

THE USE OF COVARIATES AND RANDOM EFFECTS IN EVALUATING PREDICTIVE BIOMARKERS UNDER A POTENTIAL OUTCOME FRAMEWORK

BY ZHIWEI ZHANG*, LEI NIE[†], GUOXING SOON[†] AND AIYI LIU[‡]

*Center for Devices and Radiological Health, Food and Drug Administration**,
Center for Drug Evaluation and Research, Food and Drug Administration[†]
and Eunice Kennedy Shriver National Institute
of Child Health and Human Development[‡]

Predictive or treatment selection biomarkers are usually evaluated in a subgroup or regression analysis with focus on the treatment-by-marker interaction. Under a potential outcome framework (Huang, Gilbert and Janes [*Biometrics* **68** (2012) 687–696]), a predictive biomarker is considered a predictor for a desirable treatment benefit (defined by comparing potential outcomes for different treatments) and evaluated using familiar concepts in prediction and classification. However, the desired treatment benefit is unobservable because each patient can receive only one treatment in a typical study. Huang et al. overcome this problem by assuming monotonicity of potential outcomes, with one treatment dominating the other in all patients. Motivated by an HIV example that appears to violate the monotonicity assumption, we propose a different approach based on covariates and random effects for evaluating predictive biomarkers under the potential outcome framework. Under the proposed approach, the parameters of interest can be identified by assuming conditional independence of potential outcomes given observed covariates, and a sensitivity analysis can be performed by incorporating an unobserved random effect that accounts for any residual dependence. Application of this approach to the motivating example shows that baseline viral load and CD4 cell count are both useful as predictive biomarkers for choosing antiretroviral drugs for treatment-naïve patients.

1. Introduction. Much of contemporary medical research is focused on treatment effect heterogeneity, that is, the fact that the same treatment can have different effects on different patients. The increasing awareness of treatment effect heterogeneity has motivated the development of predictive biomarkers for identifying the subpopulation of patients who would actually benefit from a new treatment [e.g., Simon (2008, 2010)]. Classical examples of predictive biomarkers include genetic markers for cancer treatment, such as the OncoType Dx multi-gene score for breast cancer [Paik et al. (2004)] and the K-RAS gene expression level for colorectal cancer [Karapetis et al. (2008)]. This article is motivated by a new

Received January 2013; revised April 2014.

Key words and phrases. Conditional independence, counterfactual, ROC regression, sensitivity analysis, treatment effect heterogeneity, treatment selection.

and growing interest in the possibility of using baseline viral load or CD4 cell count as a predictive biomarker for treating human immunodeficiency virus type 1 (HIV-1).

Evaluation of a predictive biomarker is usually based on a subgroup or regression analysis comparing treatment effects on different subpopulations defined by the biomarker [e.g., Gail and Simon (1985), Pocock et al. (2002), Russek-Cohen and Simon (1998)]. Under this approach, the performance of a predictive biomarker is measured by the interaction between marker value and treatment assignment in a regression model for the clinical outcome of interest, which will be referred to as an outcome model. For example, consider the THRIVE study, a phase 3, randomized, noninferiority trial comparing rilpivirine, a newly developed nonnucleoside reverse transcriptase inhibitor, with efavirenz in treatment-naïve adults infected with HIV-1 [Cohen et al. (2011)]. The outcome of primary interest is a binary indicator of virologic response at week 48 of treatment (see Section 4 for details). Our consideration of baseline viral load and CD4 cell count as predictive biomarkers (for choosing between efavirenz and rilpivirine) is motivated by Figure 1, which shows nonparametric estimates of treatment-specific response rates as functions of marker value, separately for each biomarker. Figure 1 suggests a qualitative interaction between treatment and each biomarker, with rilpivirine favored over efavirenz for higher values of baseline CD4 cell count and lower values of baseline viral load. Because the statistical significance in Figure 1 is not straightforward to assess, a simple logistic regression analysis that includes treatment, marker and their interaction is performed (separately for each marker), and the resulting p -value for the treatment-by-marker interaction is 0.059 for viral load and 0.031 for CD4 cell count.

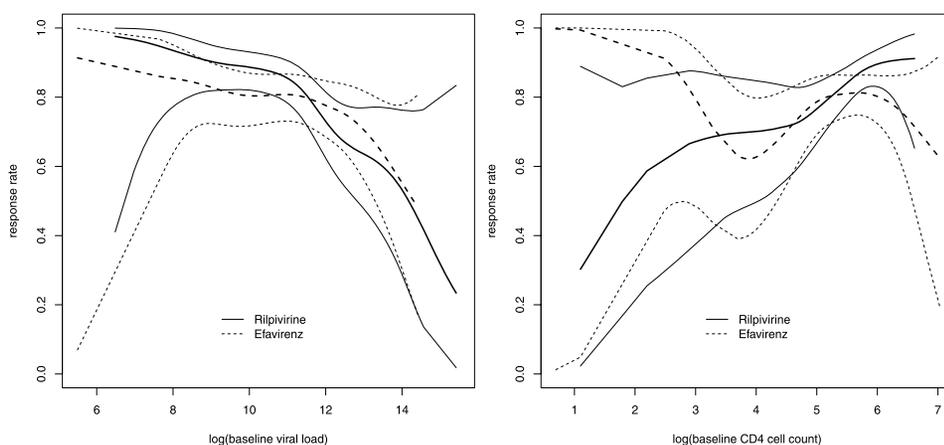


FIG. 1. *Nonparametric regression analysis of the THRIVE data: smoothed estimates (thicker lines) and pointwise 95% confidence intervals (thinner lines) for treatment-specific response rates as functions of baseline viral load (left) and CD4 cell count (right).*

While informative about possible interactions, Figure 1 is less transparent about the predictive performance of these biomarkers and their comparison. Huang, Gilbert and Janes (2012) point out that a strong interaction is not sufficient for adequate performance of a predictive biomarker, that the scale of the interaction coefficient depends on the functional form of the outcome model, and that the interaction-based approach is ill-suited for developing combination markers. These authors also propose a potential outcome framework where a predictive biomarker is—as the term suggests—treated as a predictor for a desirable treatment benefit. Note that a treatment benefit is necessarily the result of comparing potential outcomes for different treatments applied to the same patient. Under this perspective, predictive biomarkers should be evaluated using familiar measures in prediction and classification [e.g., Pepe (2003), Zhou, Obuchowski and McClish (2002), Zou et al. (2011)]. Specifically, one should consider the true and false positive rates of a binary marker and the receiver operating characteristic (ROC) curve for a continuous marker. This approach allows different markers to be compared on the same scale and facilitates the development of combination markers.

The objective of this article is to evaluate and compare baseline viral load and CD4 cell count as predictive biomarkers under the potential outcome framework. It is important to distinguish this objective from the related problems of identifying potential markers, combining several markers into a hybrid marker, choosing the cutoff point for a given marker and, more generally, developing an individualized treatment strategy. Variable selection techniques such as lasso-based methods have been used to select and combine genetic markers [e.g., Tian et al. (2012)]. Nonparametric multivariate methods [e.g., Foster, Taylor and Ruberg (2011), Qian and Murphy (2011), Su et al. (2008)] and semiparametric methods [e.g., Zhang et al. (2012)] have been used to develop a treatment rule (i.e., a set of criteria for selecting patients), which can be considered a binary hybrid marker obtained by combining multiple markers. A natural question that arises from the THRIVE study is whether or how to combine baseline viral load and CD4 cell count into a hybrid marker with improved predictive accuracy. While that question is beyond the scope of this article, we note that the potential outcome framework and the proposed methods are applicable to any given marker. Once a hybrid marker is developed, it can be evaluated and compared to the individual markers in the same framework using the same methods.

An analytical challenge for the potential outcome framework is that the desired treatment benefit, which involves potential outcomes under different treatments, is usually unobservable because each patient can receive only one treatment in a typical study. A possible exception to this limitation is a cross-over study, which has its own issues [e.g., Poulson, Gadbury and Allison (2012)] and will not be discussed in this article. In a typical clinical study such as the THRIVE study, the standard methodology in prediction and classification is not directly applicable to a predictive biomarker. To address this issue and the resulting identification problem, Huang, Gilbert and Janes (2012) make a monotonicity assumption, namely,

that one treatment dominates the other in all individual patients, and suggest a sensitivity analysis for possible departures from the monotonicity assumption. While the monotonicity assumption may be plausible in some situations such as vaccine trials, it may be less appealing as a starting point in other situations. In the THRIVE study, for example, the presence of a qualitative interaction (in the sense that neither treatment has a higher response rate for all realistic marker values) implies that neither treatment is dominant in all patients. Additionally, the approach of Huang, Gilbert and Janes (2012) is developed for a binary outcome and not readily extensible to other types of outcomes.

In this article, we propose alternative methods that do not require monotonicity or assume a binary outcome. Our first step is to account for the dependence between potential outcomes (for different treatments applied to the same patient) by adjusting for relevant covariates, such as demographic variables and baseline characteristics. If the set of measured covariates is sufficient for explaining the dependence between potential outcomes, the aforementioned performance measures can then be identified by assuming conditional independence of potential outcomes given covariates. Possible violations of this assumption can be addressed by introducing a random effect to account for residual dependence. In the next section we formulate the problem in terms of potential outcomes and provide a general rationale for the proposed approach. The proposed methods are then described in Section 3 and applied to the THRIVE study in Section 4. The article ends with a discussion in Section 5. Some technical details are provided in a supplemental article [Zhang et al. (2014)].

2. Notation and rationale. Suppose a randomized clinical trial is conducted to compare an experimental treatment with a control treatment, which may be a placebo or a standard treatment, with respect to a clinical outcome of interest, which may be discrete or continuous. Because there is only one outcome of primary interest in the THRIVE study, we will work with a scalar outcome unless otherwise noted. However, most of our methodology is readily applicable to a vector-valued outcome, with the exception of the sensitivity analysis in Section 3.2 (described fully in Web Appendix B). For a generic patient in the target population, let $Y(t)$ denote the potential outcome that will result if the patient receives treatment t (0 for control; 1 for experimental). Note that the $Y(t)$, $t = 0, 1$, cannot both be observed at a given time. Let T denote the treatment assigned randomly to a study subject, thus T is a Bernoulli variable independent of all baseline variables. Without considering noncompliance, which is negligible in the THRIVE study, we assume that T is also the actual treatment given to the subject, and write $Y = Y(T)$ for the actual outcome. Should noncompliance become a major issue, we could take an intent-to-treat perspective and compare treatment assignments or use analytical techniques to recover the actual treatment effect [e.g., van der Laan and Robins (2003)]. We assume that large values of Y are desirable. Where necessary, the subscript $i = 1, \dots, n$ will be attached to random variables to denote individual patients in the trial.

Our interest is in evaluating a predictive biomarker Z , a baseline variable which may be binary or continuous, with higher marker values supporting the use of the new treatment. The biomarker Z is intended to identify the subpopulation of patients who would benefit from the new treatment relative to the control. It can be a continuous variable as in our motivating example or a binary one such as a treatment rule developed using nonparametric multivariate methods. Let the desired treatment benefit be indicated by $B = I\{(Y(0), Y(1)) \in \mathcal{B}\}$, where $I\{\cdot\}$ is the indicator function and \mathcal{B} is the set of desirable outcomes. Note that B is by definition a comparison of the two potential outcomes. For a binary outcome, B might be an indicator for $Y(0) < Y(1)$ or $Y(0) \leq Y(1)$, depending on which treatment is preferable with identical (efficacy) outcomes. In our example, we set $B = I\{Y(0) \leq Y(1)\}$ because rilpivirine is thought to have a better safety profile and will likely be preferred over efavirenz when their efficacy outcomes are identical. For a continuous outcome, we might take $\mathcal{B} = \{(y_0, y_1) : y_1 - y_0 > \delta\}$, where δ reflects considerations of cost, clinical significance and possibly the safety profiles of the two treatments (if not incorporated into a vector-valued outcome). For an ordered categorical outcome, the definition of B may be more complicated. We shall take the definition of B as given and focus on the evaluation of Z for predicting B . The target B is an intrinsic characteristic of an individual patient, which suggests that Z can be evaluated using well-known quantities in prediction and classification [e.g., Pepe (2003), Zhou, Obuchowski and McClish (2002), Zou et al. (2011)]. For a binary marker, it makes sense to consider the true and false positive rates, defined as $\text{TPR} = P(Z = 1|B = 1)$ and $\text{FPR} = P(Z = 1|B = 0)$, respectively. For a continuous marker, it is customary to consider the ROC curve defined as

$$\text{ROC}(s) = 1 - F_{Z|B=1} \{F_{Z|B=0}^{-1}(1 - s)\}.$$

Here and in the sequel, we use F to denote a generic (conditional) distribution function, with the subscript indicating the random variable(s) concerned. The ROC curve is simply a plot of TPR versus FPR for classifiers of the form $I(Z > z)$, with the threshold z ranging over all possible values.

Because B is never observed, the existing methodology for evaluating predictors, which generally assumes that B can be observed, cannot be used directly to evaluate a predictive biomarker. Nonetheless, we note that TPR, FPR and ROC are all determined by $F_{Z|B}$, which can be recovered using Bayes' rule from the marginal distribution of Z and the conditional probability $\pi_Z(z) = P(B = 1|Z = z)$. For a binary marker, $1 - \pi_Z(0)$ and $\pi_Z(1)$ are negative and positive predictive values, respectively,

$$(1) \quad \begin{aligned} \text{TPR} &= \frac{\tau \pi_Z(1)}{\tau \pi_Z(1) + (1 - \tau) \pi_Z(0)}, \\ \text{FPR} &= \frac{\tau \{1 - \pi_Z(1)\}}{\tau \{1 - \pi_Z(1)\} + (1 - \tau) \{1 - \pi_Z(0)\}}, \end{aligned}$$

where $\tau = P(Z = 1)$. For a continuous marker, we have

$$(2) \quad \begin{aligned} F_{Z|B=1}(z_0) &= \int_{-\infty}^{z_0} \pi_Z(z) dF_Z(z) / \int_{-\infty}^{\infty} \pi_Z(z) dF_Z(z), \\ F_{Z|B=0}(z_0) &= \int_{-\infty}^{z_0} \{1 - \pi_Z(z)\} dF_Z(z) / \int_{-\infty}^{\infty} \{1 - \pi_Z(z)\} dF_Z(z). \end{aligned}$$

Since Z is fully observed, the identifiability of $F_{Z|B}$ would follow from that of $\pi_Z(z)$. Once an estimate of $\pi_Z(z)$ is available, it can be substituted into the above displays together with an empirical estimate of τ or F_Z , depending on the nature of Z .

Despite its simple appearance, $\pi_Z(z)$ is not straightforward to estimate. In fact, for any conceivable form of B , the quantity $\pi_Z(z) = P\{(Y(0), Y(1)) \in B | Z = z\}$ is not empirically identifiable because it involves the joint distribution of $Y(0)$ and $Y(1)$ given $Z = z$. Owing to randomization, it is straightforward to identify

$$F_{t|Z}(y|z) = P\{Y(t) \leq y | Z = z\} = P(Y \leq y | T = t, Z = z)$$

for each $t \in \{0, 1\}$, and to estimate it from a regression analysis for Y given T and Z . However, the dependence structure of $Y(0)$ and $Y(1)$ given $Z = z$ is not identifiable from the data [e.g., [Gadbury and Iyer \(2000\)](#)], which is also known as the fundamental problem of causal inference [[Holland \(1986\)](#)]. Because $\pi_Z(z)$ is not determined by the ‘‘marginals’’ $F_{t|Z}$ ($t = 0, 1$), its identification and estimation require additional information or assumptions about the dependence between $Y(0)$ and $Y(1)$ given $Z = z$. For a binary outcome, this can be achieved by assuming monotonicity [i.e., $Y(0) \leq Y(1)$ with probability 1] as in [Huang, Gilbert and Janes \(2012\)](#). The monotonicity assumption corresponds to maximal positive dependence of $Y(0)$ and $Y(1)$.

For a general outcome and without assuming monotonicity, we develop alternative methods by adapting the techniques of [Dodd and Pepe \(2003\)](#) and [Zhang et al. \(2013\)](#). To account for the dependence of $Y(0)$ and $Y(1)$, we start by conditioning on relevant covariates that are associated with both outcomes. Let \mathbf{X} denote a vector of such covariates measured at baseline, which may include prognostic factors and effect modifiers. In the THRIVE study, \mathbf{X} may include gender, race, and baseline age and body mass index. We include Z as a component of \mathbf{X} and write $\mathbf{X} = (Z, \mathbf{W})$, where \mathbf{W} consists of the additional baseline covariates. Writing $\pi_X(\mathbf{x}) = P\{(Y(0), Y(1)) \in B | \mathbf{X} = \mathbf{x}\}$, a conditioning argument yields

$$(3) \quad \pi_Z(z) = E\{\pi_X(\mathbf{X}) | Z = z\} = \int \pi_X(z, \mathbf{w}) F_{W|Z}(d\mathbf{w}|z).$$

Because $F_{W|Z}$ is empirically identifiable and estimable, the challenge now is to identify and estimate $\pi_X(\mathbf{x})$.

If \mathbf{X} is sufficient for explaining the dependence between $Y(0)$ and $Y(1)$, then we can expect that

$$(4) \quad Y(0) \perp Y(1) | \mathbf{X},$$

that is, that $Y(0)$ and $Y(1)$ are conditionally independent given \mathbf{X} . This assumption cannot be verified with the observed data and must be based on external information. Under assumption (4), the joint distribution of $Y(0)$ and $Y(1)$ given \mathbf{X} is determined by the “marginals”

$$F_{t|X}(y|\mathbf{x}) = \mathbb{P}\{Y(t) \leq y | \mathbf{X} = \mathbf{x}\} = \mathbb{P}(Y \leq y | T = t, \mathbf{X} = \mathbf{x}) \quad (t = 0, 1),$$

which are straightforward to identify and estimate. In reality, assumption (4) can be violated because \mathbf{X} may not explain all the dependence between $Y(0)$ and $Y(1)$. Such violations can be examined in a sensitivity analysis based on a latent variable that accounts for any residual dependence between $Y(0)$ and $Y(1)$. Under this approach, assumption (4) is relaxed as follows:

$$(5) \quad Y(0) \perp Y(1) | (\mathbf{X}, U),$$

where U is a subject-specific latent variable that is independent of \mathbf{X} . In other words, U represents what is missing from \mathbf{X} that makes assumption (4) break down. Assumption (5) alone is not sufficient to identify $\pi_X(\mathbf{x})$ because U is unobserved. However, by specifying certain quantities related to U , we can perform a sensitivity analysis based on assumption (5), as we demonstrate in the next section.

3. Methodology. We now describe methods for estimating the aforementioned performance measures (TPR, FPR and ROC). As indicated earlier, we will start by estimating $\pi_X(\mathbf{x})$ under assumptions (4) and (5). This can be done using a direct approach and an indirect approach, to be described in Sections 3.1 and 3.2, respectively. The direct and indirect approaches are based on models for $\pi_X(\mathbf{x})$ and $F_{Y|T,X}(y|t, \mathbf{x})$, respectively, which we refer to as benefit and outcome models. A benefit model is directly informative about $\pi_X(\mathbf{x})$ and thus more interpretable in the present context, while an outcome model is more familiar to practitioners and easier to estimate and validate using standard techniques. Further comments comparing the two approaches are given at the end of Section 3.2 after the approaches are described and in Section 5. In Section 3.3 we show how to convert an estimate of $\pi_X(\mathbf{x})$ into one of $\pi_Z(z)$. Estimates of the performance measures of interest are given in Section 3.4.

3.1. Direct estimation of $\pi_X(\mathbf{x})$ based on a Benefit model. A benefit model is a parametric model for $\mathbb{P}(B = 1 | \mathbf{X}) = \mathbb{P}\{(Y(0), Y(1)) \in \mathcal{B} | \mathbf{X}\}$, such as the following generalized linear model (GLM):

$$(6) \quad \pi_X(\mathbf{X}; \boldsymbol{\alpha}) = \psi(\alpha_1 + \boldsymbol{\alpha}'_X \mathbf{X}),$$

where $\boldsymbol{\alpha} = (\alpha_1, \boldsymbol{\alpha}'_X)'$ is the regression parameter and ψ is an inverse link function. Since B is binary, the probit and logit links are natural choices.

Suppose the conditional independence assumption (4) holds. To gain some intuition, consider a discrete \mathbf{X} taking values in $\{\mathbf{x}_1, \dots, \mathbf{x}_K\}$. Within each stratum

defined by $\mathbf{X} = \mathbf{x}_k$, assumption (4) implies that $Y(0)$ and $Y(1)$ are independent of each other, as if they arise from different subjects, which we assume are independent. In other words, given that $\mathbf{X}_i = \mathbf{X}_j$, the natural pair $(Y_i(0), Y_i(1))$ is identically distributed as the artificial pair $(Y_i(0), Y_j(1))$. If $T_i = 0$ and $T_j = 1$, then $(Y_i(0), Y_j(1))$ is observable as (Y_i, Y_j) , so that $\pi_X(\mathbf{x}_k) = E(B|\mathbf{X} = \mathbf{x}_k)$ can be estimated by

$$(7) \quad \frac{1}{n_{0k}n_{1k}} \sum_{i \in \mathcal{S}_{0k}} \sum_{j \in \mathcal{S}_{1k}} B_{ij},$$

where $B_{ij} = I\{(Y_i, Y_j) \in \mathcal{B}\}$, $\mathcal{S}_{tk} = \{i : T_i = t, \mathbf{X}_i = \mathbf{x}_k\}$, and n_{tk} denotes the size of \mathcal{S}_{tk} ($t = 0, 1; k = 1, \dots, K$).

The question is how to generalize this idea to a nondiscrete \mathbf{X} . To follow the logic of (7), one would need to find subjects from different treatment groups with the same value of \mathbf{X} , which becomes difficult when \mathbf{X} has continuous components (4 in our example). To overcome that problem, we borrow ideas from [Dodd and Pepe \(2003\)](#), who consider semiparametric regression for the area under the ROC curve, and work with an expanded model given by

$$(8) \quad P(B_{ij} = 1|\mathbf{X}_i, \mathbf{X}_j; \boldsymbol{\beta}) = \psi\{\beta_1 + \boldsymbol{\beta}'_X(\mathbf{X}_i + \mathbf{X}_j)/2 + \boldsymbol{\beta}'_{dX}(\mathbf{X}_j - \mathbf{X}_i)\},$$

where $\boldsymbol{\beta} = (\beta_1, \boldsymbol{\beta}'_X, \boldsymbol{\beta}'_{dX})'$, $i \in \mathcal{S}_{0k}$ and $j \in \mathcal{S}_{1k}$. The new features of model (8) relative to model (6) are introduced for the sole purpose of estimating $\boldsymbol{\alpha}$. Our research question does not pertain to the left-hand side of (8) or the regression coefficient $\boldsymbol{\beta}_{dX}$. However, assumption (4) implies that $P(B_{ij} = 1|\mathbf{X}_i = \mathbf{X}_j = \mathbf{x}) = P(B = 1|\mathbf{X} = \mathbf{x})$. Thus, when $\mathbf{X}_i = \mathbf{X}_j$, model (8) reduces to model (6) with

$$(9) \quad \boldsymbol{\alpha} = (\beta_1, \boldsymbol{\beta}'_X)'$$

In that sense, model (8) is a helper model that allows us to estimate model (6) with the observed data. Let $\mathcal{S}_t = \{i : T_i = t\}$ and let n_t denote the size of \mathcal{S}_t ($t = 0, 1$). Then the regression parameter $\boldsymbol{\beta}$ in model (8) can be estimated by solving the equation

$$(10) \quad 0 = \sum_{i \in \mathcal{S}_0} \sum_{j \in \mathcal{S}_1} \frac{\partial \pi_{ij}}{\partial \boldsymbol{\beta}} \frac{B_{ij} - \pi_{ij}}{\pi_{ij}(1 - \pi_{ij})},$$

where

$$\pi_{ij} = \psi\{\beta_1 + \boldsymbol{\beta}'_X(\mathbf{X}_i + \mathbf{X}_j)/2 + \boldsymbol{\beta}'_{dX}(\mathbf{X}_j - \mathbf{X}_i)\}.$$

As suggested by [Dodd and Pepe \(2003\)](#), equation (10) need not include all possible pairs (i, j) ; it could be based on a subset of pairs such that $\|\mathbf{X}_i - \mathbf{X}_j\| < \varepsilon$, where $\|\cdot\|$ denotes Euclidean norm, for some $\varepsilon > 0$. The choice of ε represents a bias-variance trade-off, where a larger ε leads to better efficiency and stability and also more sensitivity to the last component of model (8).

The approach just described relies heavily on the conditional independence assumption (4), which relates model (8) to model (6) through equation (9). Equation (9) does not hold when assumption (4) is violated. However, under alternative assumptions, we have

$$(11) \quad \boldsymbol{\alpha} = \gamma(\boldsymbol{\beta}_1, \boldsymbol{\beta}'_X)'$$

for a scalar γ . The key assumptions for (11) include assumption (5) and a GLM-like structure analogous to model (6):

$$(12) \quad P(B = 1 | \mathbf{X}, U; \boldsymbol{\alpha}^*) = \psi(\boldsymbol{\alpha}^*_1 + \boldsymbol{\beta}^*'_X \mathbf{X} + \boldsymbol{\alpha}^*_U U),$$

where $\boldsymbol{\alpha}^* = (\boldsymbol{\alpha}^*_1, \boldsymbol{\alpha}^*'_X, \boldsymbol{\alpha}^*_U)'$. In Section A of the supplemental article [Zhang et al. (2014)], we state additional assumptions that lead to (11) and give an expression for γ . Since $\boldsymbol{\beta}$ is identifiable and estimable using the techniques described earlier, $\boldsymbol{\alpha}$ can be estimated as soon as γ is known or estimated. Unfortunately, γ is unidentifiable from the observed data. For the probit and logit links, we show in Section A of the supplemental article [Zhang et al. (2014)] that γ can take any value greater than $2^{-1/2} \approx 0.71$. Thus, when assumption (4) is in doubt, we can perform a sensitivity analysis based on specified values of $\gamma \in (2^{-1/2}, \infty)$, with $\gamma = 1$ corresponding to conditional independence.

3.2. *Indirect estimation of $\pi_X(\mathbf{x})$ based on an outcome model.* An outcome model is a parametric model, say, $F_{Y|T,X}(y|t, \mathbf{x}; \boldsymbol{\theta})$, for the conditional distribution of Y given T and \mathbf{X} , specified up to a finite-dimensional parameter $\boldsymbol{\theta}$. Let $f_{Y|T,X}(y|t, \mathbf{x}; \boldsymbol{\theta})$ denote the associated probability density or mass function. A typical outcome model would be a GLM with the following mean structure:

$$(13) \quad E(Y|T, \mathbf{X}; \boldsymbol{\theta}) = \psi\{\theta_1 + \theta_T T + \boldsymbol{\theta}'_X \mathbf{X} + \boldsymbol{\theta}'_{TX}(T\mathbf{X})\},$$

where $\boldsymbol{\theta} = (\theta_1, \theta_T, \boldsymbol{\theta}'_X, \boldsymbol{\theta}'_{TX})'$ and ψ is an inverse link function. The parameter $\boldsymbol{\theta}$ can be estimated by maximizing the likelihood $\prod_{i=1}^n f_{Y|T,X}(Y_i|T_i, \mathbf{X}_i; \boldsymbol{\theta})$, and the resulting maximum likelihood estimate (MLE) will be denoted by $\widehat{\boldsymbol{\theta}}$. Because of randomization, the outcome model $F_{Y|T,X}(y|t, \mathbf{x}; \boldsymbol{\theta})$ implies that

$$F_{t|X}(y|\mathbf{x}) = P\{Y(t) \leq y | \mathbf{X} = \mathbf{x}\} = F_{Y|T,X}(y|t, \mathbf{x}; \boldsymbol{\theta}) \quad (t = 0, 1),$$

which can be estimated by substituting $\widehat{\boldsymbol{\theta}}$ for $\boldsymbol{\theta}$.

Under the conditional independence assumption (4), the joint distribution

$$F_{\cdot|X}(y_0, y_1|\mathbf{x}) = P\{Y(0) \leq y_0, Y(1) \leq y_1 | \mathbf{X} = \mathbf{x}\}$$

is identified as

$$F_{0|X}(y_0|\mathbf{x})F_{1|X}(y_1|\mathbf{x}) = F_{Y|T,X}(y_0|0, \mathbf{x}; \boldsymbol{\theta})F_{Y|T,X}(y_1|1, \mathbf{x}; \boldsymbol{\theta})$$

and estimated by replacing $\boldsymbol{\theta}$ with $\widehat{\boldsymbol{\theta}}$. Write

$$\widehat{F}_{\cdot|X}^{CI}(y_0, y_1|\mathbf{x}) = F_{Y|T,X}(y_0|0, \mathbf{x}; \widehat{\boldsymbol{\theta}})F_{Y|T,X}(y_1|1, \mathbf{x}; \widehat{\boldsymbol{\theta}}),$$

where the superscript CI stands for conditional independence. The corresponding estimate of $\pi_X(\mathbf{x})$ is then given by

$$\int I\{(y_0, y_1) \in \mathcal{B}\} \widehat{F}_{\cdot|X}^{\text{CI}}(dy_0, dy_1|\mathbf{x}) =: \widehat{F}_{\cdot|X}^{\text{CI}}(\mathcal{B}|\mathbf{x}).$$

When assumption (4) is in doubt, we can perform a sensitivity analysis based on assumption (5), which implies that

$$\begin{aligned} F_{\cdot|X}(y_0, y_1|\mathbf{x}) &= \int F_{\cdot|X,U}(y_0, y_1|\mathbf{x}, u) F_U(du) \\ (14) \qquad \qquad \qquad &= \int F_{0|X,U}(y_0|\mathbf{x}, u) F_{1|X,U}(y_1|\mathbf{x}, u) F_U(du), \end{aligned}$$

where we generalize the previous notation in an obvious way (with U as an additional conditioning variable). This suggests that we specify a model, say, $F_{Y|T,X,U}(y|t, \mathbf{x}, u; \boldsymbol{\theta}^*)$, for the conditional distribution of Y given (T, \mathbf{X}, U) , with a finite-dimensional parameter $\boldsymbol{\theta}^*$. Analogous to the GLM (13), we work with a generalized linear mixed model (GLMM) with

$$(15) \quad E(Y|T, \mathbf{X}, U; \boldsymbol{\theta}^*) = \psi\{\theta_1^* + \theta_T^* T + \boldsymbol{\theta}_X^{*'} \mathbf{X} + \boldsymbol{\theta}_{TX}^{*'}(T\mathbf{X}) + \theta_U^* U\},$$

where $\boldsymbol{\theta}^* = (\theta_1^*, \theta_T^*, \boldsymbol{\theta}_X^{*'}, \boldsymbol{\theta}_{TX}^{*'}, \theta_U^*)'$. The GLMM is not completely identified without additional information, and we propose a sensitivity analysis based on specified values of θ_U^* (or, rather, $|\theta_U^*|$), which is described in Section B of the supplemental article [Zhang et al. (2014)].

It is worth mentioning that the random effect U has a different interpretation here than in Section 3.1. In model (15), U represents an unobserved prognostic factor which affects both potential outcomes in the same direction; a change in U may or may not have much effect on the treatment benefit B , depending on the precise definition of B and model (15). Although one could incorporate a random treatment effect into model (15), the resulting method will likely become very complicated. In model (12), U acts like an effect modifier in that a change in U leads directly to a change in the probability of a desirable treatment benefit. (Here we use the term effect modifier in a heuristic sense which may or may not agree with an interaction-based definition.) Thus, aside from modeling assumptions, the direct and indirect approaches also differ in how they deal with departures from assumption (4). The indirect approach is designed to address violations of assumption (4) due to an unmeasured prognostic factor, whereas the direct approach is more appropriate for violations of assumption (4) due to an unmeasured effect modifier.

3.3. *Estimation of $\pi_Z(z)$.* Equations (1) and (2) suggest that evaluation of a predictive biomarker Z can be based on $\pi_Z(z) = P(B = 1|Z = z)$ and the marginal distribution of Z . Because the latter is straightforward to estimate, this section is focused on estimation of $\pi_Z(z)$ with a given estimate of $\pi_X(\mathbf{x})$, say,

$\hat{\pi}_X(\mathbf{x})$, which may be obtained using any one of the proposed methods. For a binary marker, equation (3) suggests that $\pi_Z(z)$ can be estimated by $\sum_{i=1}^n I\{Z_i = z\} \hat{\pi}_X(z, \mathbf{W}_i) / \sum_{i=1}^n I\{Z_i = z\}$. We therefore assume that Z is continuous in the rest of this section.

We propose to estimate $\pi_Z(z)$ by substituting an estimate of $F_{W|Z}$ into equation (3). One could specify a parametric model for $F_{W|Z}$, however, this can be difficult because the dimension of \mathbf{W} can be rather large (5 in our example). We therefore exploit the fact that Z is only one-dimensional and employ nonparametric regression techniques in estimating $F_{W|Z}$. Let $\xi : \mathbb{R} \rightarrow [0, \infty)$ be a kernel function and $\lambda > 0$ a bandwidth parameter. Then we can estimate $F_{W|Z}(\mathbf{w}|z)$ by

$$\frac{\sum_{i=1}^n \xi\{(Z_i - z)/\lambda\} I(\mathbf{W}_i \leq \mathbf{w})}{\sum_{i=1}^n \xi\{(Z_i - z)/\lambda\}}.$$

This, together with $\hat{\pi}_X(\mathbf{x})$, can be substituted into equation (3) to estimate $\pi_Z(z)$ as

$$\hat{\pi}_Z(z) = \frac{\sum_{i=1}^n \xi\{(Z_i - z)/\lambda\} \hat{\pi}_X(z, \mathbf{W}_i)}{\sum_{i=1}^n \xi\{(Z_i - z)/\lambda\}}.$$

An important question here is how to choose the bandwidth λ , for which we propose a cross-validation approach. The estimate $\hat{\pi}_Z(z)$ can be regarded as a nonparametric regression of $\hat{\pi}_X(Z, \mathbf{W})$ on Z , and its predictive accuracy can be assessed by comparing the “response” $\tilde{B}_i = \hat{\pi}_X(\mathbf{X}_i)$ with the estimate $\hat{B}_i = \hat{\pi}_Z(Z_i)$. We propose to partition the sample into a training set and a validation set and choose a value of λ that minimizes the average of $(\tilde{B}_i - \hat{B}_i)^2$ in the validation set with $\hat{\pi}_Z(z)$ estimated from the training set using bandwidth λ .

3.4. *Estimation of TPR, FPR and ROC.* Given $\hat{\pi}_Z(z)$ from Section 3.3, the parameters of interest can be estimated using equations (1) and (2). For a binary marker, this leads to

$$\widehat{\text{TPR}} = \frac{\hat{\tau} \hat{\pi}_Z(1)}{\hat{\tau} \hat{\pi}_Z(1) + (1 - \hat{\tau}) \hat{\pi}_Z(0)},$$

$$\widehat{\text{FPR}} = \frac{\hat{\tau} \{1 - \hat{\pi}_Z(1)\}}{\hat{\tau} \{1 - \hat{\pi}_Z(1)\} + (1 - \hat{\tau}) \{1 - \hat{\pi}_Z(0)\}},$$

where $\hat{\tau} = n^{-1} \sum_{i=1}^n Z_i$. For a continuous marker, we have

$$\widehat{\text{ROC}}(s) = 1 - \hat{F}_{Z|B=1} \{ \hat{F}_{Z|B=0}^{-1}(1 - s) \},$$

where

$$\hat{F}_{Z|B=1}(z_0) = \sum_{i=1}^n I(Z_i \leq z_0) \hat{\pi}_Z(Z_i) / \sum_{i=1}^n \hat{\pi}_Z(Z_i),$$

$$\hat{F}_{Z|B=0}(z_0) = \sum_{i=1}^n I(Z_i \leq z_0) \{1 - \hat{\pi}_Z(Z_i)\} / \sum_{i=1}^n \{1 - \hat{\pi}_Z(Z_i)\}.$$

An asymptotic analysis of these estimates can be rather tedious, especially because $\hat{\pi}_Z(z)$ involves smoothing and cross-validation. For inference, we recommend the use of bootstrap confidence intervals. To account for all variability in the estimates, the entire estimation procedure, including bandwidth selection based on cross-validation, should be repeated for each bootstrap sample.

4. Application to the THRIVE study. We now apply the methods of Section 3 to the THRIVE study introduced in Section 1, a randomized, double-blind, double-dummy, noninferiority trial at 98 hospitals or medical centers in 21 countries [Cohen et al. (2011)]. The THRIVE study compared rilpivirine with efavirenz for treating HIV-1 in treatment-naïve adults, in the presence of common background nucleoside or nucleotide reverse transcriptase inhibitors (N[t]RTIs). The study enrolled 680 adult patients who were naïve to antiretroviral therapy, with a screening viral load of at least 5000 copies per ml and viral sensitivity to N(t)RTIs. The patients were randomized in a 1 : 1 ratio to receive oral rilpivirine 25 mg once daily or efavirenz 600 mg once daily, in addition to an investigator-selected regimen of background N(t)RTIs.

The outcome of interest to us is virologic response (viral load below 50 copies/ml) at week 48 of treatment, with patient discontinuation (about 5%) counted as failure. The observed virologic response rates are 86% and 82% in the rilpivirine and efavirenz groups, respectively, and the difference between the two groups (3.5%; 95% CI: -1.7 – 8.8%) meets a prespecified noninferiority criterion based on a 12% margin ($p < 0.0001$). Thus, rilpivirine appears comparable to, if not better than, efavirenz in terms of population-level efficacy. However, Figure 1 suggests that individual patients respond differently to the two therapies and that baseline viral load and CD4 cell count could be used as predictive biomarkers. As indicated earlier, we will for safety reasons define individual-level treatment benefit as $B = I\{Y(0) \leq Y(1)\}$, where $Y(0)$ and $Y(1)$ denote potential outcomes for efavirenz and rilpivirine, respectively.

Baseline viral load and CD4 cell count are both log-transformed before entering the benefit and outcome models as covariates. In addition to these biomarkers, the covariate vector \mathbf{X} also includes gender, race (black, white or other), age and body mass index at baseline. For the direct approach of Section 3.1, the benefit model is a logistic regression model given by (6), with the aforementioned covariates as linear terms (and no interactions), and the helper model is given by (8) with the same link. For the indirect approach of Section 3.2, the outcome model is a logistic regression model similar to (13) except that interactions of \mathbf{X} with T are limited to the two biomarkers. The selection of covariates and interactions in these models is based on subject matter knowledge and not on statistical tests. Estimation of model (8) is based on the 1% of pairs (of control and experimental patients) that are most similar in terms of $\|\mathbf{X}_i - \mathbf{X}_j\|$, as suggested by Dodd and Pepe (2003). (Other proportions, from 0.001 to 1, have been attempted without yielding a material difference.) Under the indirect approach, the simplified method

in Web Appendix B is used to estimate θ^* for a given θ_U^* . In any case, estimates of $\pi_X(\mathbf{x})$ are converted into estimates of $\pi_Z(z)$ using the kernel smoothing method of Section 3.3, with a Gaussian kernel and a cross-validated bandwidth. The cross-validation is based on a 1 : 1 random partition of the sample into a training set and a validation set, and the bandwidth is chosen among $\{2^l \text{sd}(Z) : l = -5, \dots, 5\}$ using the minimum mean squared error criterion. The formulas of Section 3.4 are used to obtain estimates of ROC curves, and the trapezoidal rule is then employed in calculating the associated AUCs. The above procedure is performed for both biomarkers on the original sample as well as 200 bootstrap samples. Pointwise 90% confidence intervals for ROC curves are obtained using a simple bootstrap percentile method, and bootstrap standard errors are used for inference on AUCs (and AUC differences between the two biomarkers).

Figure 2 gives a side-by-side comparison of ROC curves for the two biomarkers, estimated using the direct approach of Section 3.1 with $\gamma = 0.71, 1, 2, 4$. The value $\gamma = 1$ corresponds to assumption (4) of conditional independence, while the value 0.71 is a theoretical lower bound. The associated AUC estimates and standard errors are shown in the upper half of Table 1. Both Figure 2 and the relevant portion of Table 1 show that the two biomarkers are useful as predictive biomarkers, with ROC curves above the diagonal and AUCs greater than 0.5 after accounting for sampling variability. The performance of each biomarker does appear to depend heavily on the value of γ ; the AUC estimate increases dramatically with increasing γ . This pattern is confirmed by additional analyses based on other values of γ (results not shown). Given the remarks at the end of Section 3.2, these results suggest that evaluation of a predictive biomarker can be rather sensitive to an unmeasured effect modifier. On the other hand, for each value of γ , the AUC estimate for baseline viral load is higher than that for baseline CD4 cell count, although the difference is not statistically significant. These results suggest that comparison of predictive biomarkers may be less sensitive to the choice of γ , even though this particular data set does not provide strong evidence that baseline viral load is better than baseline CD4 cell count as a predictive biomarker. Whether we are evaluating a single marker or comparing two markers, there is an obvious relationship between increasing γ and greater variability in the estimates, as indicated by widening confidence intervals in Figure 2 and increasing standard errors in Table 1.

Figure 3 gives another comparison of the two biomarkers based on ROC curves estimated using the indirect approach of Section 3.2 with $\theta_U^* = 0, 1.8, 4, 8$. Here, the value $\theta_U^* = 0$ corresponds to conditional independence, and the value 1.8 is a lower confidence bound (to be discussed later). The associated AUC estimates and standard errors are shown in the lower half of Table 1. These results are consistent with those from the direct approach in suggesting that both biomarkers are useful as predictive biomarkers. In particular, the results for conditional independence ($\theta_U^* = 0$) are fairly consistent with their counterparts under the direct approach (with $\gamma = 1$). Like Figure 2, Figure 3 also shows an increasing trend for

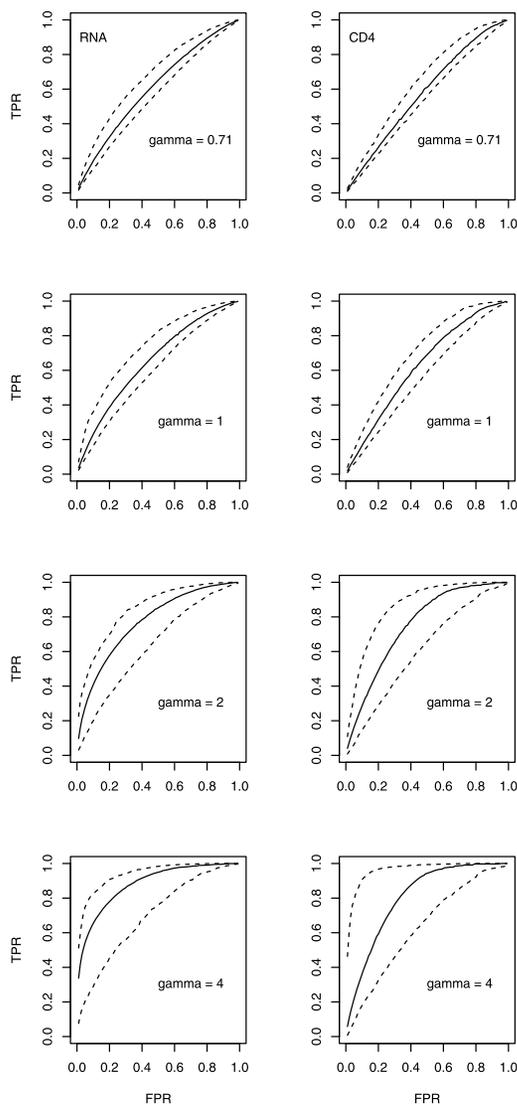


FIG. 2. ROC analysis of the THRIVE data using the direct approach: estimated ROC curves (solid lines) and 90% pointwise confidence intervals (dashed lines) for baseline viral load (RNA, left) and CD4 cell count (right), with different values of γ .

the predictive performance of each biomarker as a function of the sensitivity parameter θ_U^* . The trend is confirmed for additional values of θ_U^* , although the results are not shown. Thus, considering the remarks at the end of Section 3.2, it appears that evaluation of a predictive biomarker can also be sensitive to an unmeasured prognostic factor. The results from the indirect approach reinforce the previous observation that baseline viral load appears to perform better than baseline CD4

TABLE 1

AUC analysis of the THRIVE data: point estimates and bootstrap standard errors for the AUCs of baseline viral load (RNA) and CD4 cell count as well as their difference (RNA-CD4), obtained using the direct approach of Section 3.1 and the indirect approach of Section 3.2 with sensitivity parameters γ and θ_U^ , respectively*

Sensitivity parameter	Point estimate			Standard error		
	RNA	CD4	Diff.	RNA	CD4	Diff.
γ	<i>Direct approach</i>					
0.71	0.61	0.58	0.03	0.04	0.03	0.03
1	0.65	0.63	0.03	0.05	0.04	0.05
2	0.77	0.75	0.02	0.06	0.08	0.08
4	0.88	0.80	0.08	0.07	0.10	0.13
θ_U^*	<i>Indirect approach</i>					
0	0.64	0.59	0.05	0.04	0.04	0.05
1.8	0.65	0.60	0.06	0.05	0.05	0.05
4	0.68	0.64	0.04	0.07	0.07	0.08
8	0.72	0.70	0.02	0.09	0.09	0.11

cell count, although the differences here also fail to reach statistical significance. As is the case with the direct approach, increasing θ_U^* tends to produce greater variability in the estimates.

Although the uncertainty in comparing the two biomarkers could potentially be reduced by an increased sample size, some uncertainty will likely remain in evaluating each individual marker, given the apparent dependence on sensitivity parameters. Nonetheless, the increasing trend observed under both approaches suggests that a lower bound for predictive performance may be available under each approach. For the direct approach, the lower bound is given by $\gamma = 0.71$, as noted in Section 3.1. For the indirect approach, a lower bound for θ_U^* (and hence for the performance of each biomarker) may be available from longitudinal data (see Web Appendix B). For the THRIVE study, the lower bound for θ_U^* is estimated as 2.5 (95% CI: 1.8–3.7, based on 1000 bootstrap samples) from a GLMM analysis of repeated measurements at 24, 32, 40 and 48 weeks. Although earlier measurements (baseline through 20 weeks) are also available, we restrict our analysis to the later measurements in order to reduce misspecification bias; see Zhang et al. [(2013), Section 4] for a detailed explanation of this strategy. This GLMM analysis suggests that, under the additional assumptions given in Web Appendix B, the value $\theta_U^* = 1.8$ represents the worst case scenario for the indirect approach. The corresponding ROC and AUC estimates are better than those for $\gamma = 0.71$ under the direct approach and thus more informative as lower bounds.

Our ROC analyses under both (direct and indirect) approaches also illustrate that measures of predictive accuracy need not correlate with interactions. Although baseline CD4 cell count exhibits a more dramatic interaction in Figure 1, there is

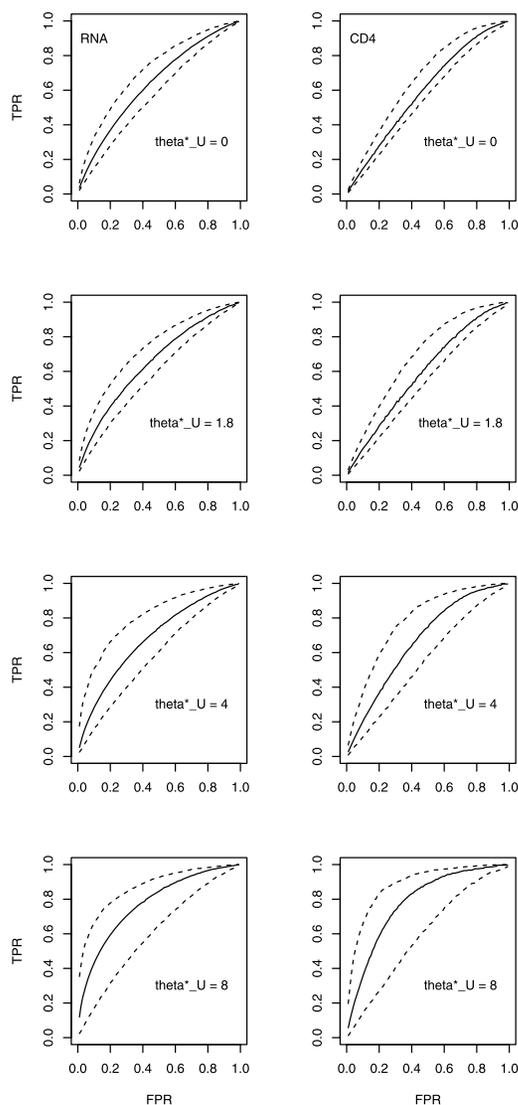


FIG. 3. ROC analysis of the THRIVE data using the indirect approach: estimated ROC curves (solid lines) and 90% pointwise confidence intervals (dashed lines) for baseline viral load (RNA, left) and CD4 cell count (right), with different values of θ_U^* .

no indication (in the same data set) that it outperforms baseline viral load as a predictive biomarker.

5. Discussion. In this article we have proposed new methods for evaluating predictive biomarkers in the potential outcome framework of Huang, Gilbert and Janes (2012). Instead of monotonicity, our starting point is conditional indepen-

dence of potential outcomes given observed covariates. Possible departures from the latter assumption can be addressed by incorporating a random effect that accounts for any residual dependence between potential outcomes. Because the random effect models are not completely identifiable, we propose a sensitivity analysis approach based on quantities related to the random effect. Our analysis of the THRIVE data reveals a great deal of sensitivity for the performance of each individual biomarker and much less sensitivity for the comparison of the two markers. Despite the uncertainty about individual biomarkers, the lower bounds on their predictive performance, available under both (direct and indirect) approaches, show that they are both useful as predictive biomarkers. For comparing the two markers, our analysis does not show much sensitivity and does not indicate a significant difference in the predictive performance of the two markers.

There does appear to be a lot of sensitivity (particularly in the upper portion) in quantifying the performance of an individual marker. Such sensitivity introduces additional uncertainty into the overall conclusion of the analysis, which is certainly undesirable. We see this as a reminder of the inherent limitation of the observed data for answering certain questions. As in many other statistical problems (e.g., missing data, censoring, confounding), the parameters of interest to us are not identifiable from the observed data without making untestable assumptions. When such assumptions are in doubt, the parameters are partially identified and the associated inference cannot be as sharp as that for point-identified parameters [Manski (2003)]. It may be disappointing to see that different assumptions (about the sensitivity parameter) can lead to a wide range for the parameter of interest. On the other hand, one could argue that the sensitivity analysis serves its purpose well by revealing the limitation of the available data and information.

In theory, point identification is achieved under the conditional independence assumption (6), which requires that the observed covariates be sufficient for explaining the dependence between the potential outcomes. Although we may never be certain that assumption (6) holds, it seems reasonable to believe that the assumption will get close to being true with a growing set of relevant covariates obtained from increasing knowledge of the disease. The more we know about the disease, the more information we have about relevant covariates, the more confidence we should be able to place in assumption (6). In practice, it may be difficult to determine when we have sufficient information to rely on assumption (6) and when we have to perform a sensitivity analysis. One possible solution would be a Bayesian approach with an informative prior on the sensitivity parameter which quantifies our uncertainty about assumption (6). As we become more confident about assumption (6), the prior will become more concentrated at or near that assumption.

Each of the direct and indirect approaches has unique advantages. The direct approach is able to accommodate complex definitions of treatment benefit involving several outcome variables of arbitrary types, as long as they are all observed.

The indirect approach is best suited for a single outcome of primary interest. Although the outcome model can include multiple outcomes in principle, their dependence structure can be difficult to specify and estimate. Even for a single outcome, the indirect approach requires a greater amount of modeling (in the sense that an outcome model implies a benefit model) and is therefore more prone to misspecification bias. On the other hand, the indirect approach is able to use all the information in the observed outcome data and therefore may have an efficiency advantage. Finally, for sensitivity analysis, the direct approach is more appropriate for an unmeasured effect modifier, and the indirect approach for an unmeasured prognostic factor. In practice, we recommend that both approaches be used in a complementary fashion, as in our analysis of the THRIVE data.

A reviewer has pointed out that some elements of personal judgment may be involved in choosing among different treatments. While this is not a major issue in the THRIVE data, where the definition of a treatment benefit is quite objective, it can certainly become a major issue in other therapeutic areas such as weight loss. For example, some patients may be willing to accept the extra risks of a surgical procedure (relative to a nonsurgical one) for an additional loss of 20 pounds, while others may not. To incorporate such personal judgment into the proposed approach, we could allow δ to vary across patients, and we would need to be able to measure δ for individual patients or at least be able to predict δ using measurable individual characteristics. In the latter case, the methodology will need to be modified to incorporate a prediction model for δ and the associated variability. It will be of interest to explore that possibility in the context of a suitable application.

Because of the complexity of the proposed methodology, a sample size formula is not yet available; however, for a given application, one could perform a simulation study to gauge the adequacy (in terms of power and precision) of a proposed sample size. Such an assessment should obviously consider the objective of the analysis (e.g., evaluating a single biomarker versus comparing two or more biomarkers). In addition, the nonparametric regression in Section 3.3 may imply a higher requirement on the sample size than do the other parts of the methodology, which are based on parametric regression techniques. A sample size that is inadequate for one-dimensional nonparametric regression may compromise the performance of the methods.

Acknowledgment. We thank the Editor (Dr. Karen Kafadar), the Associate Editor and three anonymous reviewers for their insightful and constructive comments, which have greatly improved the manuscript. The views expressed in this article do not represent the official position of the U.S. Food and Drug Administration.

SUPPLEMENTARY MATERIAL

Supplement: Technical details (DOI: [10.1214/14-AOAS773SUPP](https://doi.org/10.1214/14-AOAS773SUPP); .pdf). We provide technical details concerning the sensitivity analyses in Sections 3.1 and 3.2.

REFERENCES

- COHEN, C. J., ANDRADE-VILLANUEVA, J., CLOTET, B., FOURIE, J., JOHNSON, M. A., RUXRUNGTHAM, K., WU, H., ZORRILLA, C., CRAUWELS, H., RIMSKY, L. T., VANVEGEL, S., BOVEN, K. and THRIVE STUDY GROUP (2011). Rilpivirine versus efavirenz with two background nucleoside or nucleotide reverse transcriptase inhibitors in treatment-naive adults infected with HIV-1 (THRIVE): A phase 3, randomised, non-inferiority trial. *Lancet* **378** 229–237.
- DODD, L. E. and PEPE, M. S. (2003). Semiparametric regression for the area under the receiver operating characteristic curve. *J. Amer. Statist. Assoc.* **98** 409–417. [MR1995717](#)
- FOSTER, J. C., TAYLOR, J. M. G. and RUBERG, S. J. (2011). Subgroup identification from randomized clinical trial data. *Stat. Med.* **30** 2867–2880. [MR2844689](#)
- GADBURY, G. L. and IYER, H. K. (2000). Unit-treatment interaction and its practical consequences. *Biometrics* **56** 882–885.
- GAIL, M. and SIMON, R. (1985). Testing for qualitative interactions between treatment effects and patient subsets. *Biometrics* **41** 361–372.
- HOLLAND, P. W. (1986). Statistics and causal inference (with discussion). *J. Amer. Statist. Assoc.* **81** 945–970. [MR0867618](#)
- HUANG, Y., GILBERT, P. B. and JANES, H. (2012). Assessing treatment-selection markers using a potential outcomes framework. *Biometrics* **68** 687–696. [MR3055173](#)
- KARAPETIS, C. S., KHAMBATA-FORD, S., JONKER, D. J., O'CALLAGHAN, C. J., TU, D., TEBBUTT, N. C., SIMES, R. J., CHALCHAL, H., SHAPIRO, J. D., ROBITAILLE, S., PRICE, T. J., SHEPHERD, L., AU, H.-J., LANGER, C., MOORE, M. J. and ZALCBERG, J. R. (2008). K-ras mutations and benefit from cetuximab in advanced colorectal cancer. *N. Engl. J. Med.* **359** 1757–1765.
- MANSKI, C. F. (2003). *Partial Identification of Probability Distributions*. Springer, New York. [MR2151380](#)
- PAIK, S., SHAK, S., TANG, G., KIM, C., BAKER, J., CRONIN, M., BAEHNER, F. L., WALKER, M. G., WATSON, D., PARK, T., HILLER, W., FISHER, E. R., WICKERHAM, D. L., BRYANT, J. and WOLMARK, N. (2004). A multigene assay to predict recurrence of tamoxifen-treated, node-negative breast cancer. *N. Engl. J. Med.* **351** 2817–2826.
- PEPE, M. S. (2003). *The Statistical Evaluation of Medical Tests for Classification and Prediction*. Oxford Statistical Science Series **28**. Oxford Univ. Press, Oxford. [MR2260483](#)
- POCOCK, S. J., ASSMANN, S. E., ENOS, L. E. and KASTEN, L. E. (2002). Subgroup analysis, covariate adjustment and baseline comparisons in clinical trial reporting: Current practice and problems. *Stat. Med.* **21** 2917–2930.
- POULSON, R. S., GADBURY, G. L. and ALLISON, D. B. (2012). Treatment heterogeneity and individual qualitative interaction. *Amer. Statist.* **66** 16–24. [MR2934737](#)
- QIAN, M. and MURPHY, S. A. (2011). Performance guarantees for individualized treatment rules. *Ann. Statist.* **39** 1180–1210. [MR2816351](#)
- RUSSEK-COHEN, E. and SIMON, R. M. (1998). Evaluating treatments when a gender by treatment interaction may exist. *Stat. Med.* **16** 455–464.
- SIMON, R. (2008). Development and validation of biomarker classifiers for treatment selection. *J. Statist. Plann. Inference* **138** 308–320. [MR2412589](#)
- SIMON, R. (2010). Clinical trials for predictive medicine: New challenges and paradigms. *Clin. Trials* **7** 516–524.
- SU, X., ZHOU, T., YAN, X., FAN, J. and YANG, S. (2008). Interaction trees with censored survival data. *Int. J. Biostat.* **4** Art. 2, 28. [MR2383729](#)
- TIAN, L., ALIZADEH, A. A., GENTLES, A. J. and TIBSHIRANI, R. (2012). A simple method for detecting interactions between a treatment and a large number of covariates. Available at [arXiv:1212.2995](#).

- VAN DER LAAN, M. J. and ROBINS, J. M. (2003). *Unified Methods for Censored Longitudinal Data and Causality*. Springer, New York. [MR1958123](#)
- ZHANG, B., TSIATIS, A. A., LABER, E. B. and DAVIDIAN, M. (2012). A robust method for estimating optimal treatment regimes. *Biometrics* **68** 1010–1018. [MR3040007](#)
- ZHANG, Z., WANG, C., NIE, L. and SOON, G. (2013). Assessing the heterogeneity of treatment effects via potential outcomes of individual patients. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **62** 687–704. [MR3118329](#)
- ZHANG, Z., NIE, L., SOON, G. and LIU, A. (2014). Supplement to “The use of covariates and random effects in evaluating predictive biomarkers under a potential outcome framework.” DOI:[10.1214/14-AOAS773SUPP](#).
- ZHOU, X.-H., OBUCHOWSKI, N. A. and MCCLISH, D. K. (2002). *Statistical Methods in Diagnostic Medicine*. Wiley, New York. [MR1915698](#)
- ZOU, K. H., LIU, A., BANDOS, A. I., OHNO-MACHADO, L. and ROCKETTE, H. E. (2011). *Statistical Evaluation of Diagnostic Performance: Topics in ROC Analysis*. CRC Press, Boca Raton, FL. [MR3052764](#)

Z. ZHANG
DIVISION OF BIostatISTICS
OFFICE OF SURVEILLANCE AND BIOMETRICS
CENTER FOR DEVICES
AND RADIOLOGICAL HEALTH,
FOOD AND DRUG ADMINISTRATION
10903 NEW HAMPSHIRE AVE
SILVER SPRING, MARYLAND 20993
USA
E-MAIL: zhiwei.zhang@fda.hhs.gov

G. SOON
DIVISION OF BIOMETRICS IV
OFFICE OF BIostatISTICS
CENTER FOR DRUG EVALUATION AND RESEARCH,
FOOD AND DRUG ADMINISTRATION
10903 NEW HAMPSHIRE AVE
SILVER SPRING, MARYLAND 20993
USA
E-MAIL: guoxing.soon@fda.hhs.gov

L. NIE
DIVISION OF BIOMETRICS V
OFFICE OF BIostatISTICS
CENTER FOR DRUG EVALUATION AND RESEARCH,
FOOD AND DRUG ADMINISTRATION
10903 NEW HAMPSHIRE AVE
SILVER SPRING, MARYLAND 20993
USA
E-MAIL: lei.nie@fda.hhs.gov

A. LIU
BIostatISTICS AND BIOINFORMATICS BRANCH
DIVISION OF INTRAMURAL POPULATION
HEALTH RESEARCH
EUNICE KENNEDY SHRIVER NATIONAL INSTITUTE
OF CHILD HEALTH AND HUMAN DEVELOPMENT
NATIONAL INSTITUTES OF HEALTH
6100 EXECUTIVE BLVD.
BETHESDA, MARYLAND 20892
USA
E-MAIL: liua@mail.nih.gov