
Modeling Treatment Effect Modification in Multidrug-Resistant Tuberculosis in an Individual Patient Data Meta-Analysis

Yan Liu,¹ Mireille E. Schnitzer,^{2,3,*} Guanbo Wang,¹ Edward Kennedy,⁴ Piret Viiklepp,⁵ Mario H. Vargas,⁶ Giovanni Sotgiu,⁷ Dick Menzies^{8,9} and Andrea Benedetti^{1,8,10}

Abstract

Effect modification occurs while the effect of the treatment is not homogeneous across the different strata of patient characteristics. When the effect of treatment may vary from individual to individual, precision medicine can be improved by identifying patient covariates to estimate the size and direction of the effect at the individual level. However, this task is statistically challenging and typically requires large amounts of data. Investigators may be interested in using the individual patient data (IPD) from multiple studies to estimate these treatment effect models. Our data arise from a systematic review of observational studies contrasting different treatments for multidrug-resistant tuberculosis (MDR-TB), where multiple antimicrobial agents are taken concurrently to cure the infection. We propose a marginal structural model (MSM) for effect modification by different patient characteristics and co-medications in a meta-analysis of observational IPD. We develop, evaluate, and apply a targeted maximum likelihood estimator (TMLE) for the doubly robust estimation of the parameters of the proposed MSM in this context. In particular, we allow for differential availability of treatments across studies, measured confounding within and across studies, and random effects by study.

Keywords

Conditional average treatment effect, double robustness, individual patient data, marginal structural model, meta-analysis, multidrug-resistant tuberculosis, targeted maximum likelihood estimation

1 Introduction

Multidrug-resistant tuberculosis (MDR-TB), a form of tuberculosis (TB) with high mortality, is caused by bacteria resistant to at least the two most effective anti-TB drugs, isoniazid and rifampicin. According to the World Health Organization (WHO) (2020), a global total of 206 030 patients with MDR-TB or rifampicin-resistant (RR)-TB were detected and notified of their infection in 2019, a 10% increase compared to cases in 2018. However, the latest data reported to WHO show a treatment success rate for MDR/RR-TB of 57% globally.¹ Treating MDR-TB is challenging as a result of the heterogeneity of patients' characteristics (age, sex, HIV or other comorbidities), disease characteristics (extent and prior treatment), mycobacteria itself (different patterns of additional resistance) and characteristics of drugs (more toxicity and less effect of second-line drugs).^{2,3} Patients are typically prescribed a combination of four or more antimicrobial agents depending on the therapeutic phase and drug resistance pattern, if known.⁴ In addition, the effect of a treatment regimen may vary by an individual's characteristics and the specific combination of medications. When the effect and drug resistance pattern may vary from individual to individual, precision medicine can be improved by identifying patient covariates to estimate the size and direction of the effect at the individual level. In other words, identifying effect modifiers and assessing effect modification between different patient subgroups

¹ Department of Epidemiology, Biostatistics and Occupational Health, McGill University, Canada

² Faculty of Pharmacy, Université de Montréal, Canada

³ Department of Social and Preventive Medicine, Université de Montréal, Canada

⁴ Department of Statistics & Data Science, Carnegie Mellon University, USA

⁵ National Institute for Health Development, Estonia

⁶ Instituto Nacional de Enfermedades Respiratorias, Mexico.

⁷ Clinical Epidemiology and Medical Statistics Unit, Department of Medical, Surgical and Experimental Sciences, University of Sassari, Italy

⁸ Respiratory Epidemiology and Clinical Research Unit, Centre for Outcomes Research & Evaluation, Research Institute of the McGill University Health Centre, Canada

⁹ Montréal Chest Institute & McGill International TB Centre, Research Institute of the McGill University Health Centre, Canada

¹⁰ Department of Medicine, McGill University, Canada

Corresponding author:

Mireille Schnitzer, Faculty of Pharmacy, Université de Montréal, Montréal, Québec H3C3J7, Canada
Email: mireille.schnitzer@umontreal.ca

should be considered. However, this task is statistically challenging and typically requires large amounts of data so that treatment effects may be well-estimated for different combinations of covariate values. One may impose a working model in order to smooth (or summarize) the covariate-specific effects rather than estimate a separate effect for each possible combination of patient covariates.⁵⁻⁷ When working with observational data one must also adjust for all potential confounders of the treatment-outcome relationship, which can be accomplished via propensity scores and/or outcome regression modeling.⁸ One way to model effect modification in a simple binary treatment setting is through a marginal structural model (MSM) for the conditional average treatment effect (CATE).^{7,9} This model may be interpreted as the relationship between covariates and the expected treatment effect where the treatment effect is defined through a contrast of counterfactual outcomes. In non-meta-analytical settings, doubly robust estimators have been proposed for the estimation of a parametric MSM for the CATE⁷ as well as for nonparametric CATE models.¹⁰

Due to the large data requirements for estimating effects across patient subgroups, investigators may be interested in using individual participant data (IPD) from multiple studies to fit these treatment effect models. In our study, the data were extracted from 31 observational studies¹¹ which contrasted different treatment regimens for patients with MDR-TB, where multiple antimicrobial agents are taken concurrently by a patient over a long period. Our objective is to perform an IPD meta-analysis¹² to investigate the impact of different patient and treatment characteristics on the average treatment effect (ATE) of 14 anti-TB medications.

In this project, we propose and evaluate a targeted maximum likelihood estimator (TMLE)¹³ to model effect modification for the estimation of the parameters of the CATE MSM, in the IPD meta-analytical context. In the supplementary materials, we also describe and evaluate a novel augmented inverse probability of treatment weighted estimator (A-IPTW). For estimation of the ATE, TMLE depends on two components: an outcome regression conditional on treatment and covariates; and, weights comprised of the inverse of the propensity score where the propensity score is the probability of treatment conditional on covariates.¹⁴ TMLEs are plug in estimators with asymptotic properties that use a targeting step to optimize the bias-variance trade-off for the target parameter. In our setting, the estimator we propose allows for differential availability of

treatments across studies and random effects by study due to measured and unmeasured characteristics of the study-specific populations.

In Section 2, we describe the MDR-TB data structure and our parameters of interest. Section 3 introduces the TMLE procedure. We also describe a clustered influence function-based variance estimator. In Section 4, we present the results of simulation studies to demonstrate the properties of the estimator under different scenarios. Then in Section 5, we provide the results based on TMLE analysis of effect modification for 14 anti-TB medications using the combined IPD of the 31 observational studies.

2 Pooled Observational Studies of MDR-TB

The application data consist of IPD derived from 31 observational studies resulting in a total of 9290 MDR-TB patients. The data are available to any of the data contributors, but not publicly available. The systematic review¹¹ extended three previous systematic reviews.¹⁵⁻¹⁷ These IPD were collected from cohorts of adults with study years ranging from 1995 to 2009. The information collected for each patient includes demographics (age and sex), past TB history, clinical characteristics (pre-treatment sputum smear results for acid-fast bacilli (AFB) and culture, chest radiography, HIV infection), drug susceptibility test (DST) results, anti-microbial medications given, and outcomes. Our one-stage analytical approach pools all IPD across studies, accounting for clustering and random effects.

2.1 Data Structure

2.1.1 Outcome In the pooled dataset, the binary outcome Y represents the treatment success (the treatment was completed and cured the disease) versus treatment failure (the patient was still culture positive for MDR-TB, experienced a relapse or died).¹⁸ A patient's outcome realization is defined as lowercase y_{ij} where (i, j) refers to patient $i \in C_j$ in study j and C_j be the set of indices of patients in study j where $j \in (1, 2, \dots, J)$. In the MDR-TB data, there are pooled data from 31 studies, i.e. $J = 31$.

2.1.2 Treatments and Treatment Availabilities There are 14 antimicrobial agents observed in the data: ethambutol (EMB), ethionamide (ETO), ofloxacin (OFX), pyrazinamide (PZA), kanamycin/amikacin (KM/AM), cycloserine

(CS), capreomycin (CAP), para-aminosalicylic acid (PAS), prothionamide (PTO), streptomycin (SM), ciprofloxacin (CIP), later-generation fluoroquinolones (LgFQ), rifabutin (RIF) and group five level drugs (Gp5). LgFQ included levofloxacin, moxifloxacin, gatifloxacin and sparfloxacin.¹¹ Gp5 comprised of amoxicillin-clavulanate, macrolides (azithromycin, roxithromycin, and clarithromycin), clofazimine, thiacetazone, imipenem, linezolid, high dose isoniazid, and thioridazine.¹¹ We use k to index the k -th antimicrobial agent, where $k = 1, 2, \dots, 14$ in this case. The binary random variable $A^{(k)}$ indicates the exposure to medication k with patient realizations $a_{ij}^{(k)}$.

Not all treatments are observed in each study; we assume that a given treatment was available for a given study's participants based on whether we have observed that this treatment was taken by any patient in the given study.¹⁹ The binary variable $D^{(k)}$ is defined as the treatment availability of treatment k in a given study. Specifically, $d_{ij}^{(k)} = d_j^{(k)} = 1$ means that the treatment k is available to any subject i in study j , and is true if at least one patient in this study was prescribed the treatment k , and otherwise $d_{ij}^{(k)} = 0$.

2.1.3 Baseline Covariates and Resistance Information The baseline covariates consist of two study level covariates \mathbf{S} (the start year of MDR-TB treatment and the income group of the country of the study) and six individual level covariates \mathbf{W} (age, sex, AFB results, HIV infection, cavitation status on chest radiography and past TB history).

Resistance information based on DST is defined as the binary variable $R^{(k)}$. In this dataset, drug resistance information is available for eight medications. Thus, we denote $r^{(k)} = 1$ if the patient was found to be resistant to the treatment k and otherwise $r^{(k)} = 0$ which includes the situations where the patient's infection was susceptible to the treatment or was not known to be resistant to this treatment.

2.1.4 Observed Data Structure The observed data can be written as $\mathbf{O} = [\mathbf{S}, \mathbf{W}, \{A^{(k)}, D^{(k)}, R^{(k)}; k = 1, 2, \dots, 14\}, Y]$. We will rewrite $\mathbf{R} = \{R^{(k)}; k = 1, \dots, 14\}$, similarly for \mathbf{D} and \mathbf{A} . Then the data structure is $\mathbf{O} = (\mathbf{S}, \mathbf{D}, \mathbf{W}, \mathbf{R}, \mathbf{A}, Y)$.

2.1.5 Definition of Counterfactual Notation In this data, there is differential availability of treatments across studies. In order to define a generalized parameter of interest, we define counterfactual notation under the availability of a given treatment.¹⁹ We define counterfactual exposure to treatment $A^{(k)}\{d^{(k)} = 1\}$ as

the patient's counterfactual usage of treatment k had the patient had access. Then we define the counterfactual outcome $Y\{d^{(k)} = 1, a^{(k)}\} = Y\{a^{(k)}\}$ as the patient's outcome that would have occurred had treatment k been available and taken.

2.2 Parameter of Interest and Assumptions

As in previous work, we aim to estimate the parameters of interest defined with respect to a global population which refers to the union of super-populations specific to each study in this dataset.^{19,20} We define a non-parametric structural equation model (NPSEM), which assumes a time ordered data generating structure, in the Supplementary Materials Appendix A. Then the counterfactual likelihood is derived under NPSEM and the parameter of interest is defined with respect to this likelihood.

In this project, for a given medication k , we are interested in estimating the coefficients of potential effect modifiers $\mathbf{V}^{(k)}$ of the effect of medication k in an MSM for the CATE.⁷ To define the corresponding parameters, we recall that patients may take multiple medications concurrently. Given a medication k , we conduct the analysis by treating all other medications as confounders. For ease of notation, we define the adjustment set as

$$\mathbf{X}^{(k)} = [\mathbf{S}, \mathbf{W}, \mathbf{R}, \{D^{(k^*)}, A^{(k^*)}\}; \text{for } \forall k^* \in (1, \dots, 14) \text{ s.t. } k^* \neq k]$$

with realizations $x_{ij}^{(k)}$ for individual patients. Here, the symbol asterisk indicates treatments other than the treatment of interest. The set of potential effect modifiers investigated is a subset of the adjustment set $\mathbf{X}^{(k)}$.^{7,9} The random covariate vector is $\mathbf{V}^{(k)} = \{1, V_1^{(k)}, \dots, V_p^{(k)}\}$. Then we can model the CATE of treatment k , denoted $\psi\{\mathbf{V}^{(k)}\}$, as a linear function of potential effect modifiers such that

$$\mathbb{E}[Y\{a^{(k)} = 1\} - Y\{a^{(k)} = 0\} | \mathbf{V}^{(k)}] = \psi\{\mathbf{V}^{(k)}; \beta_{\mathbf{V}}^{(k)}\} = \{\mathbf{V}^{(k)}\}^\top \beta_{\mathbf{V}}^{(k)}$$

where the symbol \top indicates a transpose and $\beta_{\mathbf{V}}^{(k)} = \{\beta_0^{(k)}, \beta_1^{(k)}, \dots, \beta_p^{(k)}\} \in \mathbb{R}^{p+1}$. The parameter $\beta_m^{(k)}, m \neq 0$ represents the difference in the expected causal effect related to a single unit change of the effect modifier $V_m^{(k)}$ and $\beta_0^{(k)}$ is the intercept term.

Identifying the parameter of interest requires some assumptions that allow us to write the parameter in terms of distributions of the observed data.¹⁹

(a) Consistency: The counterfactual outcome $Y\{d^{(k)} = 1, a^{(k)} = 1\}$, where treatment k was available to and taken by the patient, is the same as the observed outcome for patients who in fact had access to the treatment and took it. Additionally, the counterfactual outcome had the patient not taken this treatment $Y\{a^{(k)} = 0\}$ is equal to the observed outcome for those who had not taken the treatment, either due to unavailability or other reasons. In particular, the first consistency assumption may fail if this treatment was only taken temporarily by the patient prior to a more effective antimicrobial being substituted in.

(b) Positivity: For the CATE to be estimable without extrapolation, we also need positivity assumptions. Specifically, the conditional probability of being treated for each patient given the drug's availability $Pr\{A^{(k)}\{d^{(k)} = 1\} = 1 | \mathbf{X}^{(k)} = \mathbf{x}^{(k)}, D^{(k)} = 1\}$ and the conditional probability of not being treated $Pr\{A^{(k)} = 0 | \mathbf{X}^{(k)} = \mathbf{x}^{(k)}\}$ must both be positive. This may fail if contraindications exist in the covariates $\mathbf{X}^{(k)}$, rendering treatment with $A^{(k)}$ impossible. Furthermore, $Pr\{D^{(k)} = 1 | \mathbf{S}^{(k)} = \mathbf{s}^{(k)}\}$, the probability of availability of treatment k conditional on the study-level covariates must also be positive. We also assume that the probability of treatment availability conditional on all measured covariates is only a function of the study-level covariates. Positivity is violated when, for example, certain studies occurred in a time or country where some drugs were not on the market, but only if time or country is a study-level confounder.

(c) Transportability: The counterfactual outcomes had the patient had access to the treatment and taken it are independent of treatment availability conditional on measured covariates, i.e. $Y^{(k)}\{d^{(k)} = 1, a^{(k)} = 1\} \perp D^{(k)} | \mathbf{X}^{(k)}$. This means that we can use the measured covariates to fit models where the treatment is available and use those model fits to estimate overall effects.

(d) In addition, we require unconfoundedness: that the counterfactual outcomes had the patient taken this treatment be independent of the treatment assignment conditional on the measured covariates and treatment availability. i.e. $Y^{(k)}\{a^{(k)} = 1\} \perp A^{(k)} | D, \mathbf{X}^{(k)}$. Moreover, the counterfactual outcomes had the patient not taken this treatment, $Y^{(k)}\{a^{(k)} = 0\}$, are independent of the treatment assignment given the measured covariates.

Under these assumptions, the CATE can be written as:

$$\psi\{\mathbf{V}^{(k)}\} = \mathbb{E}[\mathbb{E}\{Y|A^{(k)} = 1, \mathbf{X}^{(k)}\}|\mathbf{V}^{(k)}] - \mathbb{E}[\mathbb{E}\{Y|A^{(k)} = 0, \mathbf{X}^{(k)}\}|\mathbf{V}^{(k)}]$$

Both right-hand terms can be estimated from the observed data and thus the CATE is nonparametrically identifiable. A proof is provided in Supplementary Materials Appendix B.

3 Models and Algorithms

3.1 Outcome and Propensity Score Models

For ease of notation, we will drop the notation k for this section, and consider $A = A^{(k)}$ for a given k with $\mathbf{X} = \mathbf{X}^{(k)}$ the adjustment set for k as previously defined. The doubly robust estimators presented in this section require the estimation of two quantities, the conditional outcome expectation and propensity score. We define the former as $\bar{Q}(A, \mathbf{X}) = Pr\{Y(A = a) = 1|\mathbf{X}\}$, the probability of counterfactual treatment success conditional on the baseline covariates \mathbf{X} . The propensity score is $g(A|\mathbf{X}) = Pr(A = a|\mathbf{X})$, i.e. the conditional probability of treatment given \mathbf{X} .

In our setting with treatment availability variable D and the assumptions listed in Section 2.2, we note that

$$\bar{Q}(1, \mathbf{X}) = Pr(Y = 1|A = 1, \mathbf{X}) = Pr(Y = 1|D = 1, A = 1, \mathbf{X}).$$

Thus, we can estimate this quantity by fitting a regression model using the subgroup of subjects who received this treatment, which is necessarily a subset of those who had access to the treatment. On the other hand, $\bar{Q}(0, \mathbf{X}) = Pr(Y = 1|A = 0, \mathbf{X})$ includes both patients who did not have access to the given treatment and patients who did but did not receive this treatment.

For the propensity score, the probability of being treated for each patient is:

$$g(1|\mathbf{X}) = \underbrace{Pr(A = 1|D = 1, \mathbf{X})}_{g_1} \underbrace{Pr(D = 1|\mathbf{S})}_{g_2}.$$

The first component g_1 may be estimated by fitting a model using patients who had access to the treatment. The second part g_2 can be obtained by regressing

treatment availability on study level covariates \mathbf{S} . This second regression takes the study as the unit. Then, we can write the probability of not being treated as $g(0|\mathbf{X}) = 1 - g(1|\mathbf{X}) = 1 - g_1 \cdot g_2$.

3.2 Efficient Influence Function

The efficient influence function (EIF) for a specific parameter is the influence function that achieves the efficiency in the given space of semi-parametric models.²¹ The EIF defines the linear approximation of any efficient and regular asymptotically linear estimator. The EIF for the coefficients $\beta_{\mathbf{V}}$ in a working i.i.d. model space \mathcal{M} is

$$\mathcal{D}_{\beta_{\mathbf{V}}} = M^{-1}\mathcal{D}$$

where

$$\mathcal{D} = \left[\left\{ \frac{I_{A=1}}{g(1|\mathbf{X})} - \frac{I_{A=0}}{g(0|\mathbf{X})} \right\} \{Y - \bar{Q}(A, \mathbf{X})\} + \bar{Q}(1, \mathbf{X}) - \bar{Q}(0, \mathbf{X}) - \mathbf{V}^\top \beta_{\mathbf{V}} \right] \mathbf{V} \quad (1)$$

with normalizing matrix $M = -\mathbb{E}(\frac{\partial \mathcal{D}}{\partial \beta_{\mathbf{V}}})$. A proof is provided in Supplementary Materials Appendix C. An equivalent result was also given by Rosenblum and van der Laan.²²

3.3 TMLE

The general TMLE procedure was proposed by van der Laan and Rubin.²³ Our proposed procedure takes estimates of the conditional expected outcome $\bar{Q}(a, \mathbf{X})$ and updates them using information from the propensity score.

The first step is to produce initial estimates for $\bar{Q}(a, \mathbf{X})$ for $a = 1$ and 0 , denoted as $\bar{Q}_n(1, \mathbf{X})$ and $\bar{Q}_n(0, \mathbf{X})$, respectively. For each of $a = 1$ and 0 , we run a weighted logistic regression of Y with offset $\text{logit}\{\bar{Q}_n(a, \mathbf{X})\}$ and covariates corresponding to the set of potential effect modifiers. The weights are $A/g_n(1|\mathbf{X})$ for $a = 1$ and $(1 - A)/g_n(0|\mathbf{X})$ for $a = 0$, respectively. We then set the updated $\bar{Q}_n^*(a, \mathbf{X})$ equal to the predicted values from the above logistic regression. Finally, we fit a linear regression of $\bar{Q}_n^*(1, \mathbf{X}) - \bar{Q}_n^*(0, \mathbf{X})$ on the potential effect modifiers in order to obtain the TMLE estimates, $\hat{\beta}_{\mathbf{V}}^{TMLE}$, of the parameters of interest.

We note that this TMLE solves the equation $\sum_{j=1}^J \sum_{i \in C_j} \mathcal{D}_{ij,n}(\hat{\beta}_{\mathbf{V}}^{TMLE}) = 0$. The estimator thus has the properties of double robustness and local efficiency.

We also developed a closely related A-IPTW²⁴, which is described in the Supplementary Materials Appendix D. In finite samples, TMLE has been shown to perform better than a related A-IPTW if the positivity assumption is nearly violated or the true values of the parameters of interest are close to the parameter space boundaries.²⁵

3.4 Influence Function-based Variance Estimator

Standard errors can be estimated for the TMLE using a large-sample sandwich estimator of the efficient influence function under consistency of $\bar{Q}_n(a, \mathbf{X})$ and $g_n(a|\mathbf{X})$.²¹ Let $\beta_{\mathbf{V}_0}$ be the true value of $\beta_{\mathbf{V}}$ and $\hat{\beta}_{\mathbf{V}}$ be the TMLE estimate.

Under regularity conditions, we can write the linear approximation of the estimator as¹⁹

$$\sqrt{n}(\hat{\beta}_{\mathbf{V}} - \beta_{\mathbf{V}_0}) \approx \frac{1}{\sqrt{n}} \sum_{j=1}^J \sum_{i \in C_j} \mathcal{D}_{ij}(\beta_{\mathbf{V}_0})$$

where $\mathcal{D}_{ij}(\beta_{\mathbf{V}_0})$ is the influence function at the true values of $\bar{Q}(a, \mathbf{X}_{ij})$ and $g(a|\mathbf{X}_{ij})$ for each subject.

In order to estimate the variance of the parameter of interest while taking clustering by studies into account, we only assume independence between studies, and not individuals within the same study. Within study j , we denote the $(p+1) \times (p+1)$ dimension variance-covariance matrix of the efficient influence function as σ_j^2 . We denote the $(p+1) \times (p+1)$ dimension variance-covariance matrix of the efficient influence function of any two different subjects in study j as ρ_j . Both of these quantities can be estimated from the observed data. Then, for large J , the variance-covariance matrix of $\hat{\beta}_{\mathbf{V}}$ can be estimated using:²⁶

$$\begin{aligned} \text{Var}(\hat{\beta}_{\mathbf{V}}) &\approx \frac{1}{n^2} \text{diag} \left\{ \sum_{j=1}^J \left(\sum_{i,m \in C_j, i \neq m} \mathbb{E}[\mathcal{D}_{ij,n}(\hat{\beta}_{\mathbf{V}}) \{\mathcal{D}_{mj,n}(\hat{\beta}_{\mathbf{V}})\}^\top] \right. \right. \\ &\quad \left. \left. + \sum_{i \in C_j} \mathbb{E}[\mathcal{D}_{ij,n}(\hat{\beta}_{\mathbf{V}}) \{\mathcal{D}_{ij,n}(\hat{\beta}_{\mathbf{V}})\}^\top] \right) \right\} \\ &= \frac{1}{n^2} \text{diag} \left[\sum_{j=1}^J \left\{ n_j(n_j - 1)\rho_j + n_j\sigma_j^2 \right\} \right] \end{aligned}$$

where n_j is the size of study C_j . This variance estimator is only valid for larger numbers of clusters and consistent estimates of both the outcome and the

propensity score.¹⁹ We investigate its finite-sample performance in the simulation study. In past work, we have also found that the nonparametric clustered bootstrap performs well.^{19,26}

4 Simulation Study

We conducted simulation studies in R to develop and evaluate TMLE and A-IPTW in the implementation to model effect modification for IPD meta-analysis. We demonstrate the double robustness and finite-sample performance of TMLE and A-IPTW. Results for the A-IPTW estimator are provided in the Supplementary Materials Appendix E.

4.1 Data Generation

For each dataset, we generated two continuous study-level covariates, S_1 and S_2 , where we treated S_2 as unobserved. We also generated three individual-level covariates: W_1 continuous and W_2 and W_3 binary. In this simulation, we included three treatments where each study had access to one, two or all three. We denote the treatment availability as $D^{(k)}$, $k = 1, 2, 3$ which was generated conditional on the study level covariate S_1 . Then three treatment indicators $A^{(k)}$ were generated based on the values of S_1, W_1, W_2, W_3 and $D^{(k)}$. Subjects could take any combination of these treatments. Finally, we generated a continuous outcome Y in a model with and without random effects, conditional on individual-level covariates, all three treatments, and study-level covariate S_1 . Specifically, this model included interactions between the two effect modifiers W_1 and W_3 and $A^{(1)}$. For random effects by study we added an additional interaction term between S_2 and $A^{(1)}$. Table S1 in Supplementary Material Appendix E displays the full data generating mechanism. The observed data structure for each subject is $O = (S_1, W_1, W_2, W_3, \mathbf{D}, \mathbf{A}, Y)$.

In this simulation study, we only aimed to estimate the parameters representing effect modification of treatment $k = 1$. For both scenarios with the outcomes involving random effects or not, we drew 1000 simulations with $J \in \{10, 30, 50\}$ studies where each study contained 300 subjects. This resulted in three total sample sizes, $n = 3000, 9000, 15000$.

4.2 Analysis

In order to model the effect modification of treatment $k = 1$, we included treatment covariates $A^{(2)}$ and $A^{(3)}$ as confounders, as explained in Section 2.2. We denote the baseline covariates as $\mathbf{X}^{(1)} = \{S_1, W_1, W_2, W_3, A^{(2)}, A^{(3)}\}$. In practice we are not aware of the true set of effect modifiers so we define the potential set as $\{W_1, W_2, W_3, A^{(2)}, A^{(3)}\}$ and set $\mathbf{V}^{(1)} = \{1, W_1, W_2, W_3, A^{(2)}, A^{(3)}\}$.

As we discussed in Section 2.2, we model the CATE of treatment $A^{(1)}$ as a linear function of potential effect modifiers: $\psi\{\mathbf{V}^{(1)}; \boldsymbol{\beta}_{\mathbf{V}^{(1)}}\} = \{\mathbf{V}^{(1)}\}^\top \boldsymbol{\beta}_{\mathbf{V}^{(1)}} = \beta_0^{(1)} + \beta_1^{(1)}W_1 + \dots + \beta_5^{(1)}A^{(3)}$. Our parameters of interest are thus $\boldsymbol{\beta}_{\mathbf{V}^{(1)}}$. In the scenarios without random effects, the true values of these parameters were derived analytically. For the scenarios with random effects, we generated data with large sample sizes (10^6) and forced all $A^{(1)}$ equal to 1 and 0, respectively, to obtain both counterfactual outcomes. Then we obtained the true values of the parameters by fitting the linear regression of the differences between the two counterfactual outcomes on the potential effect modifiers. Given the data generating mechanism there were two real effect modifiers, W_1 and W_3 . When outcomes were simulated without random effects, the true corresponding coefficient values were 0.65 and 0.35, respectively. When outcomes were generated with random effects, the true values were 0.77 and 0.38 (Tables S2, S3 in the Supplementary Materials Appendix E).

We used logistic regressions to estimate each component of $g\{A|\mathbf{X}^{(1)}\}$, collectively referred to as g , and $\bar{Q}\{a, \mathbf{X}^{(1)}\}$, referred to as Q . To control the potential sources of sparsity, predicted values for g_1 and g_2 were truncated at $(\alpha, 1 - \alpha)$ where $\alpha = 0.001$. To demonstrate the double robustness property of both methods, we ran four different scenarios: 1) correctly specified parametric models for both Q and g ; 2) only the Q model correctly specified and the g model misspecified as a null model; 3) only the g model correctly specified and the Q model misspecified as a null model; 4) both Q and g misspecified as null models.

In this simulation, we applied the TMLE algorithms presented in Section 3.3. Note that the TMLE implementation we used requires that we transform the continuous Y to lie in $(0, 1)$, and then we reverse-transform at the end of the procedure.²⁷ The standard errors were estimated by the influence function sandwich estimator, first ignoring clustering and then incorporating clustering as described in Section 3.4. We compared both standard error estimates to the Monte Carlo standard errors. Then based on the clustered standard errors, we

constructed 95% Wald-type confidence intervals and computed the corresponding coverage rates which are the percentage of times that the confidence intervals contained the true parameter values.

4.3 Results

Figure 1 presents the results of TMLE for the estimation of five potential effect modifiers under the four estimation scenarios and three different sample sizes with outcomes generated without random effects. Figure 2 presents the TMLE results under random effects. The coverage rates are presented in the blue boxes of the two figures. More detailed TMLE results are provided in the Supplementary Materials Appendix E (Tables S2 - S4), along with the figures and tables of the results of the A-IPTW estimator (Tables S5 - S7 and Figures S1, S2).

From Figures 1 and 2, we see that under the first two scenarios, where the model for Q was correctly specified, the TMLE estimators had no error on average regardless of the presence of random effects. Unsurprisingly, under scenario 4 when all quantities were assigned null models, a small bias was present without random effects and a larger bias was present with random effects. In scenario 3, where the model for Q was incorrectly specified but the model for g was correct, the average error converged to zero as the number of studies increased. The estimates were also more dispersed than for the previous scenarios. This occurred because the model for g_2 is estimated using the study as the unit in the analysis, so with few studies, the sample size to fit this model is extremely small. Compared to TMLE, the A-IPTW estimator had greater average error in scenarios 3 and 4, but otherwise performed similarly (Table S8 in the Supplementary Materials Appendix E).

The coverage rates, given in the blue boxes of Figures 1 and 2, typically increased with the number of studies. For instance, with random effects with 10 studies, in scenario 1 where all models were correctly specified, the coverage rates for the five potential effect modifiers were between 79.8% and 90.1%. With 30 studies, the rates increased to 88.1% – 94.8%. Then for 50 studies, the coverage rates were 89.6% – 96.7%. Similar patterns were also obtained in the second and third scenarios.

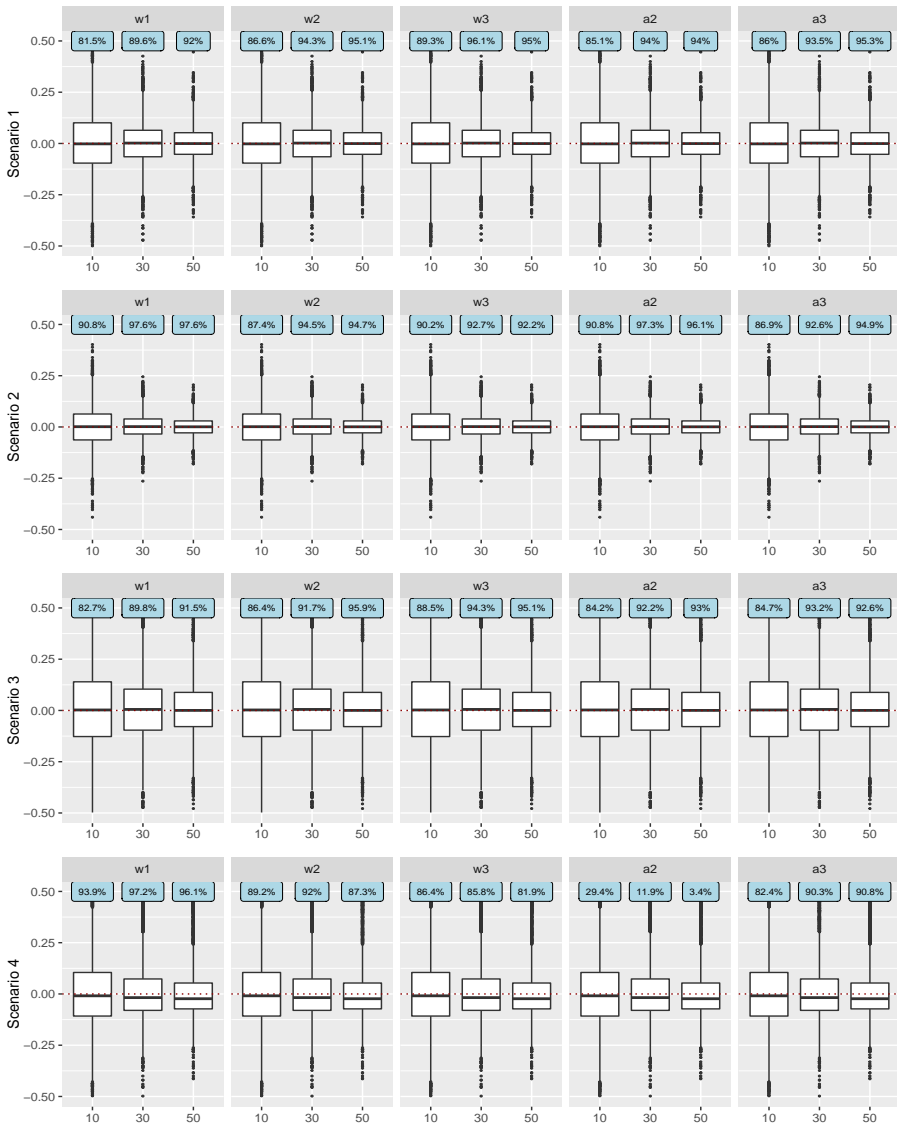


Figure 1. Error of TMLE estimates under four scenarios and three different sample sizes without random effects. The x -axis represents the number of studies. Coverage rates based on the clustered sandwich estimators of the standard error are presented in blue boxes. The four scenarios are as follows: Scenario 1 - both Q and g models are correct; Scenario 2 - Q model is correct, g is null; Scenario 3 - Q model is null, g model is correct; Scenario 4 - both Q and g models are null.

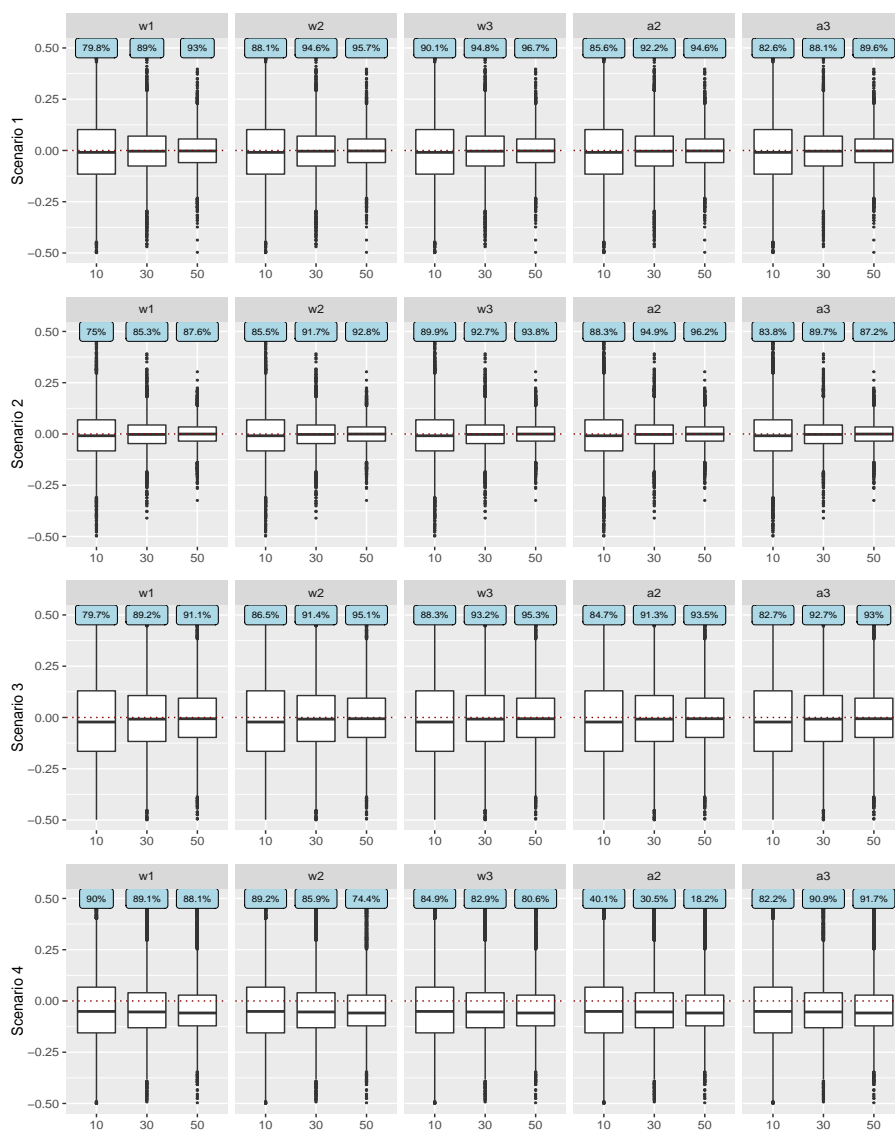


Figure 2. Error of TMLE estimates under four scenarios and three different sample sizes with random effects. The x -axis represents the number of studies for three sample sizes. Coverage rates based on the clustered sandwich estimators of the standard error are presented in blue boxes. The four scenarios are as follows: Scenario 1 - both Q and g models are correct; Scenario 2 - Q model is correct, g is null; Scenario 3 - Q model is null, g model is correct; Scenario 4 - both Q and g models are null.

5 MDR-TB Data Analysis

5.1 Descriptive Statistics of MDR-TB

The combined dataset contained the IPD from 31 observational studies with a total of 9290 patients. After removing 260 (2.8%) patients without reported outcomes, 9030 patients taking combinations of 14 different antimicrobial agents remained. Study-specific sample sizes ranged from 25 to 2182 patients (Supplementary Materials Appendix F Table S9). Missing values were present in the covariates; the cavity status variable had the most missing values (25.97%). Appendix F Table S10 and Figure S3 give summaries of the treatment-specific sample sizes and the six individual level covariates. Ofloxacin, pyrazinamide and cycloserine were the three most prescribed medications. In contrast, fewer than 1000 patients were prescribed later-generation fluoroquinolones and ciprofloxacin. Ethambutol and pyrazinamide were widely prescribed in 30 out of 31 studies, while only 14 studies had patients who took ciprofloxacin or later-generation fluoroquinolones. The number of male patients was around twice the number of female patients in all treatment groups. A total of 4892 (54.33%) patients were concentrated in the 26 – 45 year-old age group while only 154 (1.71%) were in the 0 – 17 year-old group. For the other individual level covariates, about 75% of patients were diagnosed with TB in the past. Moreover, the majority of patients had cavity (68.33%) and positive AFB (74.53%), but there were only 15% coinfecting with HIV.

In order to assess the data support needed to investigate effect modification by concurrent medication, Table 1 displays the number of patients who used a combination of any two medications, with the diagonal indicating the total number of patients taking the corresponding medication. Values ranged between 87 to 4574, indicating at least minimal data support for all pairwise combinations.

Table 1. Summary of number of patients taking any combinations for any two medications during the treatment period. The diagonal values represent the total number of patients taking each medication.

	EMB	CAP	CIP	CS	ETO	OFX	PAS	PTO	RIF	SM	PZA	KM/AM	LgFQ	Gp5
EMB	4188	734	565	1617	2472	2995	989	770	1080	634	3518	2706	272	768
CAP	734	1874	482	1719	900	1328	1386	722	325	151	1119	543	205	923
CIP	565	482	968	782	682	236	616	223	350	232	645	481	127	555
CS	1617	1719	782	5629	1941	4195	3573	3004	523	981	3234	2900	745	1833
ETO	2472	900	682	1941	3911	3175	1206	240	397	309	3433	3014	287	670
OFX	2995	1328	236	4195	3175	6464	2750	2566	465	786	4574	4191	192	1262
PAS	989	1386	616	3573	1206	2750	3937	2292	293	732	1962	1816	644	1463
PTO	770	722	223	3004	240	2566	2292	3304	154	749	1564	1532	449	1065
RIF	1080	325	350	523	397	465	293	154	1261	406	1133	332	87	195
SM	634	151	232	981	309	786	732	749	406	1366	870	192	269	339
PZA	3518	1119	645	3234	3433	4574	1962	1564	1133	870	6102	3775	436	930
KM/AM	2706	543	481	2900	3014	4191	1816	1532	332	192	3775	5015	416	1166
LgFQ	272	205	127	745	287	192	644	449	87	269	436	416	866	511
Gp5	768	923	555	1833	670	1262	1463	1065	195	339	930	1166	511	2138

5.2 Analysis and Results of MDR-TB

For each medication k , the target parameters in this application are the coefficients of the MSMs,

$$\psi\{\mathbf{V}^{(k)}; \boldsymbol{\beta}_{\mathbf{V}^{(k)}}\} = \{\mathbf{V}^{(k)}\}^\top \boldsymbol{\beta}_{\mathbf{V}^{(k)}} = \beta_0^{(k)} + \sum_{j=1}^{19} V_j^{(k)} \beta_j^{(k)}$$

where $\mathbf{V}^{(k)} = \{1, V_1^{(k)}, \dots, V_{19}^{(k)}\}$ is the set including six individual level covariates and the 13 medications excluding medication k . We standardized the continuous variable age. $\beta_0^{(k)}$ represents the baseline effect for the reference group of female patients with mean age (39), negative AFB test, no cavitation on chest radiography, no past TB history, no HIV co-infection and not taking any other of the 13 medications. The associated coefficients $\beta_j^{(k)}, j = 1, \dots, 19$ represent the difference in the CATE when varying the characteristic $V_j^{(k)}$ by one unit while holding other covariates fixed.

As noted, there are missing values in the individual-level covariates. We used multiple imputation²⁸ by chained equations with the MICE package²⁹ in R to produce 20 imputations then used Rubin’s rules to combine the estimates.³⁰ In each imputed dataset we followed the procedure described in Section 3.3 to fit a TMLE. We estimated the Q component using SuperLearner (SL)³¹ which is a methodology that uses cross validation to find an optimal convex combination of the predictions of a library of candidate algorithms defined by the user. We included the following algorithms in the SL library: generalized linear models with penalized maximum likelihood (`glmnet` function)^{32,33}, with forward stepwise variable selection (`step` function), and with a stepwise procedure based on the Akaike Information Criterion (`stepAIC` function), respectively.³⁴ Logistic regressions were used for the g models and LASSO penalties were added when the logistic regression failed to converge.

In each completed dataset, variance estimates for the coefficients were computed using the sample variance of the influence function following the expression in Section 3.4. Finally, since many comparisons made in this analysis, we performed a multiple testing adjustment of the “significance level” of the p-values to control the false discovery rate via the method of Benjamini and Hochberg³⁵ with further details given in the Supplementary Materials Appendix G.

Figures 3, 4 and 5 show the estimated coefficients, standard errors and 95% confidence intervals corresponding to the intercept and six individual-level effect modifiers and the other 13 medications. The outputs for estimated coefficients that are statistically significant after adjustment are marked in red color and with an asterisk. Corresponding tables of the numerical results are also provided in the Supplementary Materials Appendix F (Tables S11 and S12). From Figure 3 there is no evidence that the six characteristics modify the treatment effect of any drug. In Figure 4 (CS plot), patients prescribed streptomycin might have benefited less from cycloserine than patients not receiving streptomycin. In addition, taking cycloserine was associated with a greater estimated effect of ethionamide while capreomycin, kanamycin/amikacin drugs were associated with lower estimated effects (ETO plot).

The empirical distributions of the untruncated propensity scores for all drugs are provided in the Supplementary Materials Appendix F (Table S13 and Figure S4). For later-generation fluoroquinolones, we noted very large weights which likely yielded the large variability observed in Figure 5, LgFQ plot. In addition, since fewer patients were prescribed later-generation fluoroquinolones among those who were HIV positive (Appendix F Figure S3), the standard errors were inflated for the coefficient of HIV in the fluoroquinolones MSM (Figure 3, LgFQ plot). Also, among those who took rifabutin, only 87/866 subjects were also prescribed later-generation fluoroquinolones, giving rise to the large standard error (Figure 5, RIF plot).

To highlight the differences between meta-analysis and individual study results, we compared the results of Mitnick et al., conducted in 1996-2002^{36,37} for the effect modification of ethionamide. Given that all 14 medications were observed in only two studies (Appendix F Table S9), Mitnick et al. (710 subjects) and Tupasi et al. (170 subjects), we chose Mitnick et. al. to make the comparison. In Figure 6, we see that a single-study TMLE analysis of the data from the study by Mitnick et al. concluded that prothionamide was associated with greater estimated effects of ethionamide while group 5 drugs and rifabutin were associated with lower estimated effects. Both the meta-analysis and the single study results agreed that cycloserine and capreomycin were effect modifiers (positive and negative, respectively). As expected, the estimation in the meta-analysis had lower standard errors than the individual study (Figure 6). The deviation of the results suggests heterogeneity between different studies in the meta-analysis. It is

important to note that analysis of an individual study targets a parameter that is interpreted in the individual study's super-population while the meta-analysis targets the parameter interpreted in the global population. These parameters may not coincide when there is heterogeneity.

6 Discussion

In this paper, we developed one-stage doubly-robust methods for the analysis of baseline effect modification in an IPD meta-analysis. The model space that we considered was nonparametric though the parameters of interest were defined through a working linear MSM for the CATE. Our approach allowed us to analyze pooled IPD from multiple studies in order to evaluate how estimated treatment effects may vary depending on the values of patient covariates. Our past work, which proposed related methods and a TMLE for IPD meta-analysis with multiple treatments, instead estimated a treatment importance metric defined as the difference in adjusted probabilities of treatment success between the patients who used each medication and the overall population.¹⁹ Our current work proposes a novel TMLE and A-IPTW for the estimation of effect modification in the described marginal structural model in an IPD meta-analysis.

Vo et al. illustrated that in IPD meta-analysis, heterogeneity across studies can come from two sources: case-mix heterogeneity, due to effect modification, and beyond case-mix heterogeneity, due to differences in study design and measurement.³⁸ Our methods address heterogeneity by allowing for differential availability of treatments across studies and random effects by study due to measured and unmeasured characteristics of the study-specific populations.

In clinical and epidemiological research, model misspecification is always a concern when estimating treatment or exposure effects. Doubly robust methods yield consistent estimators even under misspecification of either the treatment or the outcome model. In the simulation study, we demonstrated the double robustness property of both TMLE and A-IPTW. We observed similar performance of TMLE and A-IPTW but we did not investigate near-positivity violations or other scenarios that may differentiate them in finite samples as have others.^{25,39} In addition, we demonstrated the double robustness of both methods when there are study-specific random effects for the outcome. Finally, we showed that the proposed confidence intervals, estimated using the influence

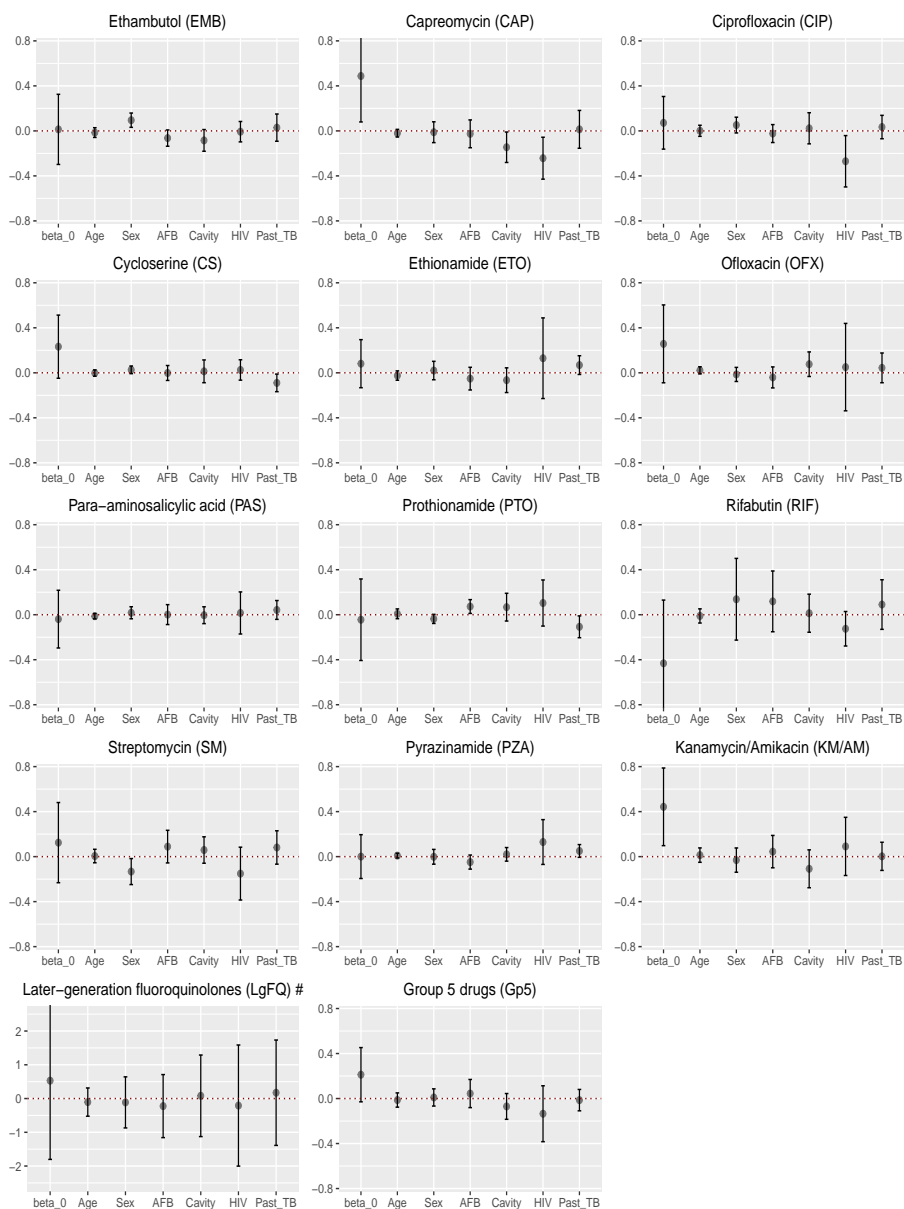


Figure 3. Estimated coefficients and the corresponding 95% confidence interval for 14 medications relative to the intercept and six demographic or clinical covariates. None of the coefficients reached statistical significance.

Larger scale for the y -axis of the LgFQ plot.

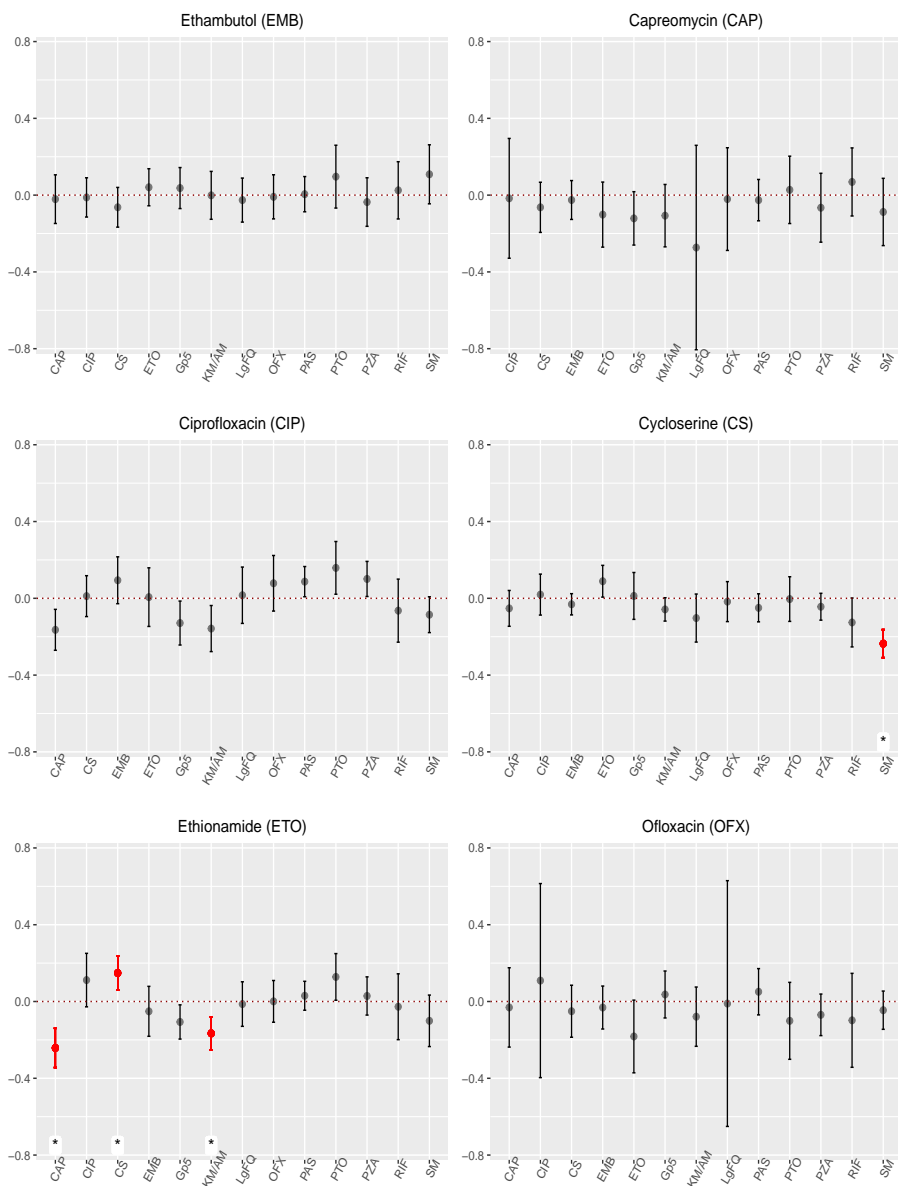


Figure 4. Estimated coefficients of potential effect modifiers and the corresponding 95% confidence intervals for EMB, CAP, CIP, CS, ETO and OFX. Significant results are shown in red and indicated with an *.

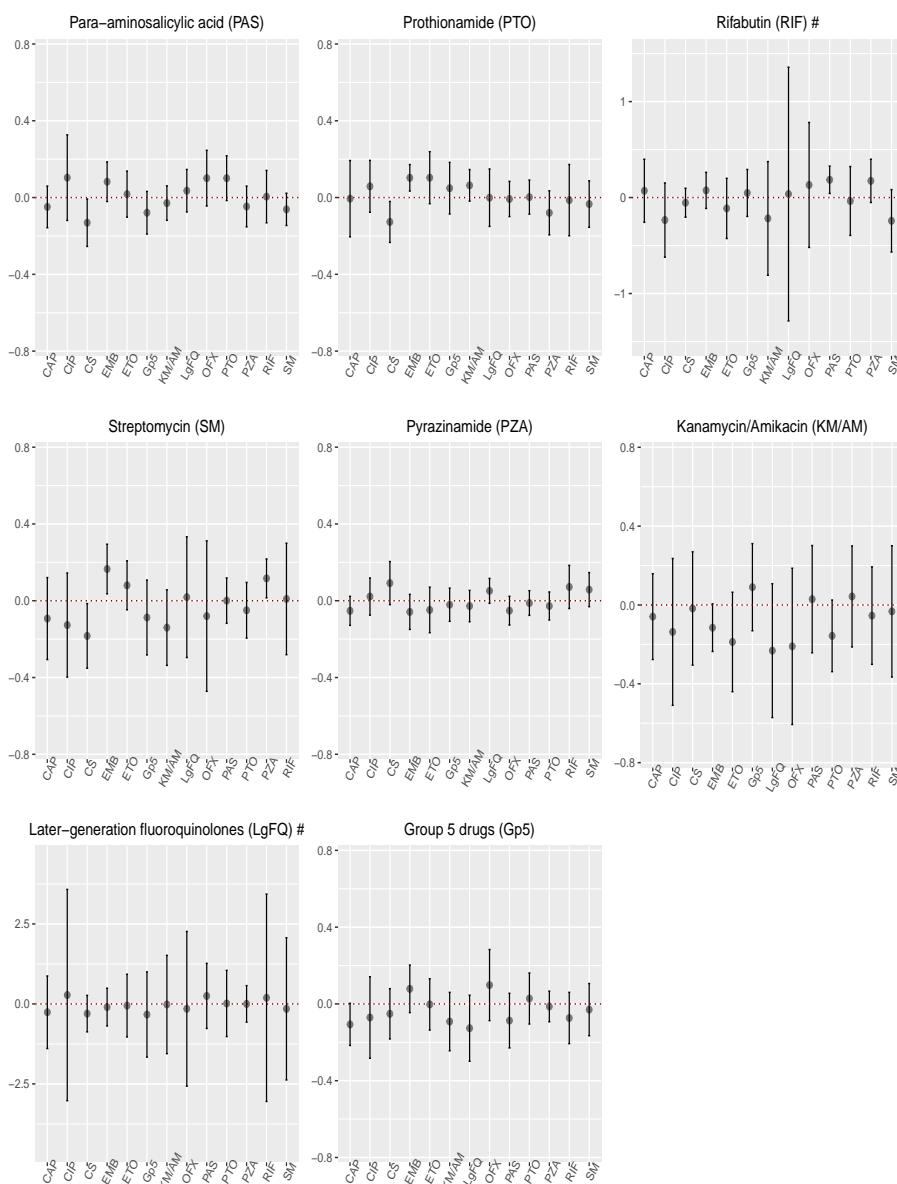


Figure 5. Estimated coefficients of potential effect modifiers and the corresponding 95% confidence intervals for PAS, PTO, RIF, SM, PZA, KM/AM, LgFQ and Gp5. None of the coefficients reached statistical significance.

Larger scale for the y -axis of the RIF and LgFQ plots.

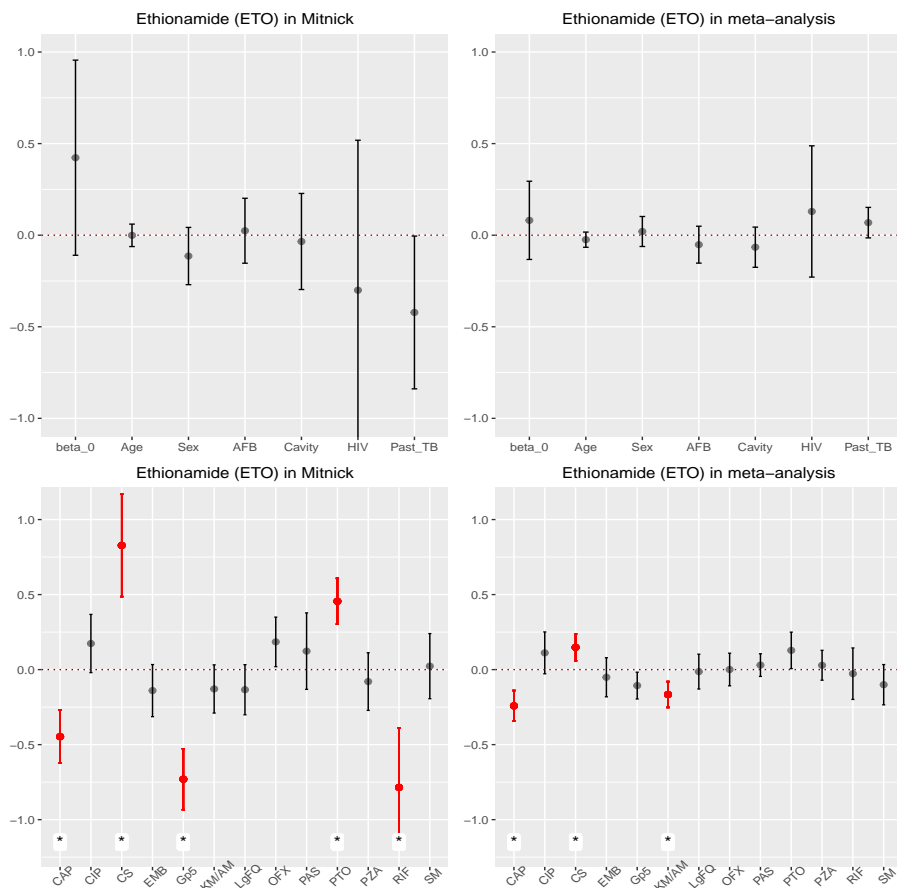


Figure 6. Estimated coefficients of potential effect modifiers and the corresponding 95% confidence intervals for ETO in the study by Mitnick et al. (left column) and in the meta-analysis (right column). Significant results are shown in red and indicated with an *.

function sandwich estimator that considers clustering by study, performed well when there were greater than 30 studies in the analysis. Indeed, a limitation of our approach is that it relies on a sufficient number of studies to estimate the generalized parameter interpreted in the global population. In particular, the ability to adjust for confounding by treatment availability depends on fitting a model for treatment availability conditional on study-level covariates, which is limited by the typically small number of studies in a meta-analysis. Indeed, we observe in the simulation study that error may persist when the number of studies

is small and the outcome regression model is incorrectly specified, even when the propensity score model components are correctly specified with parametric models. Therefore, our approach should only be undertaken when a larger number of studies are available.

The treatment of MDR-TB is challenging because of its prolonged duration, toxicity, costs and unsatisfactory outcomes.¹ Second-line TB medicines used for the treatment of drug-resistant TB include injectable drugs (capreomycin, streptomycin, kanamycin or amikacin), fluoroquinolones (ciprofloxacin, ofloxacin or later-generation fluoroquinolones), ethionamide or prothionamide, para-aminosalicylic acid and cycloserine. The WHO (2019) revised the guidelines and currently suggest that streptomycin and amikacin are to be considered only if DST results confirm susceptibility and adequate measures to monitor for adverse reactions can be ensured. In addition, capreomycin and kanamycin are not to be included in the treatment of MDR-TB patients.⁴ Currently, in general, the selection of antimicrobials for the treatment of an individual patient with TB is based, apart from drug availability, on issues such as the mycobacteria's drug-resistance pattern and assumptions about the chosen drugs (mechanism of action, potential toxicity, known pharmacological interactions, possible development of drug resistance due to previous use, etc.). However, the actual role played by each drug in the therapeutic outcome is difficult to assess because of the large number of possible combinations of antimicrobials and the potential yet undiscovered interactions affecting their pharmacokinetic and/or pharmacodynamic properties. The statistical approach proposed by the present study takes advantage of the IPD meta-analysis in order to unveil treatment effect modifiers, either regarding medications or other variables at the individual-level, with the aim to better treat and understand not only TB, but also other diseases, including those requiring the concurrent use of several drugs such as arterial hypertension, diabetes, etc.

As discussed in Section 2.2, we require unconfoundedness when estimating the causal effect. However, in practice, this condition is non-testable and we are limited to the covariates collected by the original studies. While most important confounders were collected across all studies, we did not receive DST results from many studies. However, the results were available to clinicians who may have responsively changed or added medications, making DST results an important confounder. Our analyses adjusted for all available drug sensitivity information.

Our results suggest that cycloserine may enhance the effects of ethionamide but capreomycin and kanamycin or amikacin were associated with reduced effects. In addition, streptomycin may reduce the effect of cycloserine. Our findings support the revised WHO guidelines that injectable drugs are not as beneficial as previously believed. The analyses informing the 2019 WHO guidelines (Ahmad et al.) showed that kanamycin and capreomycin were associated with worse outcomes.⁴⁰ In our results, these drugs were associated with worse outcomes in regimes with ethionamide. For the six individual characteristics, age, sex, acid fast bacilli status, HIV infection, cavitation status on chest radiography, and past TB history, they might be potential for effect modifiers since we did not find any evidence of effect modification. Being that MDR-TB an infectious disease, this finding may simply reflect that antimicrobial drugs play a much larger and definitive role for success than other characteristics at the individual level.

Since our MDR-TB data were identified from studies carried out up to 2009, we have no information about both new and repurposed effective anti-TB drugs.^{41,42} Upcoming work will apply our methods to data from more recently treated patients. Another limitation of our approach is that, because we only considered the effect of intervening on one treatment at a time, we cannot directly address how to select combinations of medications that would be expected to optimize the probability of treatment success. Previous work evaluated the causal contrasts between different regimens of concurrent medications in MDR-TB.⁴³ Future applications should directly address the more challenging question of treatment-treatment interactions on the outcome which would directly allow for the evaluation of optimal medication usage. Other ongoing work in our group involves using LASSO,³³ rather than hypothesis testing, to select the effect modifiers in the linear MSM for the CATE. This may improve upon the current work by utilizing a superior approach to variable selection.

Identifying effect modifiers is an important step for estimating subpopulation causal effects that can help guide treatment decision making for individual patients. However, such analyses require larger amounts of data than the estimation of average treatment effects. This analytic approach can generate hypotheses for drug combinations which can be tested in randomized controlled trials. We have also contributed by extending existing doubly robust methods to incorporate multiple data sources. Advances in IPD meta-analysis enable researchers to incorporate multiple sources of previously collected observational

data in their analyses in order to increase their power, which is greatly beneficial for the identification of effect modifiers.

Acknowledgements

This work was supported by the Canadian Institutes of Health Research [Project Grant to MES and AB, and a New Investigators Salary Award to MES]. The data and context were provided by the Collaborative Group for Meta-Analysis of Individual Patient Data in Multidrug-Resistant Tuberculosis. Conflict of Interest: None declared.

References

1. World Health Organization. *Global tuberculosis report 2020*. Geneva: World Health Organization, 2020.
2. Müller B, Borrell S, Rose G, et al. The heterogeneous evolution of multidrug-resistant Mycobacterium tuberculosis. *Trends Genet.* 2013; **29**:160-169.
3. Isaakidis P, Casas EC, Das M, et al. Treatment outcomes for HIV and MDR-TB co-infected adults and children: systematic review and meta-analysis. *Int J Tuberc Lung Dis* 2015; **19**: 969-978.
4. World Health Organization. *WHO consolidated guidelines on drug-resistant tuberculosis treatment*. Geneva: World Health Organization, 2019.
5. Petersen ML, Deeks SG, Martin JN, et al. History-adjusted marginal structural models for estimating time-varying effect modification. *Am J Epidemiol* 2007; **166**: 985-993.
6. Powers S, Qian J, Jung K, et al. Some methods for heterogeneous treatment effect estimation in high dimensions. *Stat Med* 2018; **37**: 1767-1787.
7. Bahamyrou A, Schnitzer ME, Kennedy EH, et al. Doubly robust adaptive LASSO for effect modifier discovery. arXiv preprint arXiv:2011.12746. 2020.
8. Hernán MA, Robins JM. Estimating causal effects from epidemiological data. *J Epidemiol* 2006; **60**: 578-586.
9. Zhao Q, Small DS, Ertefaie A. Selective inference for effect modification via the lasso. arXiv preprint arXiv:1705.08020. 2017.
10. Kennedy EH. Optimal doubly robust estimation of heterogeneous causal effects. arXiv preprint arXiv:2004.14497. 2020; Apr 29.

11. Ahuja SD, Ashkin D, Avendano M, et al. Multidrug resistant pulmonary tuberculosis treatment regimens and patient outcomes: an individual patient data meta-analysis of 9,153 patients. *PLoS Med* 2012; **9**(8): e1001300.
12. Mills EJ, Ioannidis JP, Thorlund K, et al. How to use an article reporting a multiple treatment comparison meta-analysis. *JAMA*. 2012; **308**:1246-1253.
13. van der Laan MJ, Rose S. *Targeted learning: causal inference for observational and experimental data*. New York, Springer Science & Business Media; 2011.
14. van der Laan MJ and Luedtke AR. Targeted learning of an optimal dynamic treatment, and statistical inference for its mean outcome. *UC Berkeley Division of Biostatistics Working Paper Series*. Working Paper 2014; 317.
15. Johnston JC, Shahidi NC, Sadatsafavi M, et al. Treatment outcomes of multidrug-resistant tuberculosis: a systematic review and meta-analysis. *PloS One* 2009; **4**(9): e6914.
16. Orenstein EW, Basu S, Shah NS, et al. Treatment outcomes among patients with multidrug-resistant tuberculosis: systematic review and meta-analysis. *Lancet Infect Dis* 2009; **9**: 153-161.
17. Akçakır Y. *Correlates of treatment outcomes of multidrug-resistant tuberculosis (MDR-TB): a systematic review and meta-analysis*. MSc Thesis. McGill University, CA. 2010.
18. Falzon D, Jaramillo E, Schünemann HJ, et al. WHO guidelines for the programmatic management of drug-resistant tuberculosis: 2011.
19. Wang G, Schnitzer ME, Menzies D, et al. Estimating treatment importance in multidrug-resistant tuberculosis using Targeted Learning: An observational individual patient data network meta-analysis. *Biometrics* 2020; **76**: 1007-1016.
20. Schnitzer ME, Steele RJ, Bally M, et al. A causal inference approach to network meta-analysis. *J Causal Inference*. 2016; **4**: 20160014.
21. Tsiatis, A. *Semiparametric theory and missing data*. New York, Springer Science & Business Media, 2007.
22. Rosenblum M and van der Laan MJ. Targeted maximum likelihood estimation of the parameter of a marginal structural model. *Int J Biostat* 2010; **6**(2): Article 19.
23. van der Laan MJ and Rubin D. Targeted maximum likelihood learning. *Int J Biostat* 2006; **2**(1): Article 11.
24. Scharfstein DO, Rotnitzky A, Robins JM. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *J Am Stat Assoc* 1999; **94**:1096-1120.

-
25. Porter KE, Gruber S, van der Laan MJ, et al. The relative performance of targeted maximum likelihood estimators. *Int J Biostat* 2011; **7**: 1-34.
 26. Schnitzer ME, van der Laan MJ, Moodie EE, et al. Effect of breastfeeding on gastrointestinal infection in infants: a targeted maximum likelihood approach for clustered longitudinal data. *Ann Appl Stat* 2014; **8**: 703-725.
 27. Gruber S and van der Laan MJ. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *Int J Biostat* 2010; **6**(1): Article 26.
 28. Rubin, DB. *Multiple imputation for nonresponse in surveys*. New York: Wiley, 1987.
 29. Buuren SV, Groothuis-Oudshoorn K. mice: Multivariate imputation by chained equations in R. *J. Stat. Softw* 2011; **45**: 1-67.
 30. Rubin DB. *Multiple imputation for nonresponse in surveys*. Hoboken, NJ, John Wiley & Sons, 2004.
 31. van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Stat Appl Genet Mol Biol* 2007; **6**: Article 25.
 32. Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw* 2010; **33**: 1.
 33. Tibshirani R. Regression shrinkage and selection via the lasso: a retrospective. *J R Stat Soc Series B Stat Methodol* 2011; **73**: 273-282.
 34. Venables WN, Ripley BD. *Modern applied statistics with S-PLUS*. New York, Springer Science & Business Media; 2013.
 35. Benjamini Y and Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B Stat Methodol* 1995; **57**: 289-300.
 36. Mitnick C, Bayona J, Palacios E, et al. Community-based therapy for multidrug-resistant tuberculosis in Lima, Peru. *N Engl J Med* 2003; **348**:119-128.
 37. Mitnick CD, Shin SS, Seung KJ, et al. Comprehensive treatment of extensively drug-resistant tuberculosis. *N Engl J Med* 2008 **359**:563-574.
 38. Vo TT, Porcher R, Chaimani A, et al. A novel approach for identifying and addressing case-mix heterogeneity in individual participant data meta-analysis. *Res Synth Methods* 2019; **10**: 582-596.
 39. Schnitzer ME, Moodie EE, Platt RW. Targeted maximum likelihood estimation for marginal time-dependent treatment effects under density misspecification. *Biostatistics* 2013; **14**: 1-14.
 40. Ahmad N, Ahuja SD, Akkerman OW, et al. Treatment correlates of successful outcomes in pulmonary multidrug-resistant tuberculosis: an individual patient data

-
- meta-analysis. *The Lancet* 2018; **392**: 821-834.
41. Pym AS, Diacon AH, Tang SJ, et al. Bedaquiline in the treatment of multidrug-and extensively drug-resistant tuberculosis. *Eur Respir J* 2016; **47**: 564-574.
 42. Sotgiu G, Centis R, D'Ambrosio L, et al. Efficacy, safety and tolerability of linezolid containing regimens in treating MDR-TB and XDR-TB: systematic review and meta-analysis. *Eur Respir J* 2012; **40**: 1430-1442.
 43. Siddique AA, Schnitzer ME, Bahamyrou A, et al. Causal inference with multiple concurrent medications: A comparison of methods and an application in multidrug-resistant tuberculosis. *Stat Methods Med Res* 2019; **28**: 3534-3549.