Published in final edited form as:

Biometrics. 2018 March ; 74(1): 18–26. doi:10.1111/biom.12743.

Incorporating Patient Preferences into Estimation of Optimal Individualized Treatment Rules

Emily L. Butler^{1,*}, Eric B. Laber^{2,**}, Sonia M. Davis^{1,***}, and Michael R. Kosorok^{1,3,****}

¹Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, U.S.A

²Department of Statistics, North Carolina State University, Raleigh, North Carolina, U.S.A

³Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, U.S.A

Summary

Precision medicine seeks to provide treatment only if, when, to whom, and at the dose it is needed. Thus, precision medicine is a vehicle by which healthcare can be made both more effective and efficient. Individualized treatment rules operationalize precision medicine as a map from current patient information to a recommended treatment. An optimal individualized treatment rule is defined as maximizing the mean of a pre-specified scalar outcome. However, in settings with multiple outcomes, choosing a scalar composite outcome by which to define optimality is difficult. Furthermore, when there is heterogeneity across patient preferences for these outcomes, it may not be possible to construct a single composite outcome that leads to high-quality treatment recommendations for all patients. We simultaneously estimate the optimal individualized treatment rule for all composite outcomes representable as a convex combination of the (suitably transformed) outcomes. For each patient, we use a preference elicitation questionnaire and item response theory to derive the posterior distribution over preferences for these composite outcomes and subsequently derive an estimator of an optimal individualized treatment rule tailored to patient preferences. We prove that as the number of subjects and items on the questionnaire diverge, our estimator is consistent for an oracle optimal individualized treatment rule wherein each patient's preference is known a priori. We illustrate the proposed method using data from a clinical trial on antipsychotic medications for schizophrenia.

Keywords

Competing Outcomes; Individualized Treatment Rules; Item Response Theory; Precision Medicine

*butlerel@live.unc.edu
**eblaber@ncsu.edu
***davissm@email.unc.edu
****kosorok@bios.unc.edu

6. Supplementary Materials

Web Appendices referenced in Sections 2.1, 2.2, 3 and 4 are available with this paper at the Biometrics website on the Wiley Online Library.

1. Introduction

It is widely recognized that best possible clinical care is tailored to individual patient characteristics (Sox et al., 2008; Hamburg and Collins, 2010; Collins and Varmus, 2015), including subjective factors like patient personal preference (Edwards and Elwyn, 2009). Individualized treatment rules (ITRs) formalize personalized clinical care as a function from up-to-date patient information to a recommended treatment. An optimal ITR maximizes the mean of some pre-specified clinical outcome if applied to make treatment decisions for all patients in a population of interest. This definition of optimality is mathematically convenient as it reduces estimation of an optimal ITR to a scalar optimization problem over a class of potential ITRs. However, this formulation does not directly allow for shared decision making wherein patient preferences are integrated into the decision process (Barry and Edgman-Levitan, 2012; Drake et al., 2010); on the other hand, direct preference elicitation in which the patient chooses parameters indexing a composite outcome is not feasible unless patients have undergone specialized training (Brennan, 1998; Braziunas, 2006; Lizotte et al., 2012). Thus, a common approach for preference elicitation is to administer a questionnaire populated with items that are accessible (meaningful) to a patient in a domain context yet are informative about preferences in the outcome space. These questions may ask the patient to rank or numerically score different health states and/or may ask about their attitudes or experiences with certain outcomes. We consider the simplest possible setting in which the questionnaire comprises a series of binary questions; we use item response theory (Embretson and Reise, 2013) to estimate a conditional distribution over preferences given a patient's answers to the items in the questionnaire. We use this conditional distribution to derive a preference-sensitive optimal ITR for each patient.

There is a vast literature on estimation for optimal ITRs assuming a single scalar outcome. Estimators are typically broadly categorized as regression-based (Murphy, 2005; Henderson et al., 2010; Zhao et al., 2011; Qian and Murphy, 2011; Goldberg and Kosorok, 2012; Moodie et al., 2012, 2014; Tian et al., 2014; Laber et al., 2014) or policy-search-based (Orellana et al., 2010; Zhang et al., 2012; Zhao et al., 2012; Zhang et al., 2012, 2013; Zhao et al., 2014, 2015; Laber and Zhao, 2015), though this division is somewhat superficial (Taylor et al., 2015). The estimators referenced above assume that a single outcome has been pre-specified and that this outcome faithfully represents the preferences of all patients in the population.

The literature that does not assume a fixed and known composite outcome is scarce and none addresses patient-specific preferences directly through elicitation. There are three primary approaches to estimation of optimal ITRs with multiple outcomes: (i) set-valued treatment regimes; (ii) inverse-preference elicitation; and (iii) constrained estimation. Set-valued treatment regimes (Laber et al., 2014; Lizotte and Laber, 2016) map current patient information to a subset of possible treatment options that are not estimated to be uniformly worse across all outcomes; the expectation is that a single treatment will be selected from the recommended set using patient preferences and clinical judgment (no guidance is provided on how this will be done). Thus, set-valued treatment regimes incorporate patient preferences through selection from the recommended subset but do not directly elicit preferences. Inverse-preference elicitation (Lizotte et al., 2012) attempts to communicate to

a patient the implicit composite outcomes that would be consistent with each possible treatment choice, e.g., choosing treatment 1 over treatment -1 corresponds to valuing side-effect burden at least three times as much as efficacy. Application of inverse preference elicitation requires the patient to interpret convex combinations of outcomes which may be difficult without specialized training. Another approach to estimation of optimal ITRs with multiple outcomes is to maximize the expectation of one outcome subject to constraints on a functional of the distribution of the remaining outcomes (Linn et al., 2016). This approach assumes that a single (unknown) composite outcome reflects patient preferences across the entire population.

The proposed methodology incorporates patient preferences into treatment selection in a way that is mathematically rigorous yet does not require the patient to undergo specialized training or understand complex quantitative concepts. This has at least two important clinical impacts: (i) it facilitates ‘patient-centered care’ in which patients play a key role in decision making and the evaluation of their own health outcomes; and (ii) it offers a principled means for matching patient preferences to an optimal treatment based on potentially complex outcome profiles. While many clinical and intervention scientists already seek to implement patient-centered care and to personalize treatment decisions based on patient preferences, it is recognized that the rapid introduction of new therapies and a high-degree of patient preference heterogeneity makes this a difficult task to perform using expert judgement alone (Smoller and Nierenberg, 1999; Basu and Meltzer, 2007; Frank and Zeckhauser, 2007; Hodgkin et al., 2012; Huskamp et al., 2013). The proposed methodology can be viewed as a tool for clinical decision support in the context of patient-centered care.

One therapeutic area where this method can be applied is the treatment of schizophrenia, which is a chronic, severe and debilitating mental disease. Because schizophrenia cannot be cured at this time, treatments are designed solely to alleviate symptoms. One of the main treatments for schizophrenia is the administration of antipsychotic medications. Some of these antipsychotics are associated with negative side effects, making it difficult for patients to adhere to treatment, while others have fewer side effects but are less efficacious. We obtained data from the Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE) schizophrenia trial to illustrate the application of the work presented here. In this trial, a series of atypical and typical antipsychotics were compared in a randomized, controlled, double blinded, multi-phase trial that aimed to treat schizophrenia patients. At baseline, patients answered a ten question binary response assessment which we used to measure patient attitudes across two outcomes: (1) efficacy, as measured using the positive and negative syndromes scale, and (2) aggregate side effect burden. We found our method performs as expected in balancing the patient’s attitudes between the two outcomes.

In Section 2, we formally define an optimal ITR under heterogeneous patient preferences and describe the proposed estimator and its asymptotic behavior. In Section 3, we evaluate the finite sample performance of the proposed estimator using a series of simulation experiments. In Section 4, we demonstrate how the proposed estimator can be used for clinical decision support using data from a clinical trial of schizophrenia. A discussion of open problems is provided in Section 5.

2. Optimal ITRs with Heterogeneous Patient Preferences

2.1 Setup and Notation

We assume that the observed data, $\{(W_i, X_i, A_i, Y_i, Z_i)\}_{i=1}^n$, comprises n independent and identically distributed tuples (W, X, A, Y, Z) , one per subject, where: $W \in \{0, 1\}^p$ denotes answers to items in a preference questionnaire; $X \in \mathbb{R}^m$ denotes pre-treatment patient covariates; $A \in \{-1, 1\}$ denotes the assigned treatment; $Y, Z \in \mathbb{R}$ denote outcomes of interest, coded so that higher values are better. In the context of our application to schizophrenia: W denotes subject's responses to a subset of questions on the Hogan Drug Attitude Inventory (Hogan et al., 1983); X denotes patient demographics, measures of symptom severity, and the presence/absence of comorbidities (see Section 3.4 for a complete description); A denotes perphenazine or atypical antipsychotic medication; Y denotes efficacy measured in terms of the Positive and Negative Syndromes Scale (PANSS); and Z denotes a measure of side-effect severity. In this context, an ITR, $\pi: \text{dom } W \times \text{dom } X \rightarrow \text{dom } A$, is a map from answers on a completed questionnaire and patient covariates to a recommended treatment; under π , a patient with $W = w$ and $X = x$ would be recommended treatment $\pi(w, x)$.

To define an optimal ITR, we assume that each individual in the population possesses a latent preference, denoted $H \in \mathbb{R}$, which indexes a utility function, $U(y, z; h)$, that induces a total ordering on $\text{dom } Y \times \text{dom } Z$ so that a patient with preference $H = h$ would prefer outcomes (y, z) to (y', z') if $U(y, z; h) > U(y', z'; h)$. Let $Y^*(a)$ and $Z^*(a)$ denote the potential outcomes under treatments $a \in \{-1, 1\}$ (Rubin, 1978) so that $U\{Y^*(a), Z^*(a); H\}$ is the potential utility function under treatment a . For any ITR, π , define the potential utility as $V_U(\pi) = \mathbb{E}[\sum_{a \in \{-1, 1\}} U\{Y^*(a), Z^*(a); H\} 1_{\pi(W, X)=a}]$, where the expectation is taken with respect to the joint distribution of $\{X, W, H, A, Y^*(0), Y^*(1), Z^*(0), Z^*(1)\}$. The optimal ITR, π_U^{opt} , satisfies $V_U(\pi_U^{\text{opt}}) \geq V_U(\pi)$ for all ITRs π .

We assume a utility of the form $U(y, z; h) = \Phi(h)y + \{1 - \Phi(h)\}z$ where $\Phi(\cdot)$ denotes the cumulative distribution function of a standard normal random variable. The assumption of a linear utility is common in multiobjective optimization (Fishburn, 1988; Marler and Arora, 2004; Richter and Wong, 2004; Lizotte et al., 2012). This popularity is in part because of its analytical tractability, but also because a linear utility is intuitive in the sense that it expresses the notion that a patient with preference h cares $\Phi(h)/\{1 - \Phi(h)\}$ more about Y than about Z ; nevertheless, the assumption of a constant gain in utility per unit increase in outcome may not be reasonable in some contexts. The following result shows that the optimal ITR for any rational utility (i.e., a utility that will always prefer a treatment that is better on both outcomes to one that is worse on both outcomes) can be expressed as the optimal ITR for some linear utility function. Define

$$\begin{aligned} R_Z(w, x) &= \{Z^*(1)|W=w, X=x\} - \{Z^*(-1)|W=w, X=x\}, \\ R_Y(w, x) &= \{Y^*(1)|W=w, X=x\} - \{Y^*(-1)|W=w, X=x\}, \\ R_U(w, x) &= [U\{Y^*(1), Z^*(1); H\}|W=w, X=x] - [U\{Y^*(-1), Z^*(-1); H\}|W=w, X=x], \end{aligned}$$

we assume that the foregoing quantities are well-defined for all $\mathbf{w}$ and $\mathbf{x}$. It can be shown that $\pi_U^{\text{opt}}(\mathbf{w}, \mathbf{x}) = \text{sign} \{R_U(\mathbf{w}, \mathbf{x})\}$ (Qian and Murphy, 2011).

Lemma 2.1—*Assume that $\max \{R_L(\mathbf{w}, \mathbf{x})R_Z(\mathbf{w}, \mathbf{x}), R_L(\mathbf{w}, \mathbf{x})R_Y(\mathbf{w}, \mathbf{x})\} > 0$ for all $\mathbf{x}, \mathbf{w}$. Then, there exists a real-valued random variable, H' , such that: (i) $H' \perp A, \{Y^*(a), Z^*(a)\} : a \in \{-1, 1\} | \mathbf{X}, \mathbf{W}$; and (ii) the ITR*

$$\pi_{\text{CVX}}^{\text{opt}}(\mathbf{x}, \mathbf{w}) = \arg \max_{a \in \{-1, 1\}} [\Phi(H')Y^*(a) + \{1 - \Phi(H')\}Z^*(a) | \mathbf{X} = \mathbf{x}, \mathbf{W} = \mathbf{w}]$$

satisfies $V_U(\pi_{\text{CVX}}^{\text{opt}}) = V_U(\pi_U^{\text{opt}})$.

Proof: Define $B_Y(\mathbf{w}, \mathbf{x}) = \text{sign}[R_L(\mathbf{w}, \mathbf{x})R_Y(\mathbf{w}, \mathbf{x})]$ and $B_Z(\mathbf{w}, \mathbf{x}) = \text{sign}[R_L(\mathbf{w}, \mathbf{x})R_Z(\mathbf{w}, \mathbf{x})]$, and let $\xi(\mathbf{w}, \mathbf{x}) = 0.5 + 0.5 \times 1 \{ |R_Y(\mathbf{w}, \mathbf{x})R_Z(\mathbf{w}, \mathbf{x})| > 0 \} \times |R_Z(\mathbf{w}, \mathbf{x})| / (|R_Y(\mathbf{w}, \mathbf{x})| + |R_Z(\mathbf{w}, \mathbf{x})|)$ and Set $H' =$

$$H'(W, X) = \Phi^{-1}[\xi(W, X)] \cdot [1 \{B_Y(W, X) = 1 = -B_Z(W, X), |R_Y(W, X)| \leq |R_Z(W, X)|\} - 1 \{B_Z(W, X) = 1 = -B_Y(W, X), |R_Y(W, X)| \geq |R_Z(W, X)|\}].$$

Thus, given $\mathbf{W} = \mathbf{w}$ and $\mathbf{X} = \mathbf{x}$, $H' = h'$ is completely determined so that conclusion (i) holds and $\pi_{\text{CVX}}^{\text{opt}}(\mathbf{w}, \mathbf{x}) = \text{sign}[\Phi(h')R_Y(\mathbf{w}, \mathbf{x}) + \{1 - \Phi(h')\}R_Z(\mathbf{w}, \mathbf{x})]$ which we show in the supplement equals $\pi_U^{\text{opt}}(\mathbf{w}, \mathbf{x}) = \text{sign}\{R_U(\mathbf{w}, \mathbf{x})\}$, and thus conclusion (ii) also holds.

The condition in the preceding lemma that $R_L(\mathbf{w}, \mathbf{x})$ matches the sign of $R_Y(\mathbf{w}, \mathbf{x})$ or $R_Z(\mathbf{w}, \mathbf{x})$ ensures that for any $(\mathbf{w}, \mathbf{x})$, the optimal ITR under the unknown utility is rational. To see this, note that violation of this assumption implies that there exists a pair $(\mathbf{x}, \mathbf{w})$ such that $R_L(\mathbf{w}, \mathbf{x})$ has the opposite sign of both $R_Y(\mathbf{w}, \mathbf{x})$ and $R_Z(\mathbf{w}, \mathbf{x})$ so that the optimal treatment under U , $\pi_U^{\text{opt}}(\mathbf{w}, \mathbf{x}) = \text{sign}\{R_U(\mathbf{w}, \mathbf{x})\}$, yields a worse mean outcome in both Y and Z . Violation of this assumption is unrealistic in many settings. Note that a full proof of Lemma 2.1 can be found in Web Appendix A.

2.2 Estimating an Optimal ITR

To construct an estimator of π^{opt} we first express it in terms of the underlying generative model. We assume: (C1) consistency, $(Y, Z) = \{Y^*(A), Z^*(A)\}$; (C2) positivity, there exists $\epsilon > 0$ so that $P(A = a | \mathbf{X}, \mathbf{W}) \geq \epsilon$ for each a , with probability one; (C3) ignorability, $\{Y^*(a), Z^*(a)\} : a \in \{-1, 1\} \perp A | \mathbf{X}, \mathbf{W}$; and (C4) $(A, Y, Z) \perp H | \mathbf{X}, \mathbf{W}$. The first three assumptions are standard (Robins et al., 2000; Zhang et al., 2012), whereas (C4) can be satisfied provided the conditions of Lemma (2.1) hold; nevertheless, this assumption can be weakened to $A \perp H | \mathbf{X}, \mathbf{W}$ at the expense of postulating a model for the conditional mean of $\Phi(H)Y$ and $\Phi(H)Z$ given $(\mathbf{X}, \mathbf{W}, A)$. Under (C1)–(C3), it can be shown (Schulte et al., 2014) that $\pi^{\text{opt}}(\mathbf{x}, \mathbf{w}) = \arg \max_{a \in \{-1, 1\}} \mathbb{E}[\Phi(H)Y + \{1 - \Phi(H)\}Z | \mathbf{X} = \mathbf{x}, \mathbf{W} = \mathbf{w}, A = a]$ which, under (C4), can be written as

$$\pi^{\text{opt}}(\mathbf{x}, \mathbf{w}) = \arg \max_{a \in \{-1, 1\}} [\{\Phi(H)|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}\} (Y|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}, A=a) + [1 - \{\Phi(H)|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}\} (Z|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}, A=a)]] .$$

The foregoing expression suggests the following approach for estimating π^{opt} : (i) construct estimators, say $\hat{Q}_{Y,n}(\mathbf{x}, \mathbf{w}, a)$ and $\hat{Q}_{Z,n}(\mathbf{x}, \mathbf{w}, a)$, of $Q_Y(\mathbf{x}, \mathbf{w}, a) = \mathbb{E}(Y|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}, A=a)$ and $Q_Z(\mathbf{x}, \mathbf{w}, a) = \mathbb{E}(Z|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}, A=a)$; (ii) postulate a latent preference model linking the unobservable preference H with covariates $\mathbf{X}$ and preference questionnaire items $\mathbf{W}$ and use this postulated model to construct an estimator, say $\hat{\mu}_{H,n}(\mathbf{x}, \mathbf{w})$ of $\mu(\mathbf{x}, \mathbf{w}) = \mathbb{E}\{\Phi(H)|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}\}$; and (iii) compute $\hat{\pi}_n(\mathbf{x}, \mathbf{w}) = \arg \max_{a \in \{-1, 1\}} [\hat{\mu}_{H,n}(\mathbf{x}, \mathbf{w}) \hat{Q}_{Y,n}(\mathbf{x}, \mathbf{w}, a) + \{1 - \hat{\mu}_{H,n}(\mathbf{x}, \mathbf{w})\} \hat{Q}_{Z,n}(\mathbf{x}, \mathbf{w}, a)]$. Below we describe an instantiation of this approach that appears to work well in practice and possesses a number of desirable theoretical properties.

To construct estimators of $Q_Y(\mathbf{x}, \mathbf{w}, a)$ and $Q_Z(\mathbf{x}, \mathbf{w}, a)$ we postulate linear working models of the form $Q_Y(\mathbf{x}, \mathbf{w}, a; \psi_Y) = \mathbf{x}_{Y,0}^\top \psi_{Y,0} + \mathbf{w}_{Y,1}^\top \psi_{Y,1} + a \mathbf{x}_{Y,1}^\top \psi_{Y,2} + a \mathbf{w}_{Y,1}^\top \psi_{Y,3}$ and $Q_Z(\mathbf{x}, \mathbf{w}, a; \psi_Z) = \mathbf{x}_{Z,0}^\top \psi_{Z,0} + \mathbf{w}_{Z,0}^\top \psi_{Z,1} + a \mathbf{x}_{Z,1}^\top \psi_{Z,2} + a \mathbf{w}_{Z,1}^\top \psi_{Z,3}$, where $\mathbf{x}_\ell$ and $\mathbf{w}_\ell$ for $\ell = Y, Z$ and $j = 0, 1$ are known feature vectors constructed from $\mathbf{x}$ and $\mathbf{w}$ and ψ_W, ψ_Y are unknown parameter vectors. Let $\mathbb{P}_n$ denote the empirical measure and define $\hat{\psi}_{Y,n} = \arg \min_{\psi_Y} \mathbb{P}_n \{Y - Q_Y(\mathbf{X}, \mathbf{W}, A; \psi_Y)\}^2$ and $\hat{\psi}_{Z,n} = \arg \min_{\psi_Z} \mathbb{P}_n \{Z - Q_Z(\mathbf{X}, \mathbf{W}, A; \psi_Z)\}^2$. Subsequently, we construct estimators $Q_Y(\mathbf{x}, \mathbf{w}, a; \hat{\psi}_Y)$ and $Q_Z(\mathbf{x}, \mathbf{w}, a; \hat{\psi}_Z)$ of $Q_Y(\mathbf{x}, \mathbf{w}, a)$ and $Q_Z(\mathbf{x}, \mathbf{w}, a)$.

To develop our latent preference model, we assume (C5) that $H \perp \mathbf{X} | \mathbf{W}$. This assumption simplifies our development and is justified in our application where $\mathbf{X}$ does not contain information thought to be informative about patient preferences beyond $\mathbf{W}$; however, this assumption is not necessary as one can include baseline covariates into the latent patient preference model (De Ayala, 2013). We assume (C6) that latent patient preferences are connected to items on the questionnaire through a Rasch model (Rasch, 1961, 1980) of the form $\text{logit}\{P(W_j=1|H=h)\} = \beta_{0,j} + \beta_{1,j}h, j = 1, \dots, p$ which is indexed by $\beta = (\beta_{0,1}, \beta_{1,1}, \dots, \beta_{0,p}, \beta_{1,p})$; without (C5), one could include a term for $\mathbf{X}$, e.g.,

$\text{logit}\{P(W_j=1|H=h, \mathbf{X}=\mathbf{x})\} = \beta_{0,j} + \beta_{1,j}h + \alpha_j^\top \mathbf{x}_{j,0}$, where α_j is a vector of unknown parameters and $\mathbf{x}_{j,0}$ is known feature constructed from $\mathbf{x}$. Let β^* denote the true parameter value; we use the EM algorithm to construct an estimator, $\hat{\beta}_n$, of β^* (Rizopoulos, 2006).

Given an estimator $\hat{\beta}_n$ and a postulated marginal distribution, say p_h , for the latent preferences, the conditional distribution of H given $\mathbf{W} = \mathbf{w}$ is proportional to $p(\mathbf{w}|h)p_h(h)$ which can be approximated using Metropolis Hastings. Because the conditional distribution of H is only used to construct an estimator of $\mu(\mathbf{x}, \mathbf{w}) = \mathbb{E}\{\Phi(H)|\mathbf{X}=\mathbf{x}, \mathbf{W}=\mathbf{w}\}$ and because we assume that $\mu(\mathbf{x}, \mathbf{w})$ does not depend on $\mathbf{x}$, it is computationally less burdensome to apply a method of moments type estimator. Let $\hat{h}_n(\mathbf{w})$ denote the solution to

$\sum_{j=1}^p \hat{\beta}_{n,1,j} \expit(\hat{\beta}_{n,j,0} + \hat{\beta}_{n,1,j}h) = \sum_{j=1}^p \hat{\beta}_{n,1,j} w_j$, where $\expit(u) = \exp(u)/\{1 + \exp(u)\}$. Subsequently, let $\hat{\mu}_{H,n}(\mathbf{x}, \mathbf{w}) = \Phi\{\hat{h}_n(\mathbf{w})\}$ denote our estimator of $\mu(\mathbf{w}, \mathbf{x})$. Results provided in Web Appendix A and B of the Supplemental Material show that this estimator provides qualitatively similar results as Metropolis Hastings while being significantly less

computationally intensive. Furthermore, the method of moments estimator has the advantage of not depending on an assumed marginal distribution for H .

2.2.1 Theoretical results—Let $h^*(\mathbf{w})$ denote the solution to

$$\sum_{j=1}^p \beta_{1,j}^* \text{expit}(\beta_{j,0}^* + \beta_{j,1}^* h) = \sum_{j=1}^p \beta_{j,1}^* w_j. \text{ We make the following assumptions.}$$

- (A1) The number of items satisfies $3 \leq p_n = \alpha(e^n)$.
- (A2) The estimator $\hat{h}_n(\mathbf{w})$ converges in probability to $h^*(\mathbf{w})$, pointwise for all $\mathbf{w}$.
- (A3) The estimators $\hat{Q}_{Y,n}(\mathbf{x}, \mathbf{w}, a)$ and $\hat{Q}_{Z,n}(\mathbf{x}, \mathbf{w}, a)$ converge in probability to $Q_{Y,n}(\mathbf{x}, \mathbf{w}, a)$ and $Q_{Z,n}(\mathbf{x}, \mathbf{w}, a)$, pointwise for each $\mathbf{x}$, $\mathbf{w}$, and a .

The preceding assumptions are rather mild: (A1) can be expected to hold in many applications as it is common to have $p \ll n$; (A2) holds under standard regularity conditions for methods of moments estimators; and (A3) holds under standard linear modeling assumptions.

The first result establishes consistency of the proposed estimator for the optimal ITR as the sample size diverges but the number of items remains fixed.

Theorem 2.2: Assume (A1)–(A3) and let the number of items, $p_n = p$, be fixed. Then $V_L(\pi^{\text{opt}}) - V_L(\hat{\pi}_n)$ converges to zero in probability as $n \rightarrow \infty$.

The preceding result is relevant when the number of items is small relative to the number of enrolled subjects, e.g., in the illustrative example presented in Section 3.4 there are $p = 10$ items and $n = 957$ subjects. However, the next result shows that if the number of items is allowed to diverge with the sample size then the estimated optimal ITR performs as well as an oracle that knows each patients individual preference; i.e., the oracle ITR is allowed to depend on both $(\mathbf{x}, h)$ and given by $\pi^{\text{oracle}}(\mathbf{x}, h) = \arg\max_a \mathbb{E}(U | \mathbf{X} = \mathbf{x}, H = h, A = a)$.

Theorem 2.3: Assume (A1)–(A3) and suppose $p_n \rightarrow \infty$ as $n \rightarrow \infty$. Then $V_L(\hat{\pi}_n) - V_L(\pi^{\text{oracle}})$ converges to zero in probability as $n \rightarrow \infty$.

The proof of the preceding result involves, as an intermediate step, proving consistency of the Rasch model for individual patient preferences when the number of items is nearly exponential in the number of patients. To our knowledge, this is new result in item response theory that may be of independent interest.

Building a high-quality model for the optimal ITR when the number of items is much larger than the number of subjects may require imposing additional structure on the data-generating model. In many settings, it may be reasonable to assume that the Q -functions, $Q_Y(\mathbf{w}, \mathbf{x}, a)$ and $Q_Z(\mathbf{w}, \mathbf{x}, a)$, do not depend on the items $\mathbf{w}$; e.g., the items capture patient preferences but do not contain any additional information about the mean outcomes given covariates $\mathbf{X}$ and treatment A . In this setting, the number of items does not affect estimation of the Q -functions. At the next level of generality, one might assume that the Q -functions depend on $\mathbf{W}$ only through the distribution of H given $\mathbf{W}$. In this case, one could construct an estimator of each Q -function using scalar-on-function regression (Ramsay, 2006). In the case

where Q -functions depend directly on the distribution of the items, standard methods for high-dimensional regression could be applied (e.g., Tibshirani, 1996; Fan and Li, 2001, see Song et al. (2015) for application to treatment regimes).

3. Simulation Study

In this section we examine the finite sample performance of the proposed estimation method in terms of the average value obtained. Additional simulations comparing the method of moments estimator with Metropolis Hastings are provided in Web Appendix B of the Supplemental Material.

3.1 Data Generating Model

We consider the following class of generative models. Latent preferences are drawn as *i.i.d.* so that $H \sim \text{Normal}(0, 1)$. Responses to items are drawn *i.i.d.* so that $W_j \sim \text{Bernoulli}\{\text{expit}(\beta_{0,j} + \beta_{1,j}H)\}$; the true beta parameters are chosen to ensure that when $\hat{H}_n(W)$ was regressed on W using a probit regression model, the Nagelkerke's R^2 was equal to 0.5. The treatment, covariates, and outcome data were drawn *i.i.d.* so that: $A \sim \text{Unif}\{-1, 1\}$, $X \sim N_5(\mu, I_5)$, $Y = X^\top \theta_{00} + AX^\top \theta_{01} + \varepsilon$ and $Z = X^\top \theta_{10} + AX^\top \theta_{11} + \delta$ where $\varepsilon, \delta \sim N(0, 1)$, $\theta_{00} = (2.5, .2, .25, -.7, -2.5, 2.4)$, $\theta_{01} = (1.7, -2.3, 4.5, 6, -7.3, -1.6)$, $\theta_{10} = t + q^* \theta_{00}$ and $\theta_{11} = t + q^* \theta_{01}$. We set $q = -2$ and $t = 3$ so that the outcomes favor different treatments about 85% of the time; other settings provide qualitatively similar results and are provided in the Supplemental Material.

3.2 Simulation Results

The simulations were each repeated 500 times for all combinations of $n = 25, 50, 100, 200$ and $p = 5, 10, 20, 50$. Let $Q_U(x, w, a) = \mathbb{E}(U|X = x, W = w, A = a)$ and $\hat{Q}_U(x, w, a) = \hat{\mu}_{H,n}(x, w) \hat{Q}_{Y,n}(x, w, a) + \{1 - \hat{\mu}_{H,n}(x, w)\} \hat{Q}_{Z,n}(x, w, a)$ to be the proposed estimator of $Q_U(x, w, a)$. In our first simulation experiment we compute the average squared difference between the estimated value function $\hat{V}_U(\hat{\pi}_n) = \mathbb{P}_n \max_a \hat{Q}_U(X, W, a)$ and the true value function of the optimal ITR, $V_U(\pi^{\text{opt}})$. The results are displayed in Figure 1. It can be seen that the quality of the estimated value improves as either n or p increases. The averaged squared distance between $\hat{V}_Y(\hat{\pi}_Y) = \mathbb{P}_n \max_a \hat{Q}_Y(X, W, a)$ and $V_Y(\pi_Y^{\text{opt}})$ is displayed in Figure 2; as expected, the quality of the approximation improves as n increases but is insensitive to changes in p . Figure 3 shows an analogous figure for the average difference between $\hat{V}_Z(\hat{\pi}_Z)$ and $V_Z(\pi_Z^{\text{opt}})$.

We also compared $\hat{\pi}_n$ with π^{opt} and π^{oracle} in terms of treatment recommendations. Figures 4 and 5 show the average percentage of disagreement of $\hat{\pi}_n$ with π^{opt} and π^{oracle} . It can be seen that the level of agreement is generally high and, as anticipated by the theoretical results in the preceding section, the agreement improves as n and p increase.

4. Case Study

The Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE) schizophrenia trial was designed to compare new antipsychotic drugs to conventional ones in a randomized,

controlled, double blinded, multi-phase trial (Lieberman et al., 2005; Stroup et al., 2003; Davis et al., 2003, 2011). The trial was intended for patients who were already being treated for schizophrenia but who might benefit from a medicinal change. The patients not only received antipsychotic treatments, but were also offered psychosocial treatment with their families. In the first phase of the trial, patients were randomized to one of 5 medications: 4 of which were atypical antipsychotics (olanzapine, quetiapine, risperidone and ziprasidone) and one conventional antipsychotic (perphenazine). The primary study analysis (Lieberman et al., 2005) compared all treatments individually, and identified substantial variability amongst the atypicals. However, for the purposes of this illustration, we dichotomize treatment into two groups: atypical antipsychotics and perphenazine. Patients who discontinued their treatment would move to phase 2 and would be randomized to another treatment according to the protocol. The structure of the first phase of this trial makes it an ideal application for the method presented in this paper. At baseline, patients answered a 10 question, binary-response assessment, which we use as a measure of the patient's preference across two outcomes: (i) efficacy measured using the Positive and Negative Syndromes Scale (PANSS) (Kay et al., 1987); and (ii) side effect burden measured as the sum of side effect and adverse event indicators. Detailed descriptions of these outcomes are provided in Web Appendix C of the Supplemental Material. Outcomes were measured at the end of the first phase except for adverse events which were monitored throughout the phase. The end of the first phase occurred when the patient underwent a treatment change (at the discretion of the patient and clinician) or 18 months, whichever came first.

The patient preference information was collected using a 10 question Drug Attitude Inventory (Hogan et al., 1983) assessment that each patient answered at the baseline. We opted to exclude one question ("I take medications of my own free choice") because it cannot be interpreted as a trade-off between two outcomes of interest. The remaining 9 questions are given in Table 1 and were coded so that (1) favors efficacy and (0) favors side effect burden.

We selected tailoring covariates based on clinical expertise and prior analyses (Shortreed and Moodie, 2012; Shortreed et al., 2014). The variables included in our analysis are gender, age, race (white, black, other), BMI, diastolic blood pressure, systolic blood pressure, baseline clinical severity of schizophrenia, age at first antipsychotic medication, and any antipsychotic medication they were taking at the baseline (olanzapine, quetiapine, risperidone, ziprasidone, perphenazine, haldol, or any long-term injectable antipsychotic).

The coefficients indexing the estimated optimal Q -functions are displayed in Table 2. The main effect of treatment has an opposite sign in the two Q -functions as well as several interactions involving treatment; thus, there is a trade-off between the two outcomes that must be made in choosing a treatment and furthermore, this trade-off varies across patient characteristics. The estimated optimal treatment allocations for each treatment by outcome is displayed in Table 3. While efficacy appears to favor atypical antipsychotics, side effect burden tends to favor perphenazine. As expected, the composite outcome occupies a middle ground between the two marginal outcomes; the optimal treatment for the composite outcome and efficacy are the same 92% of the time, the optimal treatment for the composite outcome and side effect burden are the same 64% of the time, and the optimal treatment for

efficacy and side effect burden are the same 56% of the time. Another view of these data is presented in Table 4 which shows the fraction of overlap between the proposed estimated optimal ITR and the optimal ITR based only on efficacy or side effect burden.

5. Discussion

Clinical and intervention scientists must often balance multiple, possibly competing outcomes when designing a personalized treatment plan for a patient. We combined a latent trait model with Q -learning to incorporate individual patient preferences into an evidence-based treatment plan formalized as an ITR. The proposed estimator is consistent for the optimal ITR and possesses an oracle property if the number of items on the preference questionnaire is allowed to diverge with the sample size; incidentally, to establish these results we developed new theory for the Rasch model that may be of interest beyond the application we consider here. We view the proposed methodology as an important first step toward mathematical formalizations of shared decision making in the context of precision medicine.

There are number of interesting and important extensions of the proposed methodology. The extension to settings with multiple treatment periods is of particular interest. This extension is challenging as a patient's preferences may change over time in response to treatment received and interim outcomes experienced, e.g., a patient who experiences severe side effects may develop a strong aversion to them in the future. Furthermore, in applications where interventions are delivered on a fine time scale, e.g., mobile-health, one must judiciously choose the timetable for administering the preference questionnaire to balance information gain with patient burden.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank the editor, associate editor and referees for their helpful comments which led to an improved paper. This work was supported in part by National Science Foundation Graduate Research Fellowship DGE-1144081 and National Institutes of Health Grant P01 CA142538.

References

- Barry MJ, Edgman-Levitan S. Shared decision making the pinnacle of patient-centered care. *New England Journal of Medicine*. 2012; 366:780–781. [PubMed: 22375967]
- Basu A, Meltzer D. Value of information on preference heterogeneity and individualized care. *Medical Decision Making*. 2007; 27:112–127. [PubMed: 17409362]
- Braziunas, D. Tech Rep. Department of Computer Science, University of Toronto; 2006. Computational approaches to preference elicitation.
- Brennan PF. Improving health care by understanding patient preferences. *Journal of the American Medical Informatics Association*. 1998; 5:257–262. [PubMed: 9609495]
- Collins FS, Varmus H. A new initiative on precision medicine. *New England Journal of Medicine*. 2015; 372:793–795. [PubMed: 25635347]

- Davis SM, Koch GG, Davis CE, LaVange LM. Statistical approaches to effectiveness measurement and outcome-driven re-randomizations in the clinical antipsychotic trials of intervention effectiveness (catie) studies. *Schizophrenia Bulletin*. 2003; 29:73–80. [PubMed: 12908662]
- Davis SM, Stroup TS, Koch GG, Davis CE, Rosenheck RA, Lieberman JA. Time to all-cause treatment discontinuation as the primary outcome in the clinical antipsychotic trials of intervention effectiveness (catie) schizophrenia study. *Statistics in Biopharmaceutical Research*. 2011; 3:253–265.
- De Ayala, RJ. The theory and practice of item response theory. Guilford Publications; 2013.
- Drake RE, Deegan PE, Rapp C. The promise of shared decision making in mental health. 2010
- Edwards, A., Elwyn, G. Shared decision-making in health care: Achieving evidence-based patient choice. Oxford University Press; 2009.
- Embretson, SE., Reise, SP. Item response theory. Psychology Press; 2013.
- Fan J, Li R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*. 2001; 96:1348–1360.
- Fishburn, PC. Nonlinear preference and utility theory. Vol. 5. Johns Hopkins University Press; Baltimore: 1988.
- Frank RG, Zeckhauser RJ. Custom-made versus ready-to-wear treatments: Behavioral propensities in physicians choices. *Journal of Health Economics*. 2007; 26:1101–1127. [PubMed: 18031852]
- Goldberg Y, Kosorok MR. Q-learning with censored data. *Annals of Statistics*. 2012; 40:529–560. [PubMed: 22754029]
- Hamburg MA, Collins FS. The path to personalized medicine. *New England Journal of Medicine*. 2010; 363:301–304. [PubMed: 20551152]
- Henderson R, Ansell P, Alshibani D. Regret-regression for optimal dynamic treatment regimes. *Biometrics*. 2010; 66:1192–1201. [PubMed: 20002404]
- Hodgkin D, Volpe-Vartanian J, Merrick EL, Horgan CM, Nierenberg AA, Frank RG, et al. Customization in prescribing for bipolar disorder. *Health economics*. 2012; 21:653–668. [PubMed: 21506194]
- Hogan TP, Awad A, Eastwood R. A self-report scale predictive of drug compliance in schizophrenics: reliability and discriminative validity. *Psychological medicine*. 1983; 13:177–183. [PubMed: 6133297]
- Huskamp HA, O'Malley AJ, Horvitz-Lennon M, Taub AL, Berndt ER, Donohue JM. How quickly do physicians adopt new drugs? the case of second-generation antipsychotics. *Psychiatric Services*. 2013
- Kay SR, Flszbein A, Opfer LA. The positive and negative syndrome scale (panss) for schizophrenia. *Schizophrenia bulletin*. 1987; 13:261. [PubMed: 3616518]
- Laber EB, Linn KA, Stefanski LA. Interactive model building for q-learning. *Biometrika*. 2014:asu043.
- Laber EB, Lizotte DJ, Ferguson B. Set-valued dynamic treatment regimes for competing outcomes. *Biometrics*. 2014; 70:53–61. [PubMed: 24400912]
- Laber EB, Zhao Y. Tree-based methods for individualized treatment regimes. *Biometrika*. 2015; 102:501–514. [PubMed: 26893526]
- Lieberman JA, Stroup TS, McEvoy JP, Swartz MS, Rosenheck RA, Perkins DO, et al. Effectiveness of antipsychotic drugs in patients with chronic schizophrenia. *New England Journal of Medicine*. 2005; 353:1209– 1223. [PubMed: 16172203]
- Linn, KA., Laber, EB., Stefanski, L. Estimation of dynamic treatment regimes for complex outcomes: balancing benefits and risks. In: Kosorok, MR., Moodie, EE., editors. *Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine*. Vol. chapter 15. SIAM; New York: 2016. p. 249-260.
- Lizotte DJ, Bowling M, Murphy SA. Linear fitted-q iteration with multiple reward functions. *The Journal of Machine Learning Research*. 2012; 13:3253–3295. [PubMed: 23741197]
- Lizotte DJ, Laber EB. Multi-objective markov decision processes for data-driven decision support. Under review. 2016; 1:1–25.

- Marler RT, Arora JS. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*. 2004; 26:369–395.
- Moodie EE, Chakraborty B, Kramer MS. Q-learning for estimating optimal dynamic treatment rules from observational data. *Canadian Journal of Statistics*. 2012; 40:629–645. [PubMed: 23355757]
- Moodie EE, Dean N, Sun YR. Q-learning: Flexible learning about useful utilities. *Statistics in Biosciences*. 2014; 6:223–243.
- Murphy SA. A generalization error for q-learning. *Journal of machine learning research: JMLR*. 2005; 6:1073. [PubMed: 16763665]
- Orellana L, Rotnitzky A, Robins JM. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part i: main content. *The International Journal of Biostatistics*. 2010; 6
- Qian M, Murphy SA. Performance guarantees for individualized treatment rules. *Annals of Statistics*. 2011; 39:1180–1210. [PubMed: 21666835]
- Ramsay, JO. *Functional data analysis*. Wiley Online Library; 2006.
- Rasch, G. *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*. Vol. 4. University of California Press; Berkeley, CA: 1961. On general laws and the meaning of measurement in psychology; p. 321-333.
- Rasch, G. *Probabilistic models for some intelligence and attainment tests*. Danish National Institute for Educational Research; 1980.
- Richter MK, Wong KC. Concave utility on finite sets. *Journal of economic theory*. 2004; 115:341–357.
- Rizopoulos D. ltm: An r package for latent variable modelling and item response theory analyses. *Journal of Statistical Software*. 2006; 17:1–25.
- Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000:550–560. [PubMed: 10955408]
- Rubin DB. Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*. 1978:34–58.
- Schulte PJ, Tsiatis AA, Laber EB, Davidian M. Q-and a-learning methods for estimating optimal dynamic treatment regimes. *Statistical science: a review journal of the Institute of Mathematical Statistics*. 2014; 29:640. [PubMed: 25620840]
- Shortreed SM, Laber E, Stroup TS, Pineau J, Murphy SA. A multiple imputation strategy for sequential multiple assignment randomized trials. *Statistics in Medicine*. 2014; 33:4202–4214. [PubMed: 24919867]
- Shortreed SM, Moodie EE. Estimating the optimal dynamic antipsychotic treatment regime: evidence from the sequential multiple-assignment randomized clinical antipsychotic trials of intervention and effectiveness schizophrenia study. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. 2012; 61:577–599.
- Smoller JW, Nierenberg AA. Small numbers, big impact. *Harvard review of psychiatry*. 1999; 7:109–113. [PubMed: 10471249]
- Song R, Kosorok M, Zeng D, Zhao Y, Laber E, Yuan M. On sparse representation for optimal individualized treatment selection with penalized outcome weighted learning. *Stat*. 2015; 4:59–68. [PubMed: 25883393]
- Sox, H., McNeil, B., Eden, J., Wheatley, B., et al. *Knowing What Works in Health Care: A Roadmap for the Nation*. National Academies Press; 2008.
- Stroup TS, McEvoy JP, Swartz MS, Byerly MJ, Glick ID, Canive JM, et al. The national institute of mental health clinical antipsychotic trials of intervention effectiveness (catie) project: schizophrenia trial design and protocol development. *Schizophrenia Bulletin*. 2003; 29:15–31. [PubMed: 12908658]
- Taylor JMG, Cheng W, Foster JC. Reader reaction to a robust method for estimating optimal treatment regimes by zhang et al.(2012). *Biometrics*. 2015; 71:267–273.
- Tian L, Alizadeh AA, Gentles AJ, Tibshirani R. A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association*. 2014; 109:1517–1532. [PubMed: 25729117]

- Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1996;267–288.
- Zhang B, Tsiatis AA, Davidian M, Zhang M, Laber E. Estimating optimal treatment regimes from a classification perspective. *Stat.* 2012; 1:103–114. [PubMed: 23645940]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012; 68:1010–1018. [PubMed: 22550953]
- Zhang B, Tsiatis AA, Laber EB, Davidian M. Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*. 2013; 100:681–694.
- Zhao Y, Zeng D, Laber EB, Kosorok MR. New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*. 2014
- Zhao Y, Zeng D, Laber EB, Song R, Yuan M, Kosorok MR. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika*. 2015; 102:151–168. [PubMed: 25937641]
- Zhao Y, Zeng D, Rush AJ, Kosorok MR. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*. 2012; 107:1106–1118. [PubMed: 23630406]
- Zhao Y, Zeng D, Socinski MA, Kosorok MR. Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics*. 2011; 67:1422– 1433. Received April 2016. [PubMed: 21385164]

Figure 1.

Averaged squared difference between $\hat{V}_L(\hat{\pi}_n)$ and $V_L(\pi^{\text{opt}})$. Lighter shades of gray correspond to larger error; it can be seen that the approximation improves as sample size or number of items increases.

Figure 2.

Averaged squared difference between $\hat{V}_Y(\hat{\pi}_Y)$ and $V_Y(\pi_Y^{\text{opt}})$. Lighter shades of gray correspond to larger error; it can be seen that the quality of the approximation improves with sample size but not the number of items.

Figure 3.

Averaged squared difference between $\hat{V}_Z(\hat{\pi}_Z)$ and $V_Z(\pi_Z^{\text{opt}})$. Lighter shades of gray correspond to larger error; it can be seen that the quality of the approximation improves with sample size but not the number of items.

Figure 4.

Average percentage disagreement between $\hat{\pi}_n$ and π^{opt} . Lighter shades of gray correspond to higher disagreement; it can be seen that agreement generally increases as n and p increase.

Figure 5. Average percentage disagreement between $\hat{\pi}_n$ and π^{oracle} . Lighter shades of gray correspond to higher disagreement; it can be seen that agreement generally increases as n and p increase.

Table 1

Patient preference questions

Question	Coding
"For me, the good things about medication outweigh the bad"	Yes (1) No (0)
"I feel weird, like a 'zombie', on medication"	Yes (0) No (1)
"Medications make me feel more relaxed"	Yes (1) No (0)
"Medications make me feel tired and sluggish"	Yes (0) No (1)
"I take medications only when Im sick"	Yes (0) No (1)
"I feel more normal on medication"	Yes (1) No (0)
"It is unnatural for my mind and body to be controlled by medications"	Yes (0) No (1)
"My thoughts are clearer on medication"	Yes (1) No (0)
"By staying on medications I can prevent being sick"	Yes (1) No (0)

Table 2

Coefficients of the Q-Function for Y and Z: CATIE Case Study

	Q_Y	Q_Z
Intercept	-5.51	13.04
Treatment	9.97	-2.09
Male (ref: Female)	-5.91	0.19
Black (ref: Other)	3.08	-0.80
White (ref: Other)	4.02	-0.14
Age	-0.11	-0.01
Baseline Clinical Severity of Schizophrenia	3.29	-0.19
Age at First Antipsychotic Medication	0.24	0.02
Olanzapine (at baseline, pre-randomization)	-3.77	-0.08
Quetiapine (at baseline, pre-randomization)	-5.91	0.24
Risperidone (at baseline, pre-randomization)	-1.48	0.06
Ziprasidone (at baseline, pre-randomization)	-0.94	0.60
Perphenazine (at baseline, pre-randomization)	-3.01	1.17
Haldol (at baseline)	-3.72	0.38
Decanoate (at baseline)	-7.68	0.73
BMI	-0.04	-0.01
Diastolic Blood Pressure	0.05	0.01
Systolic Blood Pressure	-0.10	0.0008
Treatment*Male (ref: Female)	-5.05	-0.21
Treatment*Black (ref: Other)	-1.86	-0.42
Treatment*White (ref: Other)	-2.49	0.25
Treatment*Age	0.05	0.01
Treatment*Baseline Clinical Severity of Schizophrenia	0.21	-0.01
Treatment*Age at First Antipsychotic Medication	-0.04	0.02
Treatment*Olanzapine (at baseline, pre-randomization)	-1.07	0.26
Treatment*Quetiapine (at baseline, pre-randomization)	-5.22	-0.08
Treatment*Risperidone (at baseline, pre-randomization)	-2.78	0.22
Treatment*Ziprasidone (at baseline, pre-randomization)	3.97	0.62
Treatment*Perphenazine (at baseline, pre-randomization)	-1.04	0.03
Treatment*Haldol (at baseline, pre-randomization)	-1.38	0.20
Treatment*Decanoate (at baseline, pre-randomization)	-0.13	1.60
Treatment*BMI	0.02	0.05
Treatment*Diastolic Blood Pressure	-0.01	0.01
Treatment*Systolic Blood Pressure	-0.05	-0.02

Table 3

Treatment recommendations

	Perphenazine	Atypical Antipsychotics
$\hat{\pi}_Y$: Efficacy	36%	64%
$\hat{\pi}_Z$: Side Effect Burden	67%	33%
$\hat{\pi}_\pi$: $\Phi(E)Y + \{1 - \Phi(E)\}Z$	41%	59%

Table 4

Percent of agreement in treatment recommendations

	$\hat{\pi}_Y$: Efficacy	$\hat{\pi}_Z$: Side Effect Burden
Both recommend perphenazine	35%	36%
Only $\hat{\pi}_n$ recommends perphenazine	6%	5%
Only $\hat{\pi}_n$ recommends atypical antipsychotic	2%	31%
Both recommend atypical antipsychotic	57%	28%