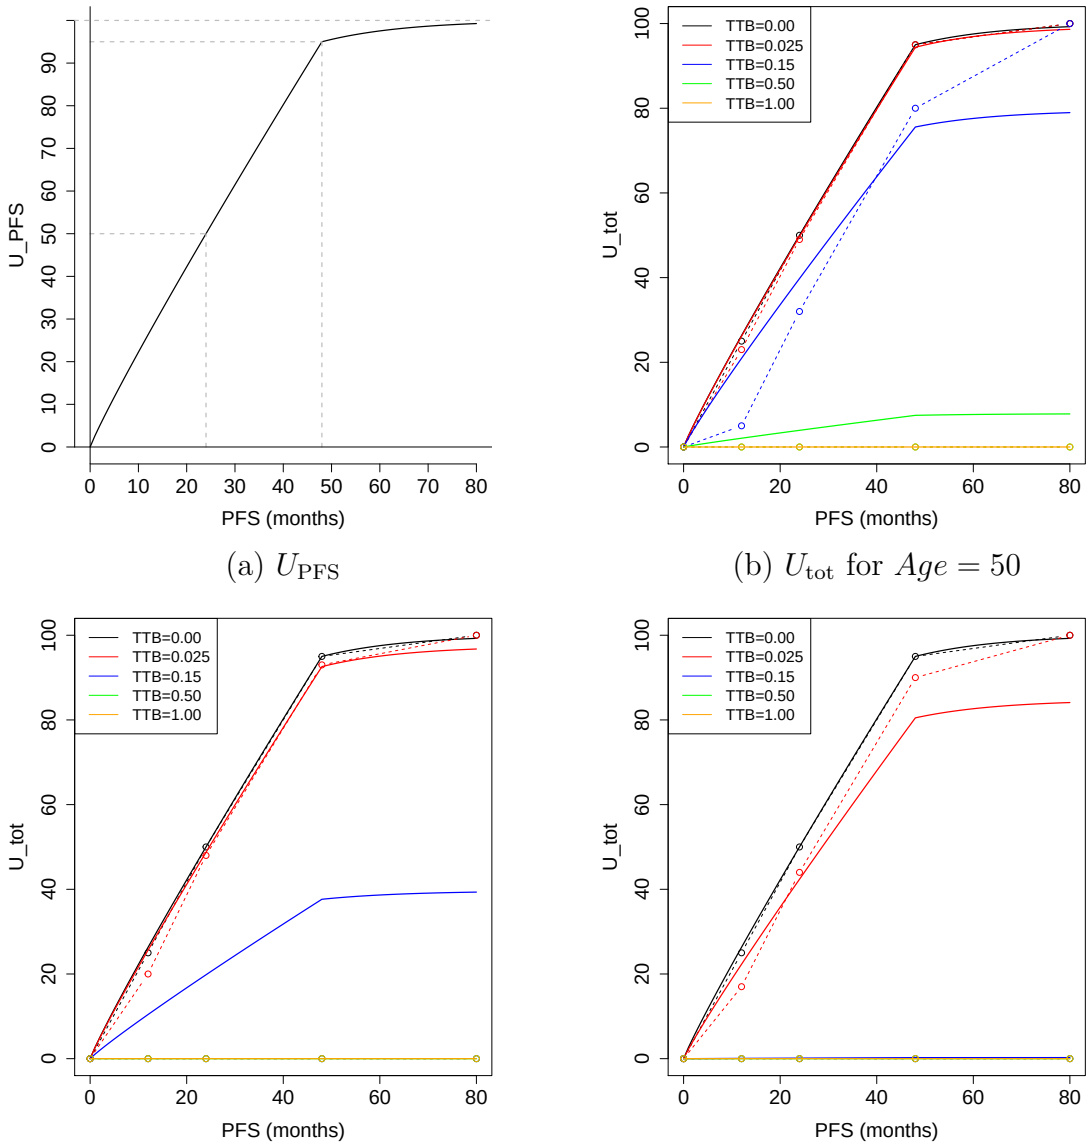

(a) U_{PFS}

(b) U_{tot} for $Age = 50$

(c) U_{tot} for $Age = 65$

(d) U_{tot} for $Age = 85$

Figure 2: Illustration of utility functions. Panel (a) has $U_{\text{PFS}}(t)$, the utility function of progression free survival (PFS) time with t = PFS time. Panels (b)-(d) have the total utility functions $U_{\text{tot}}(t, q, Age)$, where t = PFS time and q = total toxicity burden (TTB).

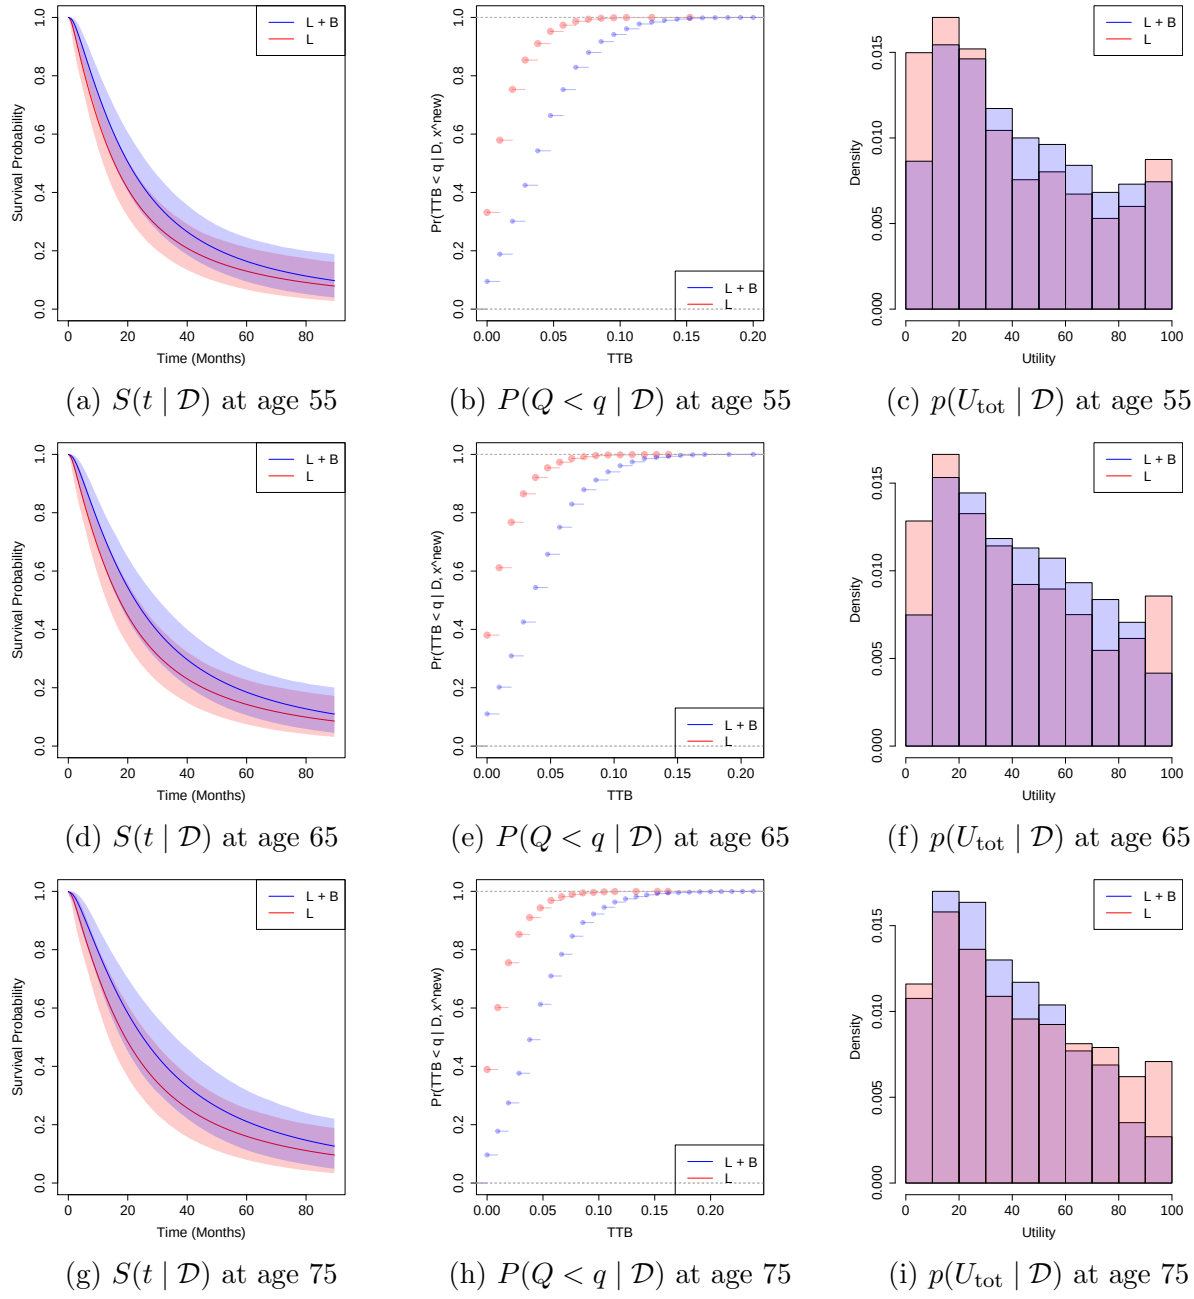

(a) $S(t \mid \mathcal{D})$ at age 55 (b) $P(Q < q \mid \mathcal{D})$ at age 55 (c) $p(U_{\text{tot}} \mid \mathcal{D})$ at age 55

(d) $S(t \mid \mathcal{D})$ at age 65 (e) $P(Q < q \mid \mathcal{D})$ at age 65 (f) $p(U_{\text{tot}} \mid \mathcal{D})$ at age 65

(g) $S(t \mid \mathcal{D})$ at age 75 (h) $P(Q < q \mid \mathcal{D})$ at age 75 (i) $p(U_{\text{tot}} \mid \mathcal{D})$ at age 75

Figure 3: [Breast Cancer Trial Data] Estimated posterior predictive (PP) survival functions $S(t \mid \mathcal{D})$ where $\mathcal{D}$ denotes data, are given in panels (a), (d), and (g); estimates of PP cumulative distributions of total toxicity burden (TTB, Q) in panels (b), (e), and (h); and PP distribution estimates of total utility U_{tot} in (c), (f) and (i). The top, middle, and bottom rows are correspond to ages (x_1^{new}) 55, 65, and 75 years. Covariates $(x_2^{\text{new}}, x_3^{\text{new}}) = (0, 0)$, which indicate absence of measurable disease at baseline and the patient's disease free interval prior to trial entry $\leq$ 24 months, are fixed. In each panel, red and blue represent treatments L ($\tau = 0$) and $L + B$ ($\tau = 1$), respectively.

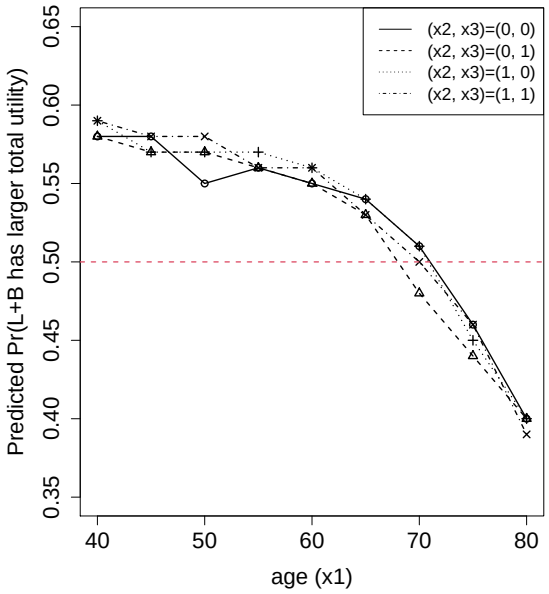

(a) $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L+B)$

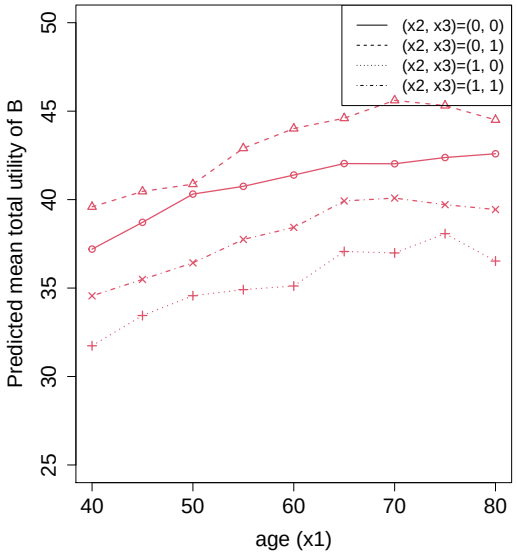

(b) $\hat{\bar{u}}_{\text{tot}}(L, \boldsymbol{x}^{\text{new}})$

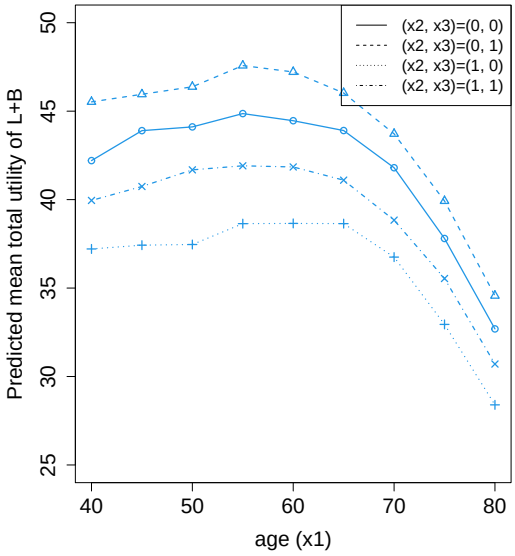

(c) $\hat{\bar{u}}_{\text{tot}}(L+B, \boldsymbol{x}^{\text{new}})$

Figure 4: [Breast Cancer Trial Data] (a) Estimated posterior predictive (PP) probability $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, L, L+B)$ that treatment letrozole plus bevacizumab (L+B) has greater utility than treatment letrozole plus placebo (L) for different cases of $\boldsymbol{x}^{\text{new}} = (x_1^{\text{new}}, x_2^{\text{new}}, x_3^{\text{new}})$, where x_1^{new} = age, x_2^{new} = presence of measurable disease at baseline, and x_3^{new} = disease free interval prior to trial entry > 24 months for an unobserved patient. Panels (b) and (c) give posterior predictive mean utility estimates $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ for treatments $\tau = L$+B and L, respectively, for different values of $\boldsymbol{x}^{\text{new}}$.

Supplementary Materials: Utility Based Bayesian Personalized Treatment Selection for Advanced Breast Cancer

1 Posterior Computation

1.1 Details of MCMC Simulation

Recall that $\boldsymbol{\theta} = (w_m, \boldsymbol{\beta}_m, \sigma_t^2, \boldsymbol{\alpha}_{m,k}, \boldsymbol{e}, \bar{\alpha}_{k,p}, v_p^2, \Omega)$ is the vector of all model parameters, and $\tilde{\boldsymbol{\theta}} = (M, \bar{\boldsymbol{\beta}}, a_t, b_t, \tau_p^2, \bar{\bar{\alpha}}_p, v_\alpha^2, a_v, b_v, a_s, \Omega^0)$ is the vector of all fixed hyper-parameters. Given $\tilde{\boldsymbol{\theta}}$ and data $\mathcal{D}$, the joint posterior of $\boldsymbol{\theta}$ and the patient specific random effects $\boldsymbol{s} = \{\boldsymbol{s}_i, i = 1, \ldots, n\}$ is in (7) of the main text,

$$p(\boldsymbol{\theta}, \boldsymbol{s} \mid \mathcal{D}, \tilde{\boldsymbol{\theta}}) \;\propto\; \left\{ \prod_{i=1}^{n} p(\tilde{t}_i^o, \delta_i, \boldsymbol{z}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i, \boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \times p(\boldsymbol{s}_i \mid \boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \right\} p(\boldsymbol{\theta} \mid \tilde{\boldsymbol{\theta}}), \tag{1}$$

where the joint likelihood of the observed data for the i^{th} patient is the product

$$\begin{aligned} p(\tilde{t}_i^o, \delta_i, \boldsymbol{z}_i \mid \tau_i, \boldsymbol{x}_i, \boldsymbol{s}_i, \boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) &= \{f(\tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0}, \boldsymbol{\beta}, \sigma_t^2)\}^{\delta_i} \{1 - F(\tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0}, \boldsymbol{\beta}, \sigma_t^2)\}^{1-\delta_i} \\ &\quad \times \prod_{k=1}^{K} p(z_{ik} \mid \tau_i, \boldsymbol{x}_i, s_{i,k}, \boldsymbol{\alpha}_k, \boldsymbol{u}_k). \end{aligned}$$

For computational convenience, when fitting the model we approximate the DDP by truncating the number of mixture components of f and g_k to M;

$$\begin{aligned} \tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0} &\overset{indep}{\sim} \sum_{m=1}^{M} w_m \phi_1(\tilde{t}_i \mid \eta_{i,m,0}(\tau_i, \boldsymbol{x}_i) + s_{i,0}, \sigma_t^2), \\ \tilde{z}_{i,k} \mid \tau_i, \boldsymbol{x}_i, s_{i,k} &\overset{indep}{\sim} \sum_{m=1}^{M} w_m \phi_1(\tilde{z}_{i,k} \mid \eta_{i,m,k}(\tau_i, \boldsymbol{x}_i) + s_{i,k}, \sigma_z^2), \end{aligned} \tag{2}$$

where $w_m / \prod_{m'=1}^{m}(1 - w_{m'}) \overset{iid}{\sim} \text{Be}(1, \xi)$, $m = 1, \ldots, M-1$ and $w_M = 1 - \sum_{m=1}^{M-1} w_m$. To aid in the posterior computation, as is common in mixture models, we introduce auxiliary variables, $r_i \in \{1, \ldots, M\}$, $i = 1, \ldots, n$, which indicate a mixture component for logarithm transformed survival times $\tilde{t}_i$ and probit scores $\tilde{z}_{i,k}$, and let $r_i = m$ with probability w_m. We then define the distribution of $\tilde{t}_i$ and $\tilde{z}_{i,k}$ conditional on the auxiliary variables;

$$\begin{aligned} \tilde{t}_i \mid \tau_i, \boldsymbol{x}_i, s_{i,0}, r_i &\overset{indep}{\sim} \text{N}(\eta_{i,r_i,0}(\tau_i, \boldsymbol{x}_i) + s_{i,0}, \sigma_t^2), \\ \tilde{z}_{i,k} \mid \tau_i, \boldsymbol{x}_i, s_{i,k}, r_i &\overset{indep}{\sim} \text{N}(\eta_{i,r_i,k}(\tau_i, \boldsymbol{x}_i) + s_{i,k}, \sigma_z^2), \end{aligned}$$

where $\eta_{r_i,0}(\tau_i, \boldsymbol{x}_i) = \boldsymbol{\beta}'_{r_i} \tilde{\boldsymbol{x}}$, and $\eta_{i,r_i,k}(\tau_i, \boldsymbol{x}_i) = \boldsymbol{\alpha}'_{k,r_i} \tilde{\boldsymbol{x}}_i$. Also, given r_i, we have

$$\begin{aligned} p(\tilde{t}_i^o, \delta_i \mid \tau_i, \boldsymbol{x}_i, s_{i,k}, r_i) &= \{\phi_1(\tilde{t}_i^o \mid \eta_{i,r_i,0}(\tau_i, \boldsymbol{x}_i) + s_{i,0})\}^{\delta_i} \{1 - \Phi_1(\tilde{t}_i^o \mid \eta_{i,r_i,0}(\tau_i, \boldsymbol{x}_i) + s_{i,0})\}^{(1-\delta_i)}, \\ \text{P}(z_{i,k} = j \mid \tau_i, \boldsymbol{x}_i, s_{i,k}, r_i) &= \int_{u_{k,j}}^{u_{k,j+1}} \phi_1(\tilde{z} \mid \eta_{r_i,k}(\tau_i, \boldsymbol{x}_i) + s_{i,k}, \sigma_z^2) d\tilde{z}. \end{aligned}$$

We then extend the random parameter vector by including $\boldsymbol{c} = (c_1, \ldots, c_n)$, $\boldsymbol{\theta}' = (\boldsymbol{\theta}, \boldsymbol{c})$, and use Markov chain Monte Carlo (MCMC) simulation to generate posterior samples of $(\boldsymbol{\theta}', \boldsymbol{s})$. The full conditionals are given below.

1. $\boldsymbol{s}_i$, $i = 1, \ldots, n$

   The full conditionals of $\boldsymbol{s}_i$ is

   $$p(\boldsymbol{s}_i \mid \mathcal{D}, \boldsymbol{\theta}') \propto \exp\left(-\frac{1}{2}\boldsymbol{s}_i'\Omega\boldsymbol{s}_i\right) p(t_i^o, \delta_i, \boldsymbol{z}_i \mid \boldsymbol{\theta}', \boldsymbol{s}_i).$$

   We update $\boldsymbol{s}_i$ using a random walk Metropolis-Hasting algorithm.

2. Ω

   The full conditionals of Σ_s is

   $$\Sigma_s \mid \mathcal{D}, \boldsymbol{\theta}'_{-\Sigma_s}, \boldsymbol{s} \sim \text{Inv-Wishart}(a_s + n, \sum_{i=1}^{n} \boldsymbol{s}_i'\boldsymbol{s}_i + \Omega^0),$$

   where $\boldsymbol{\theta}'_{-\Omega}$ is a vector of all random parameters, excluding Ω.

3. $\alpha_{m,k,p}$, $k = 1, \ldots, K$, $m = 1, \ldots, M$ and $p = 0, \ldots, P$

   The full conditionals of $\alpha_{m,k,p}$ is

   $$p(\alpha_{m,k,p} \mid \mathcal{D}, \boldsymbol{\theta}'_{-\alpha_{m,k,p}}, \boldsymbol{s}) \propto \prod_{i=1|r_i=m}^{n} p(z_{i,k} \mid \boldsymbol{\theta}', \boldsymbol{s}_i) \exp\left\{-\frac{(\alpha_{m,k,p} - \bar{\alpha}_{k,p})^2}{2v_p^2}\right\},$$

   where $\boldsymbol{\theta}'_{-\alpha_{m,k,p}}$ is a vector of all random parameters, excluding $\alpha_{m,k,p}$. We update $\alpha_{m,k,p}$ using a random walk Metropolis-Hasting algorithm.

4. v_p^2, $p = 0, \ldots, P$

The full conditionals of v_p^2 is

$$v_p^2 \mid \mathcal{D}, \boldsymbol{\theta}'_{-v_p^2} \sim \text{IG}\left(a_v + \frac{KM}{2}, b_v + \frac{\sum_{k,m}(\alpha_{k,m,p} - \bar{\alpha}_{k,p})^2}{2}\right),$$

where $\boldsymbol{\theta}'_{v_p^2}$ is a vector of all random parameters, excluding v_p^2.

5. $\bar{\alpha}_{k,p}$, $k = 1, \ldots, K$ and $p = 0, \ldots, P$

   The full conditionals of $\bar{\alpha}_{k,p}$ is

$$\bar{\alpha}_{k,p} \mid \boldsymbol{\theta}'_{-\bar{\alpha}_{k,p}} \sim \text{N}\left(\left(\frac{1}{v_\alpha^2} + \frac{M}{v_p^2}\right)^{-1}\left(\frac{\bar{\bar{\alpha}}_p}{v_\alpha^2} + \frac{\sum_{m=1}^{M}\alpha_{m,k,p}}{v_p^2}\right), \left(\frac{1}{v_\alpha^2} + \frac{M}{v_p^2}\right)^{-1}\right)$$

   where $\boldsymbol{\theta}'_{-\bar{\alpha}_{k,p}}$ is a vector of all random parameters, excluding $\bar{\alpha}_{k,p}$.

6. $e_{k,j}$, $k = 1, \ldots, K$, $j = 1, \ldots, J-1$

   The full conditionals of $e_{k,j}$ is

$$p(e_{k,j} \mid \mathcal{D}, \boldsymbol{\theta}_{-e_{k,j}}, \boldsymbol{s}) \propto \prod_{i=1 \mid z_{i,k} \geq j}^{n} p(z_{i,k} \mid \boldsymbol{\theta}, \boldsymbol{s}) e_{k,j}^{a_e - 1} \exp\left(-b_e e_{k,j}\right),$$

   where $\boldsymbol{\theta}'_{-e_{k,j}}$ is a vector of all random parameters, excluding $e_{k,j}$. We update $e_{k,j}$ using a random walk Metropolis-Hasting algorithm.

7. $\beta_{m,p}$, $m = 1, \ldots, M$ and $p = 0, \ldots, P$

   The full conditionals of $\beta_{m,p}$ is

$$p(\beta_{m,p} \mid \mathcal{D}, \boldsymbol{\theta}'_{-\beta_{m,p}}, \boldsymbol{s}) \propto \prod_{i=1 \mid r_i = m}^{n} p(\tilde{t}_i^o, \delta_i \mid \boldsymbol{\theta}', \boldsymbol{s}_i) \exp\left\{-\frac{(\beta_{m,p} - \bar{\beta}_p)^2}{2\tau^2}\right\},$$

   where $\boldsymbol{\theta}'_{-\beta_{m,p}}$ is a vector of all random parameters, excluding $\beta_{m,p}$. We update $\beta_{m,p}$ using a random walk Metropolis-Hasting algorithm.

8. σ_t^2

   The full conditionals of σ_t^2 is

$$p(\sigma_t^2 \mid \mathcal{D}, \boldsymbol{\theta}'_{-\sigma_t^2}, \boldsymbol{s}) \propto \prod_{i=1}^{n} p(\tilde{t}_i^o, \delta_i \mid \boldsymbol{\theta}', \boldsymbol{s}_i)(\sigma_t^2)^{a_t - 1} \exp\left(-\frac{b_t}{\sigma_t^2}\right),$$

   where $\boldsymbol{\theta}'_{-\sigma_t^2}$ is a vector of all random parameters, excluding σ_t^2. We update σ_t^2 using a random walk Metropolis-Hasting algorithm.

9. w_m, $m = 1, \ldots, (M-1)$

   The full conditionals of w_m is

$$w_m / \prod_{m'=1}^{m} (1 - w_{m'}) \mid \mathcal{D}, \boldsymbol{\theta}_{-w_m} \sim \text{Be}\left(\sum_{i=1}^{n} 1(r_i = m) + 1, \sum_{i=1}^{n} 1(s_i > m) + M\right),$$

   where $\boldsymbol{\theta}'_{-w_m}$ is a vector of all random parameters, excluding w_m.

10. r_i, $i = 1, \ldots, n$

    The full conditionals of r_i is

$$p(r_i = m \mid \mathcal{D}, \boldsymbol{\theta}_{-r_i}, \boldsymbol{s}) \propto w_m p(\tilde{t}_i^o, \delta_i, \boldsymbol{z}_i \mid \boldsymbol{\theta}'_{-s_i}),$$

    where $\boldsymbol{\theta}'_{-s_i}$ is a vector of all random parameters, excluding s_i.

The Markov chain parameters in the proposal distributions are automatically tuned by adaptive Metropolis-Hastings algorithms (Roberts and Rosenthal, 2009). We checked mixing and convergence of the Markov chain and did not find any evidence of converging to a wrong distribution.

1.2 Computation for Decision Making

We generate a posterior sample of $\boldsymbol{\theta}$ values through posterior MCMC simulation as described in § 1.1, and use it to evaluate quantities needed for subject-specific decision making. Suppose that we consider a subject with $\boldsymbol{x}^{\text{new}} = (x_1^{\text{new}}, \ldots, x_P^{\text{new}})$ assuming a particular τ is given to the patient. One of the covariates is Age, and we have $U_{\text{tot}}(\boldsymbol{y}, \boldsymbol{x}) = U_{\text{tot}}(\boldsymbol{y}, Age)$ with $\boldsymbol{y} = (t, q)$, where t and q are progression free survival (PFS) time and total toxicity burden (TTB), respectively. The posterior predictive distribution of $\tilde{t}$ and $\boldsymbol{z}$ given $\boldsymbol{x}^{\text{new}}$ and τ, provides a basis of the decision making;

$$p(\tilde{t}, \boldsymbol{z} \mid \boldsymbol{x}^{\text{new}}, \mathcal{D}) = \int_{\Theta} \int_{\mathbb{R}^{K+1}} p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \boldsymbol{\theta}, \boldsymbol{s}^{\text{new}}) p(\boldsymbol{s}^{\text{new}} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid \mathcal{D}) d\boldsymbol{s}^{\text{new}} d\boldsymbol{\theta}. \tag{3}$$

Through random variable transformation, we may derive the posterior predictive distribution of $U_{\text{tot}}(t, q, Age^{\text{new}})$ for the new patient from (3), where $t = \exp(\tilde{t})$ and $q = \sum_{k=1}^{K} z_k / \{(J-1)K\}$; for any $u \in [0, U_{\max} = 100]$,

$$P(U_{\text{tot}} \leq u \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D}) = \sum_{\boldsymbol{z}} \int_{A_u(\boldsymbol{z})} p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D}) d\tilde{t}. \tag{4}$$

where $A_u(\boldsymbol{z}) = \{\tilde{t} \mid U_{\text{tot}}(t, q) < u, \text{ where } t = \exp(\tilde{t}) \text{ and } q = \sum z_k / \{(J-1)K\}$. We approximate $p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ and $p(U_{\text{tot}} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ numerically using the posterior samples of $\boldsymbol{\theta}$ simulated from MCMC, and use them to approximate the predictive mean utility;

$$\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}}) = \sum_{z_1=0}^{J-1} \cdots \sum_{z_K=1}^{J-1} \int_{\mathbb{R}} U_{\text{tot}}(t, q, Age^{\text{new}}) p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D}) d\tilde{t}, \tag{5}$$

The computation proceeds as follows;

1. For each $v_1 = 1, \ldots, V_1$ and $v_2 = 1, \ldots, V_2$

   (a) Simulate $\boldsymbol{s}^{(v_1, v_2)}$ from $p(\boldsymbol{s} \mid \boldsymbol{\theta}^{(v_1)})$.

(b) Simulate $(\tilde{t}^{(v_1,v_2)}, \boldsymbol{z}^{(v_1,v_2)})$ from $p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \boldsymbol{s}^{(v_1,v_2)}, \boldsymbol{\theta}^{(v_1)})$.

(c) Evaluate $U_{\text{tot}}(t^{(v_1,v_2)}, q^{(v_1,v_2)}, Age^{\text{new}})$, where $t^{(v_1,v_2)} = \exp(\tilde{t}^{(v_1,v_2)})$ and $q^{(v_1,v_2)} = \sum_{k=1}^{K} z_k^{(v_1,v_2)}/\{(J-1)K\}$.

Using the obtained Monte Carlo sample of U_{tot}, we approximate $\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ in (5) for the patient with covariate $\boldsymbol{x}^{\text{new}}$ and treatment τ

$$\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}}) \approx \frac{1}{V_1 V_2} \sum_{v_1=1}^{V_1} \sum_{v_2=1}^{V_2} U_{\text{tot}}(t^{(v_1,v_2)}, q^{(v_1,v_2)}, Age^{\text{new}}). \tag{6}$$

An alternative way of comparing a pair of treatments (τ, τ'), $\tau \neq \tau'$ of treatments is to compute the posterior probability $\Delta(\boldsymbol{x}^{\text{new}})$ that treatment τ' has a larger mean total utility than treatment τ for a patient with prognostic covariates $\boldsymbol{x}^{\text{new}}$,

$$\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau') = \Pr\left(U_{\text{tot}}(t(\tau), q(\tau), Age^{\text{new}}) < U_{\text{tot}}(t(\tau'), q(\tau'), Age^{\text{new}}) \mid \tau, \tau', \boldsymbol{x}^{\text{new}}, \mathcal{D}\right). \tag{7}$$

Similar to $\bar{u}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$, we numerically approximate $\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau')$ using a MCMC sample of $\boldsymbol{\theta}$ as follows;

1. For each $v_1 = 1, \ldots, V_1$ and $v_2 = 1, \ldots, V_2$

   (a) Simulate $\boldsymbol{s}^{(v_1,v_2)}$ from $p(\boldsymbol{s} \mid \boldsymbol{\theta}^{(v_1)})$.

   (b) Simulate $(\tilde{t}^{(v_1,v_2)}(\tau), \boldsymbol{z}^{(v_1,v_2)}(\tau))$ from $p(\tilde{t}, \boldsymbol{z} \mid \tau, \boldsymbol{x}^{\text{new}}, \boldsymbol{s}^{(v_1,v_2)}, \boldsymbol{\theta}^{(v_1)})$. Similarly, simulate $(\tilde{t}^{(v_1,v_2)}(\tau'), \boldsymbol{z}^{(v_1,v_2)}(\tau'))$ from $p(\tilde{t}, \boldsymbol{z} \mid \tau', \boldsymbol{x}^{\text{new}}, \boldsymbol{s}^{(v_1,v_2)}, \boldsymbol{\theta}^{(v_1)})$.

   (c) Compute

      i. $t^{(v_1,v_2)}(\tau) = \exp(\tilde{t}^{(v_1,v_2)}(\tau))$ and $q^{(v_1,v_2)}(\tau) = \sum_{k=1}^{K} z_k^{(v_1,v_2)}(\tau)/\{(J-1)K\}$

      ii. $t^{(v_1,v_2)}(\tau') = \exp(\tilde{t}^{(v_1,v_2)}(\tau'))$ and $q^{(v_1,v_2)}(\tau') = \sum_{k=1}^{K} z_k^{(v_1,v_2)}(\tau')/\{(J-1)K\}$

   (d) Evaluate

$$\Delta(^{(v_1,v_2)}\boldsymbol{x}^{\text{new}}, \tau, \tau') = 1(U_{\text{tot}}(t^{(v_1,v_2)}(\tau), q^{(v_1,v_2)}(\tau), Age^{\text{new}}) < U_{\text{tot}}(t^{(v_1,v_2)}(\tau'), q^{(v_1,v_2)}(\tau'), Age^{\text{new}})),$$

where $1(a)$ is the indicator function.

Using the obtained Monte Carlo sample of $\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau')$, we approximate the probability $\Delta(\boldsymbol{x}^{\text{new}})$ for the patient with covariate $\boldsymbol{x}^{\text{new}}$ and treatments (τ, τ') as follows;

$$\Delta(\boldsymbol{x}^{\text{new}}, \tau, \tau') \approx \frac{1}{V_1 V_2} \sum_{v_1=1}^{V_1} \sum_{v_2=1}^{V_2} \Delta^{(v_1,v_2)}(\boldsymbol{x}^{\text{new}}, \tau, \tau'). \tag{8}$$

2 Illustrative Examples of Utilities

A first example is a utility function for a single outcome that does not vary with $\boldsymbol{x}$. Utilities of $y =$ ordinal post-operative morbidity (POM) category were elicited from a surgeon planning a randomized trial comparing two treatments for patients with esophageal cancer undergoing chemoradiation and surgery (Murray et al., 2018). As often done when dealing with ordinal variables, y was defined as integer values used nominally to identify the six POM levels, with $y = 0$ for normal recovery, 1 for minor complications, 2 for complications requiring pharmaceutical intervention, 3 for complications requiring surgical, endoscopic, or radiological intervention, 4 for life-threatening complications requiring intensive care, and 5 for death. To reflect relative desirability and obtain a meaningful quantitative representation of the POM categories, utilities were established by first fixing $U(0) = 100$ and $U(5) = 0$ and then eliciting the intermediate values from the surgeon, $U(1) = 80$, $U(2) = 65$, $U(3) = 25$, and $U(4) = 10$. These numerical utilities reflect the subjective view that POM categories with scores $y \leq 2$ are much more desirable than those with scores $y \geq 3$, that $y = 1$ and 2 have similar utilities of 80 and 65, and that $y = 4$ is nearly as bad as death. Although one utility was used for all $\boldsymbol{x}$, the distribution of y may vary by $\boldsymbol{x}$ and τ, potentially leading to decisions of choosing optimal treatments that vary with $\boldsymbol{x}$. The estimate $\hat{\beta}_\tau$ of the effect of τ or estimated mean raw POM score $\overline{y} = \sum_{y=0}^{5} y\, p(y \mid \tau, \boldsymbol{x}, \hat{\boldsymbol{\theta}})$, where $\hat{\boldsymbol{\theta}}$ is an estimate of $\boldsymbol{\theta}$, that a typical statistical analysis reports, may not be meaningful and potentially misleading for treatment

selection because the numerical values of y are nominal. To use the elicited utility for treatment comparison, a frequentist approach may use the mean utility estimate $\overline{U}(\tau, \boldsymbol{x}, \hat{\boldsymbol{\theta}}) = \sum_{y=0}^{5} U(y)p(y \mid \tau, \boldsymbol{x}, \hat{\boldsymbol{\theta}})$ of giving treatment τ to a patient with covariates $\boldsymbol{x}$. The Bayesian framework allows one to use the posterior predictive distribution of y, $p(y \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ for a new patient with covariates $\boldsymbol{x}^{\text{new}}$. Similarly to $\overline{U}(\tau, \boldsymbol{x}^{\text{new}}, \hat{\boldsymbol{\theta}})$, the posterior predictive mean utility $\bar{u}(\tau, \boldsymbol{x}^{\text{new}}) = \sum_{y=0}^{5} U(y)p(y \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ can be used as a quantitative criterion for comparison, with $\bar{u}(\tau, \boldsymbol{x}^{\text{new}})$ computed for each τ and $\tau^{\text{opt}} = \arg\max_{\tau} \bar{u}(\tau, \boldsymbol{x}^{\text{new}})$ chosen as an optimal treatment for the patient. In contrast with $\overline{U}(\hat{\tau}, \boldsymbol{x}, \boldsymbol{\theta})$, $\bar{u}(\tau, \boldsymbol{x}^{\text{new}})$ accounts for uncertainty about $\boldsymbol{\theta}$ through the posterior distribution.

A utility function that characterizes the trade-off between two binary outcomes, $\boldsymbol{y} = (y_1, y_2)$, is useful in settings where y_1 is the indicator of efficacy and y_2 is the indicator of toxicity. As done for the one-dimensional POM score above, one may begin by first assigning $U(1,0) = 100$ to the best possible outcome pair (efficacy, no toxicity), and $U(0,1) = 0$ to the worst possible outcome pair (toxicity, no efficacy). Subject to the constraints $0 < U(0,0), U(1,1) < 100$ required for consistency, one may assign numerical values between 0 and 100 to the events (0,0) and (1,1) that reflect their comparative desirability. This will depend on the medical setting and particular definitions of y_1 and y_2. For example, $U(0,0) = 40$ and $U(1,1) = 60$ reflect a large cost for toxicity relative to efficacy, whereas $U(0,0) = 20$ and $U(1,1) = 75$ reflect greater value of efficacy as a trade-off for toxicity. Similar to the earlier example, we can use the posterior predictive mean utility $\bar{u}(\tau, x^{\text{new}}) = \sum_{y_1=0}^{1} \sum_{y_2=0}^{1} U(\boldsymbol{y})p(\boldsymbol{y} \mid \tau, x^{\text{new}}, \mathcal{D})$ to choose an optimal treatment for a future patient with prognostic state x^{new}. This may be elaborated by defining different utilities for patients with Good or Poor prognosis, i.e., $U(\boldsymbol{y}, x)$ with x denoting prognostic state. One possibility is $U(0,0,\mathit{Good}) = 50$, $U(1,1,\mathit{Good}) = 80$, and $U(0,0,\mathit{Poor}) = 40$, $U(1,1,\mathit{Poor}) = 70$.

The next example is a utility for a bivariate ordinal outcome. A phase I-II trial was conducted to optimize radiation dose for palliative treatment of pediatric patients with

diffuse intrinsic pontine gliomas. The design was described by Thall and Nguyen (2012), and trial results were reported by Amsbaugh et al. (2019). Efficacy was defined as the number of improvements in clinical symptoms, radiographic appearance of the tumor, or quality of life, so $y_1 = 0$, 1, 2, or 3. Toxicity was defined in terms of fatigue, nausea/vomiting, headache, skin inflammation/desquamation, blindness, and brain edema or necrosis, as $y_2 =$ low, moderate, high, or severe. Both outcomes were scored within 42 days from the start of therapy. The elicited utilities are given in Table 1. Initially there were three toxicity levels, but consideration of the most severe toxicities caused the clinicians to split the High/Severe category into two levels, with Severe including brain necrosis and blindness. To see the usefulness of the numerical utilities in Table 1, suppose that, instead of conducting a phase I-II trial, a phase I trial were conducted based only on "dose-limiting toxicity" (DLT) defined by the binary variable y_{DLT} =1 for High or Severe toxicity, and 0 otherwise, while ignoring the patient's efficacy score y_2. Then $(y_1, y_2) =$ (Moderate,0) would be scored as $y_{DLT} = 0$, while $(y_1, y_2) =$ (High,3) would be scored as $y_{DLT} = 1$, despite the fact that they have the same utility of 25. This illustrates how, once two competing outcomes, such as efficacy and toxicity, are considered together and numerical utilities are assigned to quantify the desirability of each possible combination, it becomes obvious why a conventional procedure based on one outcome but ignoring the other may lead to decisions and actions that do not make sense. This will be the case in our illustrative data analysis, given below, with regard to whether PFS and TTB are considered together or separately.

3 Additional Details of Utility Function Elicitation for Breast Cancer Data

Recall that the total utility of outcomes $\boldsymbol{y} = (t, q)$ for a patient with Age is constructed as

$$U_{\text{tot}}(t, q, Age) = U_{\text{PFS}}(t)U_{\text{TTB}}(q, Age), \tag{9}$$

where $0 < U_{\text{PFS}}(t) \leq U_{\max}$, with $U_{\max} = 100$, and $0 < U_{\text{TTB}}(q, Age) \leq 1$. The construction of $U_{\text{PFS}}(t, Age)$ is shown in Eq (9). We constructed

$$U_{\text{TTB}}(q, Age) = \exp\{-q^2/(2g^2(Age))\}, \tag{10}$$

where $g(\tilde{x}) = \exp(c_0 + c_1 Age)$. We then calibrated the values of c_0 and c_1 for a close approximation of $U_{\text{tot}}(t, q, Age)$ to the elicited utilities as follows; Tab 2 has elicited $U_{\text{tot}}(t, q, Age)$ at $Age = 40, 50, 65$, and 85 for a grid of t and q, provided by clinicians. We set $g(Age)$=0.5, 0.15, 0.1 and 0.05 for $Age = 40, 50, 65$ and 85, respectively, which are the largest values of q having a non-zero utility. We then regress $\log(g)$ on Age (to ensure $g > 0$), and set c_0 and c_1 to be the obtained regression coefficients. The resulting function is $g = \exp(0.823 - 0.047Age)$. Since $c_1 < 0$, $g(Age)$ decreases in Age, and $U_{\text{TTB}}(q, Age)$ decreases faster as q increases for a large value of Age. In other words, occurrences of AEs are worse for older patients, and an increase in q more greatly decrease U_{tot} for older patients.

In general, $U_{\text{tot}}(t, q, Age)$ can be derived in many different ways. Different functions for $U_{\text{TTB}}(q, Age)$ were explored as well, such as exponential functions of q and power functions of q. We emphasize that close collaboration with subject experts is essential in constructing a reasonable U_{tot}.

4 Additional Details of the Breast Cancer Data Analysis

We compare posterior predictive distributions of individual outcomes, PFS (t) and TTB (q), for treatments L and L+B for more values of $\boldsymbol{x}^{\text{new}}$. Panel (a) of Figs 1 illustrates estimated posterior predictive probability, $\hat{\text{P}}(t(L, \boldsymbol{x}^{\text{new}}) < t(L+B, \boldsymbol{x}^{\text{new}}) \mid \mathcal{D})$, that treatment L+B has greater PFS than treatment L for different cases of $\boldsymbol{x}^{\text{new}} = (x_1^{\text{new}}, x_2^{\text{new}}, x_3^{\text{new}})$, where x_1^{new}=age, x_2^{new} = measurable disease, and x_3^{new} = disease free interval prior to trial entry > 24 months for an unobserved patient. Panels (b) and (c) of the figure show posterior predictive median PFS estimates $\hat{t}(\tau, \boldsymbol{x}^{\text{new}})$ for treatments τ =L and L+B, respectively, for different values of $\boldsymbol{x}^{\text{new}}$. Similarly, panel (a) of 2 illustrates estimated posterior predictive probabilities, $\hat{\text{P}}(q(L, \boldsymbol{x}^{\text{new}}) < q(L+B, \boldsymbol{x}^{\text{new}}) \mid \mathcal{D})$, that treatment L+B has greater TTB (q) than treatment L for different cases of $\boldsymbol{x}^{\text{new}}$. Panels (b) and (c) of the figure give posterior predictive mean TTB estimates $\hat{t}(\tau, \boldsymbol{x}^{\text{new}})$ for treatments τ =L and L+B, respectively, for different values of $\boldsymbol{x}^{\text{new}}$.

5 Additional Details of Simulation Study

Recall that we have the number of patients to be $n = 350$ and $K = 20$ toxicity types each having $J = 6$ grades. We include three covariates $(\tau_i, x_{i,1}, x_{i,2})$, where τ, x_1 and x_2 are a binary treatment indicator, age, and a binary indicator of disease measurability, respectively, so our $P = 2$. We generated patient-specific frailty vectors $\boldsymbol{s}_i^{\text{TR}} \overset{iid}{\sim} \text{N}_{K+1}(\boldsymbol{0}, \Omega^{\text{TR}})$, where Ω^{TR} assumes that variances of $s_{i,k}^{\text{TR}}$ are 0.05 and correlations between $s_{i,k}^{\text{TR}}$ and $s_{i,k'}^{\text{TR}}$ are 0.5 if $k, k' > 0$ and $k \neq k'$, and -0.5 if k or $k' = 0$. We simulated t_i and $\tilde{z}_{i,k}$ from a mixture of two normals. We first simulated the membership of patient i, $\gamma_i^{\text{TR}} \overset{iid}{\sim}$ Categorical$(w_1^{\text{TR}}, w_2^{\text{TR}})$,

where $w_1^{\text{TR}} = 0.4$ and $w_2^{\text{TR}} = 0.6$, and generated $\tilde{t}_i$ and $\tilde{z}_{i,k}^{\text{TR}}$ conditional on γ_i^{TR}; if $\gamma_i^{\text{TR}} = 1$,

$$\begin{aligned}
\tilde{t}_i \mid \boldsymbol{x}_i, s_{i,0}^{\text{TR}} &\overset{indep}{\sim} \text{N}(4.0 - 0.5\tau_i + 0.5x_{i,1} + 0.5\tau_i x_{i,1} + s_{i,0}^{\text{TR}}, 0.6), \\
\tilde{z}_{i,k} \mid \boldsymbol{x}_i, s_{i,k}^{\text{TR}}, \boldsymbol{\alpha}^{\text{TR}} &\overset{indep}{\sim} \text{N}(\alpha_{1,k,0}^{\text{TR}} + \alpha_{1,k,1}^{\text{TR}}\tau_i + \alpha_{1,k,2}^{\text{TR}}x_{i,1} + \alpha_{1,k,3}^{\text{TR}}\tau_i x_{i,1} + s_{i,k}^{\text{TR}}, 4),
\end{aligned}$$

where $\alpha_{1,k,0}^{\text{TR}} \overset{iid}{\sim} \text{U}(-4.5, -4.0)$, $\alpha_{1,k,1}^{\text{TR}} \overset{iid}{\sim} \text{U}(0.5, 1.0)$, $\alpha_{1,k,2}^{\text{TR}} \overset{iid}{\sim} \text{U}(0.3, 0.8)$, and $\alpha_{1,k,3}^{\text{TR}} \overset{iid}{\sim} \text{U}(0, 0.5)$ were simulated for all k. Otherwise, we generated

$$\begin{aligned}
\tilde{t}_i \mid \boldsymbol{x}_i, s_{i,0}^{\text{TR}} &\overset{indep}{\sim} \text{N}(1.5 + 0.8\tau_i - 0.5x_{2,i} + s_{i,0}^{\text{TR}}, 0.6) \\
\tilde{z}_{i,k} \mid \boldsymbol{x}_i, s_{i,k}^{\text{TR}}, \boldsymbol{\alpha}^{\text{TR}} &\overset{indep}{\sim} \text{N}(\alpha_{2,k,0}^{\text{TR}} + \alpha_{2,k,1}^{\text{TR}}\tau_i + \alpha_{2,k,2}^{\text{TR}}x_{i,2} + s_{i,k}^{\text{TR}}, 4),
\end{aligned}$$

where $\alpha_{2,k,0}^{\text{TR}} \overset{iid}{\sim} \text{U}(-4.5, -4.0)$, $\alpha_{2,k,1}^{\text{TR}} \overset{iid}{\sim} \text{U}(1.0, 1.5)$, and $\alpha_{2,k,2}^{\text{TR}} \overset{iid}{\sim} \text{U}(0.5, 1.0)$ were simulated for all k. We simulated censoring times as $c_i \overset{iid}{\sim} \text{N}(4.63, 2)$ and let $t_i^o = \min(t_i, c_i)$ and $\delta_i = 1(c_i < t_i)$, resulting in 22.86% of t_i's being censored. For ordinal outcomes $z_{i,k}$, we also simulated cutoff points $u_{k,j}^{\text{TR}} \overset{iid}{\sim} \text{N}(\bar{u}_j, 0, 01^2)$, $j = 2, \ldots, J-1$, where $\bar{u}_j = 0.7, 1.3, 2.0, 3.0, 4.0$. Finally, we let $z_{i,k} = j$ if $u_{k,j-1}^{\text{TR}} < \tilde{z}_{i,k}^{\text{TR}} \leq u_{k,j}^{\text{TR}}$.

References

Amsbaugh, M., Mahajan, A., Thall, P., McAleer, M., Paulino, A., Grosshans, D., Khatua, S., Ketonen, L., Fontanilla, S., and McGovern, S. (2019). A phase 1/2 trial of reirradiation for diffuse intrinsic pontine gliomas. *International Journal of Radiation Oncology, Biology, Physics*, 104:144–148.

Murray, T. A., Yuan, Y., Thall, P. F., Elizondo, J. H., and Hofstetter, W. L. (2018). A utility-based design for randomized comparative trials with ordinal outcomes and prognostic subgroups. *Biometrics*, 74(3):1095–1103.

Roberts, G. O. and Rosenthal, J. S. (2009). Examples of adaptive mcmc. *Journal of Computational and Graphical Statistics*, 18(2):349–367.

Thall, P. F. and Nguyen, H. Q. (2012). Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of biopharmaceutical statistics*, 22(4):785–801.

		Toxicity Severity			
		Mild	Moderate	High	Severe
	0	50	25	10	0
Efficacy	1	85	50	15	5
Score	2	92	60	20	7
	3	100	75	25	10

Table 1: Elicited utilities of the 16 possible outcomes (y_1, y_2) = (Toxicity Severity, Efficacy Score) in the DIPG radiation dose-finding trial.

(a) Age 40 years

TTB (q)	PFS (t)				
	0	12	24	48	∞
0.000	0	25	50	95	100
0.025	0	25	50	95	100
0.050	0	23	49	95	100
0.100	0	20	48	93	100
0.150	0	17	44	90	100
0.500	0	5	32	80	100
1.000	0	0	0	0	0

(b) Age 50 years

PFS (t)				
0	12	24	48	∞
0	25	50	95	100
0	23	49	95	100
0	20	48	93	100
0	17	44	90	100
0	5	32	80	100
0	0	0	0	0
0	0	0	0	0

(c) Age 65 years

TTB (q)	PFS (t)				
	0	12	24	48	∞
0.000	0	25	50	95	100
0.025	0	20	48	93	100
0.050	0	17	44	90	100
0.100	0	5	32	80	100
0.150	0	0	0	0	0
0.500	0	0	0	0	0
1.000	0	0	0	0	0

(d) Age 85 years

PFS (t)				
0	12	24	48	∞
0	25	50	95	100
0	17	44	90	100
0	5	32	80	100
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0

Table 2: [**Elicited Utilities**]. Values of $U_{\text{tot}}(\boldsymbol{y}, \tilde{x})$ at age x=40, 50, 65, and 85 years. Outcomes $\boldsymbol{y} = (t, q)$ consist of PFS (t) and scaled TTB (q).

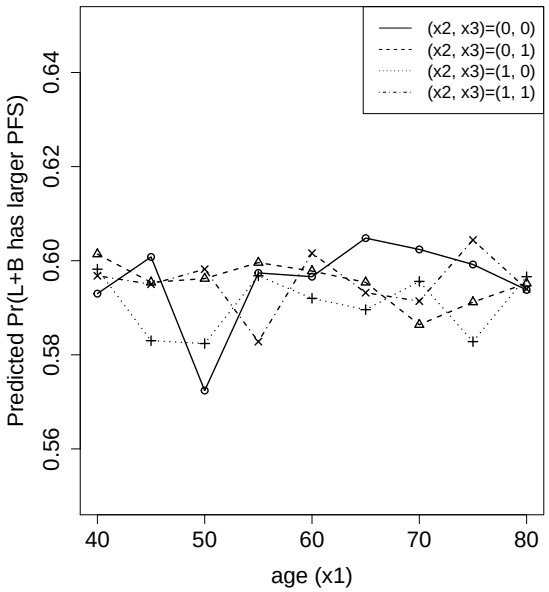

(a) $\hat{\mathrm{P}}(t(L, \boldsymbol{x}^{\mathrm{new}}) < t(L+B, \boldsymbol{x}^{\mathrm{new}}) \mid \mathcal{D})$

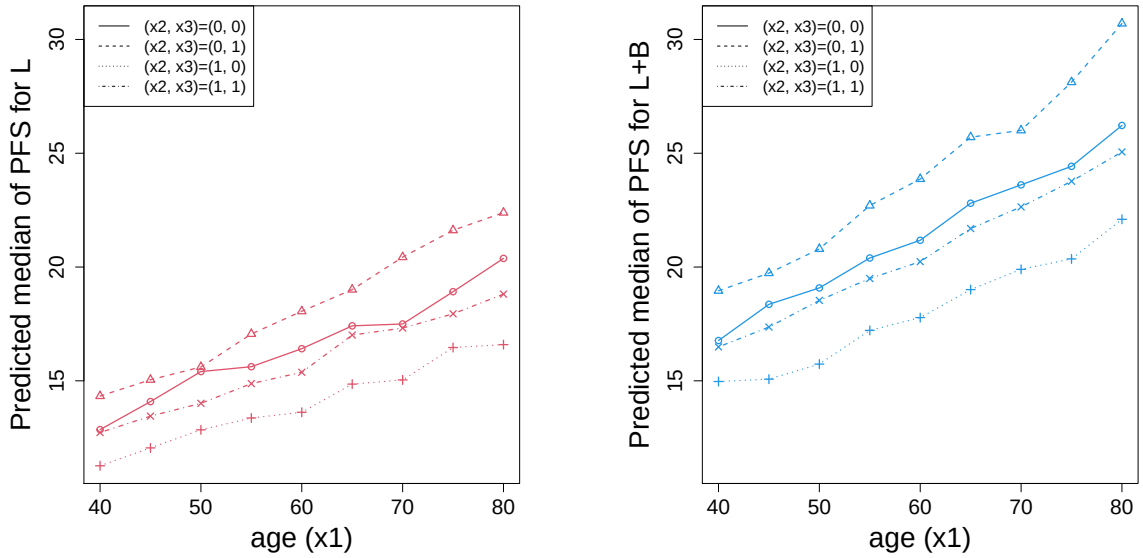

(b) $\hat{t}(\tau, \boldsymbol{x}^{\mathrm{new}})$ with $\tau = L$

(c) $\hat{t}(\tau, \boldsymbol{x}^{\mathrm{new}})$ with $\tau = L + B$

Figure 1: [Breast Cancer Trial Data] (a) Estimated posterior predictive (PP) probability, $\hat{\mathrm{P}}(t(L, \boldsymbol{x}^{\mathrm{new}}) < t(L+B, \boldsymbol{x}^{\mathrm{new}}) \mid \mathcal{D})$, that treatment L+B has greater PFS than treatment L for different cases of $\boldsymbol{x}^{\mathrm{new}} = (x_1^{\mathrm{new}}, x_2^{\mathrm{new}}, x_3^{\mathrm{new}})$, where x_1^{new}=age, x_2^{new} = measurable disease, and x_3^{new} = disease free interval prior to trial entry > 24 months for an unobserved patient. Panels (b) and (c) give posterior predictive median PFS estimates $\hat{t}(\tau, \boldsymbol{x}^{\mathrm{new}})$ for treatments $\tau =L$ and L+B, respectively, for different values of $\boldsymbol{x}^{\mathrm{new}}$.

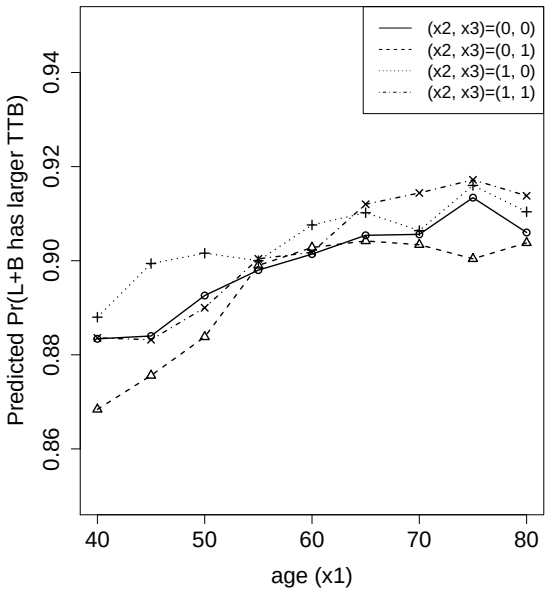

(a) $\hat{\mathrm{P}}(q(L, \boldsymbol{x}^{\mathrm{new}}) < q(L+B, \boldsymbol{x}^{\mathrm{new}}) \mid \mathcal{D})$

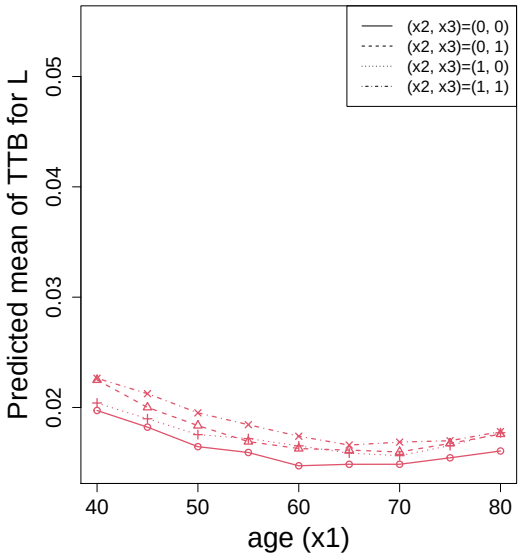

(b) $\hat{q}(\tau, \boldsymbol{x}^{\mathrm{new}})$ with $\tau = L$

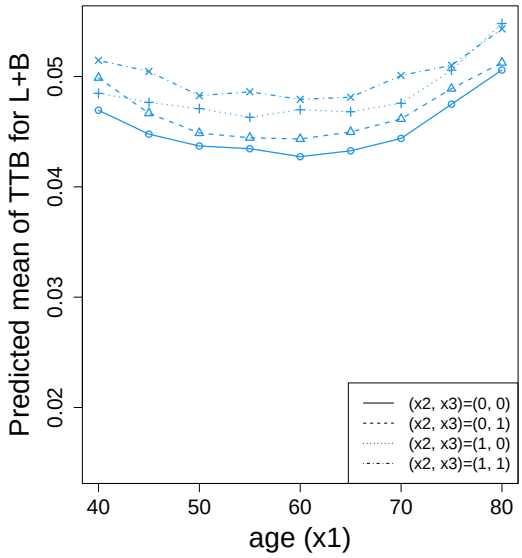

(c) $\hat{t}(\tau, \boldsymbol{x}^{\mathrm{new}})$ for $\tau = L + B$

Figure 2: [Breast Cancer Trial Data] (a) Estimated posterior predictive (PP) probability, $\hat{\mathrm{P}}(q(L, \boldsymbol{x}^{\mathrm{new}}) < q(L+B, \boldsymbol{x}^{\mathrm{new}}) \mid \mathcal{D})$, that treatment L+B has greater TTB (q) than treatment L for different cases of $\boldsymbol{x}^{\mathrm{new}} = (x_1^{\mathrm{new}}, x_2^{\mathrm{new}}, x_3^{\mathrm{new}})$, where x_1^{new}=age, x_2^{new} = measurable disease, and x_3^{new} = disease free interval prior to trial entry > 24 months for an unobserved patient. Panels (b) and (c) give posterior predictive mean TTB estimates $\hat{t}(\tau, \boldsymbol{x}^{\mathrm{new}})$ for treatments $\tau =L$ and L+B, respectively, for different values of $\boldsymbol{x}^{\mathrm{new}}$.

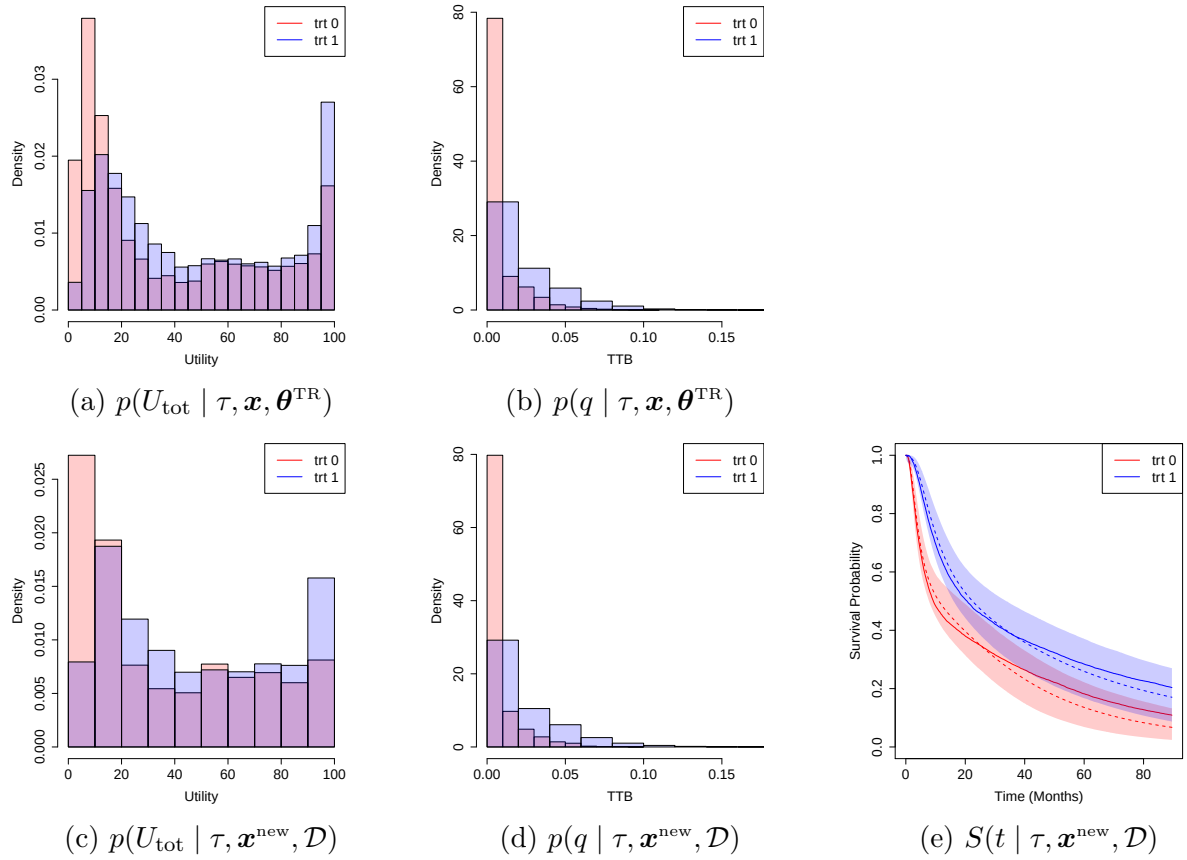

(a) $p(U_{\text{tot}} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$ (b) $p(q \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$

(c) $p(U_{\text{tot}} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ (d) $p(q \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$ (e) $S(t \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$

Figure 3: [Simulation Study] $p(U_{\text{tot}} \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$, $p(q \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$ and $S(t \mid \tau, \boldsymbol{x}, \boldsymbol{\theta}^{\text{TR}})$ (solid lines) are shown in panels (a), (b) and (e), respectively. $x_1 = 65$ and $x_2 = 0$ are assumed, and each panel gives distributions with $\tau = 0$ in red solid line and $\tau = 1$ in blue solid line. Posterior predictive distributions, $p(U_{\text{tot}} \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$, $p(q \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$, and $S(t \mid \tau, \boldsymbol{x}^{\text{new}}, \mathcal{D})$, for an unobserved patient with $\boldsymbol{x}^{\text{new}} = (65, 0)$ are illustrated in panels (c), (d) and (e). In panel (e), the shaded areas represent 95% pointwise credible intervals, and the dashed lines are posterior mean estimates.

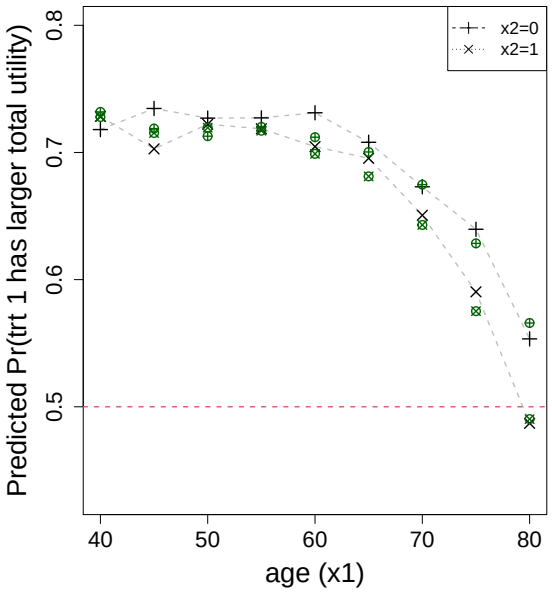

(a) $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, 0, 1)$

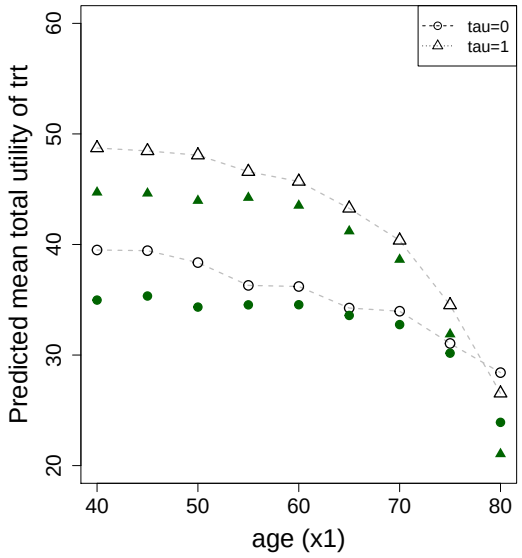

(b) $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ for $x_2^{\text{new}} = 0$

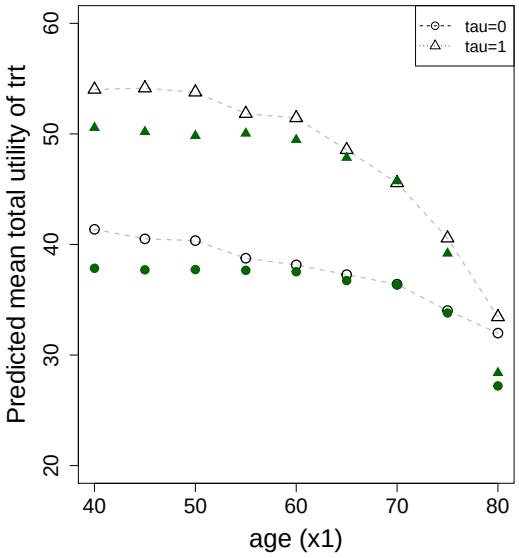

(c) $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ for $x_2^{\text{new}} = 1$

Figure 4: [Simulation Study] (a) Estimated posterior predictive probabilities, $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, 0, 1)$, that U_{tot} with $\tau = 1$ is greater than that with $\tau = 0$, for a patient with prognostic covariates $\boldsymbol{x}^{\text{new}}$. A grid from 40 years to 80 years by 5 years is used for the covariate $x_1^{\text{new}} = Age$, and different values of x_2^{new} are used. In panels (b) and (c), the posterior expected utilities are computed for treatments $\tau = 0$ and 1, as functions of Age. In the plots, solid dark green symbols represent their corresponding true values.

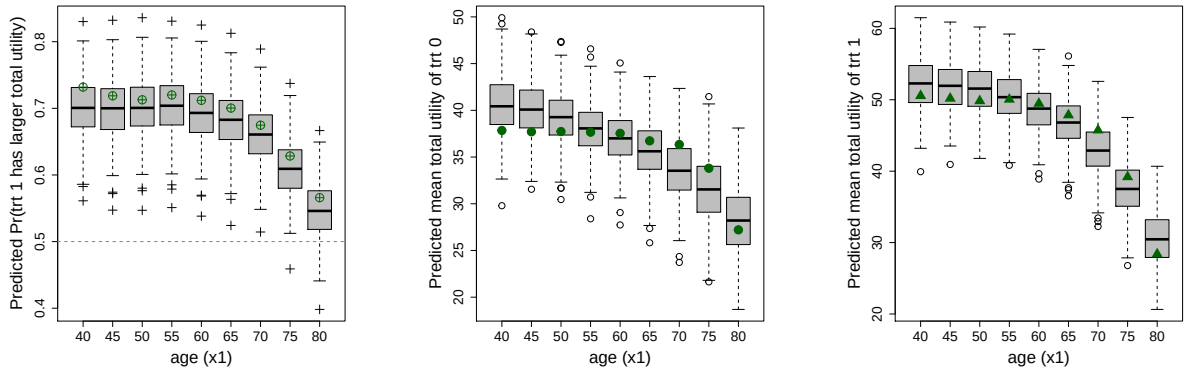

(a) $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, 0, 1)$, $x_2^{\text{new}} = 0$ (b) $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ $(\tau, x_2^{\text{new}}) = (0, 0)$ (c) $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ $(\tau, x_2^{\text{new}}) = (1, 0)$

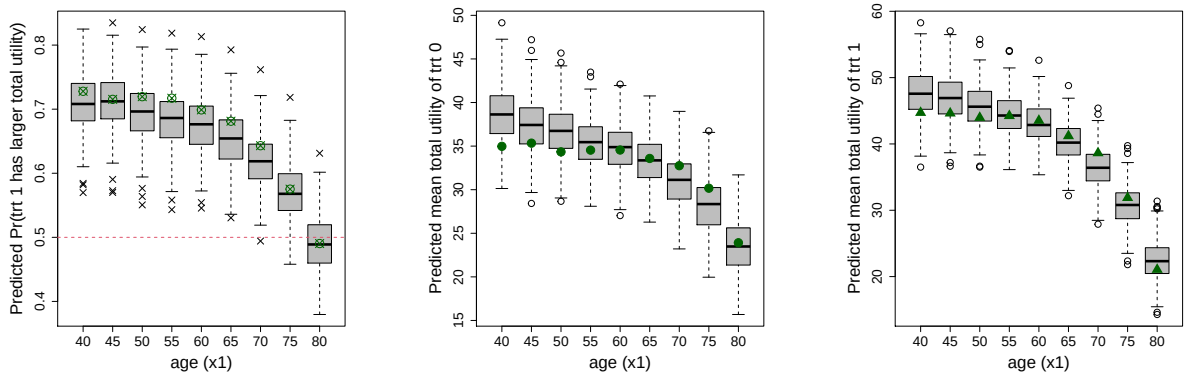

(d) $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, 0, 1)$, $x_2^{\text{new}} = 1$ (e) $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ $(\tau, x_2^{\text{new}}) = (0, 1)$ (f) $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$ $(\tau, x_2^{\text{new}}) = (1, 1)$

Figure 5: [Simulation Study] Distributions of estimated posterior predictive probabilities, $\hat{\Delta}(\boldsymbol{x}^{\text{new}}, 0, 1)$, based on 100 simulated datasets, are shown in panels (a) and (d) for $x_2^{\text{new}} = 0$ and 1, respectively. A grid of $x_1^{\text{new}} = Age$ is used for the evaluation. Distributions of posterior predictive mean utility estimates, $\hat{\bar{u}}_{\text{tot}}(\tau, \boldsymbol{x}^{\text{new}})$, are in panels (b), (c), (e) and (f) for combinations of (τ, x_2^{new}) with $\tau, x_2^{\text{new}} \in \{0, 1\}$.