

Research Article

Kellie N. Ottoboni and Jason V. Poulos*

Estimating population average treatment effects from experiments with noncompliance

<https://doi.org/10.1515/jci-2018-0035>

Received Mar 13, 2018; accepted Aug 03, 2020

Abstract: Randomized control trials (RCTs) are the gold standard for estimating causal effects, but often use samples that are non-representative of the actual population of interest. We propose a reweighting method for estimating population average treatment effects in settings with noncompliance. Simulations show the proposed compliance-adjusted population estimator outperforms its unadjusted counterpart when compliance is relatively low and can be predicted by observed covariates. We apply the method to evaluate the effect of Medicaid coverage on health care use for a target population of adults who may benefit from expansions to the Medicaid program. We draw RCT data from the Oregon Health Insurance Experiment, where less than one-third of those randomly selected to receive Medicaid benefits actually enrolled.

Keywords: Causal inference, external validity, health insurance, observational studies, noncompliance, randomized controlled trials

2020 Mathematics Subject Classification: 62D20; 62P20; 62P25

1 Introduction

Randomized control trials (RCTs) are the gold standard for estimating the causal effect of a treatment. An RCT may give unbiased estimates of sample average treatment effects, but external validity is an issue when RCT participants are unrepresentative of the actual population of interest. For example, participants in an RCT in which individuals volunteer to sign up for health insurance may be in poorer health at baseline than the overall population. External validity is particularly relevant to policymakers who require information on how the treatment effect would generalize to the broader population.

A new research frontier in causal inference focuses on developing methods for extrapolating RCT results to a population [1–3]. Existing approaches to this problem are based in settings where there is full compliance with treatment; however, noncompliance is a prevalent issue in RCTs. Noncompliance occurs when individuals who are assigned to the treatment group do not accept the treatment. For individuals assigned to control, we are unable to observe who would have complied had they been assigned treatment. Noncompliance biases the intention-to-treat (ITT) estimate of the effect of treatment assignment toward zero.

We propose a reweighting method for estimating complier-average causal effects for the target population from RCT data with noncompliance, and refer to this estimator as the Population Average Treatment Effect on Treated Compliers (PATT-C). We model compliance in the RCT in order to predict the likely compliers in the RCT control group. Assuming that the response surface is the same for compliers in the RCT and population members who received treatment, we then predict the response surface for all RCT compliers and use the predicted values from the response surface model to estimate the potential outcomes of population members who received treatment, given their covariates.

***Corresponding Author: Jason V. Poulos:** Department of Statistical Science, Duke University and SAMSI, Durham, North Carolina, 27708, United States of America; Email: jason.poulos@duke.edu

Kellie N. Ottoboni: Department of Statistics, University of California, Berkeley, California, 94720, United States of America

Our approach for estimating PATT-C differs from previous reweighting methods because we only need to estimate the potential outcomes for RCT compliers and we cannot observe who in the control group would have complied had they been assigned treatment. In Stuart et al. [2], for example, a propensity score model is used to predict participation in the RCT, given pretreatment covariates common to both the RCT and population data. Individuals in the RCT and population are then weighted according to the inverse of the estimated propensity score. Hartman et al. [3] propose a method of reweighting the responses of individuals in an RCT according to the covariate distribution of the population. Reweighting methods typically leverage exchangeability of potential outcomes between the covariate-adjusted treated and controls in the RCT. In our approach, the potential outcomes between the complier treated and complier controls are not exchangeable by design, since we need to assume we know the compliance model.

When estimating the average causal effect for compliers from an RCT, researchers typically scale the estimated ITT effect by the compliance rate, assuming that there is only single crossover from treatment to control.¹ When extrapolating RCT results to a population, one might simply reweight the ITT effect according to the covariate distribution of the population and then divide by the proportion of treated compliers in the population in order to yield a population average effect of treatment on treated compliers. However, we do not observe the population compliance rate, and it is likely to differ across subgroups based on pretreatment covariates. By explicitly modeling compliance, our approach allows researchers to decompose population estimates by covariate group, which is useful for policymakers in evaluating the efficacy of policy interventions for subgroups of interest in a population.

We apply the PATT-C estimator to measure the effect of Medicaid coverage on health care use for a target population of adults who may benefit from government-backed expansions to the Medicaid program. We are particularly interested in measuring the effect of Medicaid on emergency room (ER) use because it is the main delivery system through which the uninsured receive health care, and the uninsured could potentially receive higher quality health care through primary care visits. An important policy question is whether Medicaid expansions will decrease ER utilization and increase primary care visits by the previously uninsured. We draw RCT data from a large-scale health insurance experiment where less than one-third of those randomly selected to receive Medicaid benefits actually enrolled.

The paper proceeds as follows: Section 2 presents the proposed estimator and the necessary assumptions for its identifiability; Section 3 describes the estimation procedure; Section 4 reports the estimator's performance in simulations; Section 5 uses the estimator to identify the effect of extending Medicaid coverage to the low-income adult population in the U.S.; Section 6 discusses the results and offers direction for future research.

2 Estimator

We are interested in using the results of an RCT to estimate complier-average causal effects for a target population. Compliance with treatment in the population is not assigned at random. It may depend on unobserved variables, thus confounding the effect of treatment received on the outcome of interest. RCTs are needed to isolate the effect of treatment received.

Ideally, we would take the results of an RCT and reweight the sample such that the reweighted covariates match the those in the population. In practice, one rarely knows the true covariate distribution in the target population. Instead, we consider data from an observational study in which participants are representative of the target population. The proposed estimator combines RCT and observational data to overcome these issues.

¹ Alternative methods for compliance-adjusting sample estimates include estimating sharp bounds to the ITT effect in the presence of noncompliance [4, 5]; adjustment for treatment noncompliance using principal stratification [6, 7]; and maximum-likelihood and Bayesian inferential methods [8].

2.1 Assumptions

Let Y_{isd} be the potential outcome for individual i in group s and treatment received d . Let $S_i \in \{0, 1\}$ denote the sample assignment, where $s = 0$ is the population and $s = 1$ is the RCT. Let $T_i \in \{0, 1\}$ indicate treatment assignment and $D_i \in \{0, 1\}$ indicate whether treatment was actually received. Treatment is assigned at random in the RCT, so we observe both D_i and T_i when $S_i = 1$. For compliers in the RCT, $D_i = T_i$.

The absence of T_i in the subscript of the potential outcomes notation implicitly assumes the exclusion restriction for noncompliers, which precludes treatment assignment from having an effect on the potential outcomes of noncompliers in the RCT [9]. Also implicit in our notation is the assumption of no interference between units, which prevents the potential outcomes for any unit from varying with treatments assigned to other units [10].

Let W_i be individual i 's observable pretreatment covariates that are related to the sample selection mechanism for membership in the RCT, treatment assignment in the population, and complier status. Let C_i be an indicator for individual i 's compliance with treatment, which is only observable for individuals in the RCT treatment group. In the population, we suppose that treatment is made available to individuals based on their covariates W_i .² Individuals with $T_i = 0$ do not receive treatment, while those with $T_i = 1$ may decide whether or not to accept treatment. For individuals in the population, we only observe D_i — not T_i .

Assumption 1. *Consistency under parallel studies:*

$$Y_{i0d} = Y_{i1d} \quad \forall i, d = \{0, 1\}.$$

Assumption (1) requires that each individual i has the same response to treatment received, whether i is in the RCT or not. Compliance status C_i is not a factor in this assumption because we assume that compliance is conditionally independent of sample and treatment assignment for all individuals with covariates W_i .

Assumption 2. *Conditional independence of compliance and sample and treatment assignment:*

$$C_i \perp\!\!\!\perp S_i, T_i \mid W_i, \quad 0 < \mathbb{P}(C_i = 1 \mid W_i) < 1.$$

Assumption (2) implies that $\mathbb{P}(C_i = 1 \mid S_i = 1, T_i = 1, W_i) = \mathbb{P}(C_i = 1 \mid S_i = 1, T_i = 0, W_i)$, which is useful when predicting the probability of compliance as a function of covariates W_i . Together, Assumptions (1) and (2) ensure that potential outcomes do not differ based on sample assignment or receipt of treatment.

Assumption (3) ensures the potential outcomes for treatment are independent of sample assignment for individuals with the same covariates W_i and treatment assignment.

Assumption 3. *Strong ignorability of sample assignment for treated:*

$$(Y_{i01}, Y_{i11}) \perp\!\!\!\perp S_i \mid (W_i, T_i = 1, C_i = 1), \quad 0 < \mathbb{P}(S_i = 1 \mid W_i, T_i = 1, C_i = 1) < 1.$$

We make a similar assumption for the potential outcomes under control.

Assumption 4. *Strong ignorability of sample assignment for controls:*

$$(Y_{i00}, Y_{i10}) \perp\!\!\!\perp S_i \mid (W_i, T_i = 1, C_i = 1), \quad 0 < \mathbb{P}(S_i = 1 \mid W_i, T_i = 1, C_i = 1) < 1.$$

Note that Assumptions (3) and (4) imply strong ignorability of sample assignment for treated and control non-compliers since Assumption (2) states that compliance is also independent of sample and treatment assignment, conditional on W_i . However, we are interested only on modeling the response surfaces for compliers.

² We use the same W_i across all identifying assumptions, which implicitly assumes that the observable covariates that determine sample selection also determine population treatment assignment and complier status.

Restrictive exclusion criteria in RCTs can result in a sample covariate distribution that differs substantially from the population covariate distribution, thereby reducing the external validity of RCTs [11]. High rates of exclusion pose a threat to Assumptions (3) and (4) if exclusion increases the likelihood that there are unobserved differences between the RCT and target population that are correlated with potential outcomes. For example, the RCT in our application required enrolled participants to recertify their eligibility status every six months during the study period. The exclusion of participants who failed to recertify because their household income exceeded a given cutoff threatens strong ignorability if the factors that contributed to the failure to recertify are correlated with unobservables that are also correlated with potential outcomes. While we cannot directly test the strong ignorability assumptions, bias arising from violations of these assumptions would cause the placebo tests described in Section 5.4 to fail.

We include an additional assumption to identify the average causal effect for compliers, as in Angrist et al. [12].

Assumption 5. *One-sided noncompliance:*

$$\mathbb{P}(D_i | T_i = 0) = 0, \quad \forall i.$$

Assumption (5) ensures that noncompliance is one-sided; i.e., individuals assigned to control are not allowed to receive treatment. It explicitly rules out the existence of individuals who always receive treatment (i.e., never-takers) and those who receive the opposite of their treatment assignment (i.e., defiers).

2.2 Causal diagram

Figure 1 shows the assumptions needed to identify PATT-C in the form of a causal diagram [13]. The missing arrow between T_i (or S_i) and C_i signifies that treatment (or sample) assignment may only depend on compliance status through covariates W_i , by Assumption (2). Likewise, the missing arrow between Y_{isd} and T_i (or S_i) signifies that potential outcomes may only depend on treatment (or sample) assignment through covariates, by Assumptions (3) and (4).

Confounding arcs represent potential back-door paths that contain only unobserved variables. The existence of back-door paths from D_i to Y_{isd} through S_i and C_i imply that the average causal effect of D_i on Y_{isd} cannot be identified by just conditioning on W_i ; that is, we cannot ignore the role of S_i and C_i . From the internal validity standpoint, the role of W_i is critical: in the presence of unobserved confounders, there is a back-door pathway from T_i back to W_i and into Y_{isd} .

2.3 PATT-C

PATT-C is the average causal effect of taking up treatment estimated on individuals who received treatment in the population. This interpretation follows directly from Assumptions (1) and (2), which ensure that potential outcomes do not differ based on sample assignment or receipt of treatment. It is written as follows:

$$\tau_{\text{PATT-C}} = \mathbb{E}(Y_{i01} - Y_{i00} | S_i = 0, D_i = 1). \quad (1)$$

Theorem (1) relates the treatment effect in the RCT to the treatment effect in the population (proof given in Appendix A).

Theorem 1. *Under Assumptions (1) – (5),*

$$\tau_{\text{PATT-C}} = \mathbb{E}_{01} [\mathbb{E}(Y_{i11} | S_i = 1, D_i = 1, W_i)] - \mathbb{E}_{01} [\mathbb{E}(Y_{i10} | S_i = 1, T_i = 0, C_i = 1, W_i)], \quad (2)$$

where $\mathbb{E}_{01} [\mathbb{E}(\cdot | \dots, W_i)]$ denotes the expectation with respect to the distribution of W_i for population members who received treatment.

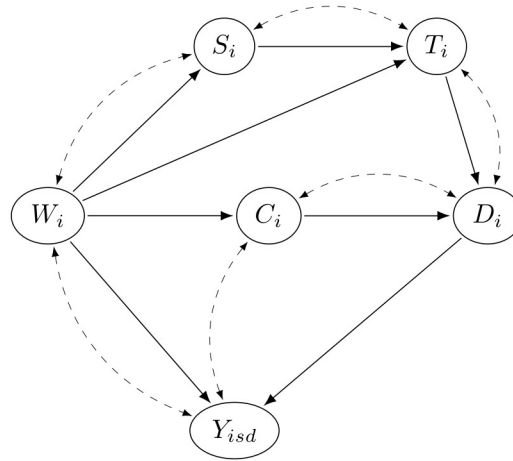


Figure 1: Causal diagram representing the effect of treatment received, D_i , on potential outcomes, Y_{isd} . Solid unidirectional arrows represent causal links between observed variables. Dashed bidirectional arrows represent confounding arcs containing only latent variables, which are implicit in the confounding arcs.

3 Estimation procedure

There are two challenges in turning Theorem (1) into an estimator of $\tau_{\text{PATT-C}}$ in practice. First, we must estimate the inner expectation over potential outcomes of compliers in the RCT. In the empirical example, we use an ensemble of algorithms to estimate the response surface for compliers in the RCT, given their covariates. Thus, the first term in the expression for $\tau_{\text{PATT-C}}$ is estimated by the weighted average of points on the response surface, evaluated for each treated population member's potential outcome under treatment. The second term is estimated by the weighted average of points on the response surface, evaluated for each treated population member's potential outcome under control.

The second challenge is that we cannot observe which individuals are included in the estimation of the second term. In the RCT control group, C_i is unobservable, as they always receive no treatment ($D_i = 0$). We must estimate the second term of Eq. (2) by predicting who in the control group would be a complier had they been assigned to treatment. Explicitly modeling compliance allows us to decompose PATT-C estimates by subgroup according to covariates common to both RCT and observational datasets. This approach also accounts for settings where the compliance rate differs between the sample and population, as well as across subgroups.

The procedure for estimating $\tau_{\text{PATT-C}}$ using Theorem (1) is as follows:

- S.1** Using the group assigned to treatment in the RCT ($S_i = 1, T_i = 1$), train a model (or an ensemble of models) to predict the probability of compliance as a function of covariates W_i .
- S.2** Using the model from S.1, predict who in the RCT assigned to control *would have* complied to treatment had they been assigned to the treatment group.³
- S.3** For the observed compliers assigned to treatment and predicted compliers assigned to control, train a model to predict the response using W_i and D_i , which gives $\mathbb{E}(Y_{i1d} \mid S_i = 1, D_i = d, W_i)$ for $d \in \{0, 1\}$.

³ We dichotomize the predicted values based on an “optimal” cut-point that minimizes the Euclidean distance between the ROC curve and the (0,1) corner of the ROC plane. This method, which is commonly used in the epidemiology literature [14], minimizes misclassification error by finding the cut-point closest to where the true positive rate (i.e., the proportion of sample compliers correctly identified) is 1 and the inverse of the true negative rate (i.e., the share of noncompliers correctly identified) is 0. We locate a prediction threshold of 45% that balances the true positive and true negative rates at about 70% each.

S.4 For all individuals who received treatment in the population ($S_i = 0, D_i = 1$), estimate their potential outcomes using the model from S.3, which gives Y_{i1d} for $d \in \{0, 1\}$. The mean counterfactual Y_{i11} minus the mean counterfactual Y_{i10} is the estimate of $\tau_{\text{PATT-C}}$.

Observe that Assumptions (3) and (4) are particularly important for estimating $\tau_{\text{PATT-C}}$: the success of the proposed estimator hinges on the assumption that the response surface is the same for compliers in the RCT and target population. If this does not hold, then the potential outcomes Y_{i10} and Y_{i11} for target population individuals cannot be estimated using the model from S.3. Section 5.3 discusses whether the strong ignorability assumptions are plausible in the empirical application.

3.1 Modeling assumptions

In addition to the identification assumptions, we require additional modeling assumptions for the estimation procedure. As pointed out in Section 2.1, we require that W_i is complete because if any relevant elements of W_i are not controlled, then there is a backdoor pathway from T_i back to W_i and into Y_{isd} . Additionally, we assume that the compliance model is accurate in predicting compliance in the training sample of RCT participants assigned to treatment and also generalizable to RCT participants assigned to control (S.1 and S.2). We describe below the method of evaluating the generalizability of the compliance model.

3.2 Ensemble method

In the empirical application, we use the super learner weighted ensemble method [15, 16] for the estimation steps S.1 and S.3. The super learner combines algorithms with a convex combination of weights based on minimizing cross-validated error, and typically outperforms single algorithms selected by cross-validation.

We choose a variety of candidate algorithms to construct the ensemble, with a preference for algorithms that tend to perform well in supervised classification tasks. We also have a preference for algorithms that have a built-in variable selection property. The idea is that we input the same W_i and each candidate algorithm selects the most important covariates for predicting compliance status or potential outcomes.⁴ We select three types of candidate algorithms: additive regression models [17]; gradient boosted regression [18]; L1 or L2-regularized linear models (i.e., lasso or ridge regression, respectively) [19]; and ensembles of decision trees (i.e., random forests) [20]. Lasso is particularly attractive because it tends to shrink all but one of the coefficients of correlated covariates to zero.

4 Simulations

We conduct a simulation study comparing the performance of the PATT-C estimator against its unadjusted analogue, which we refer to as the Population Average Treatment Effect on the Treated (PATT):

$$\tau_{\text{PATT}} = \mathbb{E}_{01} [\mathbb{E}(Y_{i11} | S_i = 1, T_i = 1, W_i)] - \mathbb{E}_{01} [\mathbb{E}(Y_{i10} | S_i = 1, T_i = 1, W_i)]. \quad (3)$$

Eq. (3) identifies the population-average causal effect of treatment assignment, adjusted according to the covariate distribution of population members who received treatment. PATT is estimated following the same estimation procedure as PATT-C, except that in estimation step S.3 the response curve is estimated on all RCT participants, regardless of compliance status, conditional on their covariates and actual treatment received. Identical to S.4 in the estimation procedure for PATT-C, we then use the response model to estimate the outcomes of population members who received treatment given their covariates. Like the PATT-C estimator, the

⁴ A potential concern when predicting potential outcomes is that the algorithm might shrink the treatment received predictor to zero, which would result in no difference between counterfactual potential outcomes.

PATT estimator crucially relies on the assumption that the response surface is the same for RCT participants and population members who received treatment.

We compare the population estimators against the sample Complier Average Causal Effect (CACE) [21], which is commonly referred to the Local Average Treatment Effect (LATE) in the econometrics literature [12, 22]. In the context of program evaluation, it is more relevant than the ITT estimator because only RCT participants who received treatment would have their outcomes affected by treatment in the presence of a nonnegative treatment effect.

CACE is defined as the average causal effect of treatment received restricted to sample compliers:

$$\tau_{\text{CACE}} = \mathbb{E}(Y_{i11} - Y_{i10} \mid S_i = 1, C_i = 1). \quad (4)$$

In other words, CACE is the treatment effect for RCT participants who would comply regardless of treatment assignment. However, we are unable to observe the compliance status of RCT participants assigned to control because we do not know if they would have complied if they had been assigned to treatment. A generalization of the instrumental variables estimator of the CACE in the presence of noncompliers is given by:

$$\hat{\tau}_{\text{CACE}} = \frac{\mathbb{E}(Y_{i11} - Y_{i10} \mid S_i = 1)}{\mathbb{P}(T_i = D_i = 1 \mid S_i = 1)}, \quad (5)$$

which is identified under Assumption (5). The estimator is equivalent to scaling the ITT effect by the sample proportion of treated compliers [23].

4.1 Simulation design

The simulation is designed so that the effect of treatment is heterogeneous and depends on covariates which are different in the RCT and target population. The design satisfies the conditional independence assumptions in Figure 1.

In the simulation, RCT eligibility, complier status, and treatment assignment in the population depend on multivariate normal covariates $(W_i^1, W_i^2, W_i^3, W_i^4)$ with means $(0.5, 1, -1, -1)$ and covariances $\text{Cov}(W_i^1, W_i^2) = \text{Cov}(W_i^1, W_i^4) = \text{Cov}(W_i^2, W_i^4) = \text{Cov}(W_i^3, W_i^4) = 1$ and $\text{Cov}(W_i^1, W_i^3) = \text{Cov}(W_i^2, W_i^3) = 0.5$. The first three covariates are observed by the researcher and W_i^4 is unobserved. $U_i, V_i, R_i,$ and Q_i are standard normal error terms. $U_i, V_i, R_i, Q_i,$ and $(W_i^1, W_i^2, W_i^3, W_i^4)$ are mutually independent.

The equation for selection into the RCT is

$$S_i = \mathbb{I}(e_2 + g_1 W_i^1 + g_2 W_i^2 + g_3 W_i^3 + e_4 W_i^4 + R_i > 0).$$

The parameter e_2 varies the fraction of the population eligible for the RCT and e_4 varies the degree of confounding with sample selection. We set the constants $g_1, g_2,$ and g_3 to be 0.5, 0.25, and 0.75, respectively.

Complier status is determined by

$$C_i = \mathbb{I}(e_3 + h_2 W_i^2 + h_3 W_i^3 + e_5 W_i^4 + Q_i > 0),$$

where e_3 varies the fraction of compliers in the population, and e_5 varies the degree of confounding with treatment assignment. We set the constants h_2 and h_3 to 0.5.

For individuals in the population ($S_i = 0$), treatment is assigned by

$$T_i = \mathbb{I}(e_1 + f_1 W_i^1 + f_2 W_i^2 + e_6 W_i^4 + V_i > 0),$$

where e_1 varies the fraction eligible for treatment in the population and e_6 varies the degree of confounding with sample selection. We set the constants f_1 and f_2 to 0.25 and 0.75, respectively. For individuals in the RCT ($S_i = 1$), treatment assignment T_i is a sample from a Bernoulli distribution with probability $p = 0.5$.

Finally, the response is determined by

$$Y_{isd} = a + bD_i + c_1 W_i^1 + c_2 W_i^2 + c_3 W_i^4 + dU_i,$$

where we set a , c_1 , c_3 , and d to 1 and c_2 to 2. The treatment effect b is heterogeneous:

$$b = \begin{cases} 1, & \text{if } W_i^1 > 0.75 \\ -1, & \text{if } W_i^1 \leq 0.75 \end{cases}$$

We generate a population of 30,000 individuals and randomly sample 5,000. Those among the 5,000 who are eligible for the RCT ($S_i = 1$) are selected. Similarly, we sample 5,000 individuals from the population and select those who are not eligible for the RCT ($S_i = 0$) to be our observational study participants.⁵ We set each individual's treatment received D_i according to their treatment assignment and complier status and observe their responses Y_{isd} . In this design, the manner in which S_i , T_i , D_i , C_i , and Y_{isd} are simulated ensures that Assumptions (1) – (5) hold.

In the assigned-treatment RCT group ($S_i = 1$, $T_i = 1$), we train a gradient boosted regression on the covariates to predict who in the control group ($S_i = 1$, $T_i = 0$) would comply with treatment ($C_i = 1$), which is unobservable. These individuals *would have* complied had they been assigned to the treatment group. For this group of observed compliers to treatment and predicted compliers from the control group of the RCT, we estimate the response surface by training another gradient boosted regression on features (W_i^1 , W_i^2 , W_i^3) and D_i .

4.2 Simulation results

We vary each of the parameters e_1 , e_2 , e_3 , e_4 , e_5 , and e_6 along a grid of five random standard normal values in order to generate different combinations of rates of compliance, treatment eligibility, RCT eligibility in the population, and confounding. For each possible combination of the six parameters, and holding all other parameters constant, we compute over 10 simulation runs the average root mean squared error (RMSE) between the true population average treatment effect and the PATT-C, PATT, or CACE estimates. Averaging across combinations, PATT yields the highest average RMSE (1.06), followed by the CACE (0.89), and the PATT-C (0.76). Since PATT does not adjust for compliance, we would expect bias relative to the PATT-C in settings with non-compliance.

Figures 2 and 3 show the average RMSE of the estimators as a function of the population compliance rate and the share of population members eligible to participate in the RCT or the population treatment rate, respectively. The PATT estimator does not correct for bias resulting from noncompliance in the population and consequently performs poorly when the population compliance rate is relatively low (i.e., $\leq 60\%$). The PATT-C estimator corrects for noncompliance in the population and thus outperforms the PATT in low-compliance settings. The CACE corrects for noncompliance in the sample, and underperforms compared to the PATT-C due to differences between the sample and population.

Figure 4 compares the average RMSE of the estimators at varying levels of compliance in the population. The error for each of the estimators decreases as a greater share of population members comply with treatment. PATT-C outperforms both PATT and CACE in terms of minimizing RMSE when the population compliance rate is below 90%. The PATT outperforms the PATT-C only at a 90% compliance rate, and the CACE outperforms the PATT when the population compliance rate is at 60% or below.

In the Appendix, Figures B1, B2, and B3 plot the relationships between estimation error and the degrees of confounding in the mechanisms that determine compliance, treatment assignment, and sample selection, respectively. The estimation error of PATT-C is comparatively less invariant to increases in the degree of confounding in the three mechanisms compared to its unadjusted counterpart. The estimation error of CACE is generally more variable than that of the population estimators due to CACE's inability to account for differences between the sample and target population.

⁵ This set-up mimics the reality that a population census is usually impossible.

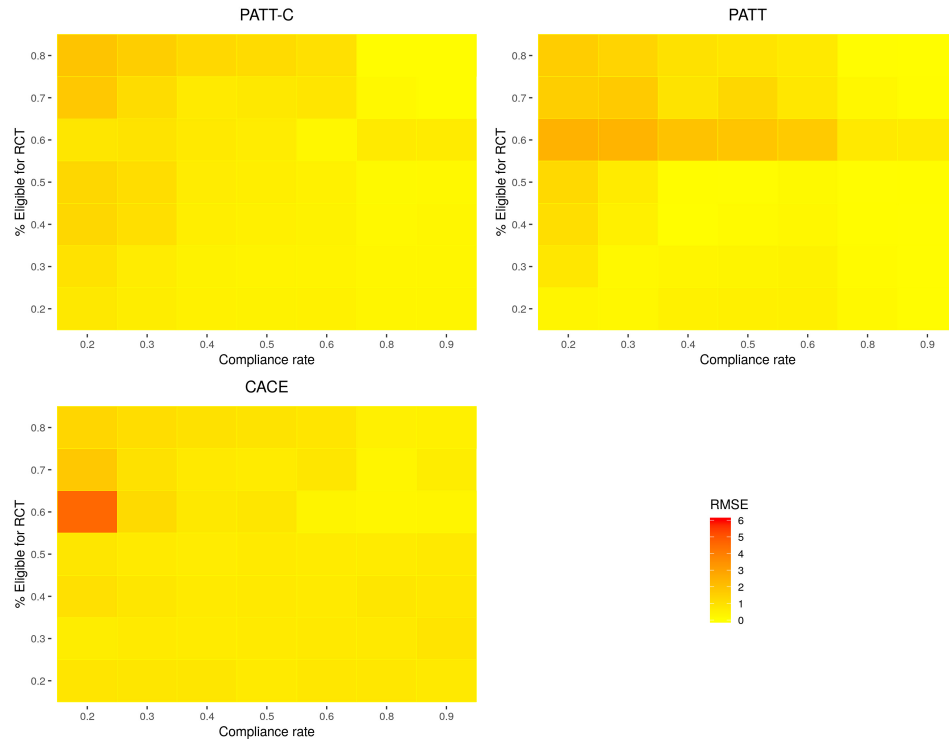


Figure 2: Average RMSE, binned by compliance rate and percent eligible for the RCT. *Notes:* Darker tiles correspond to higher errors and white tiles correspond to missing simulated data.

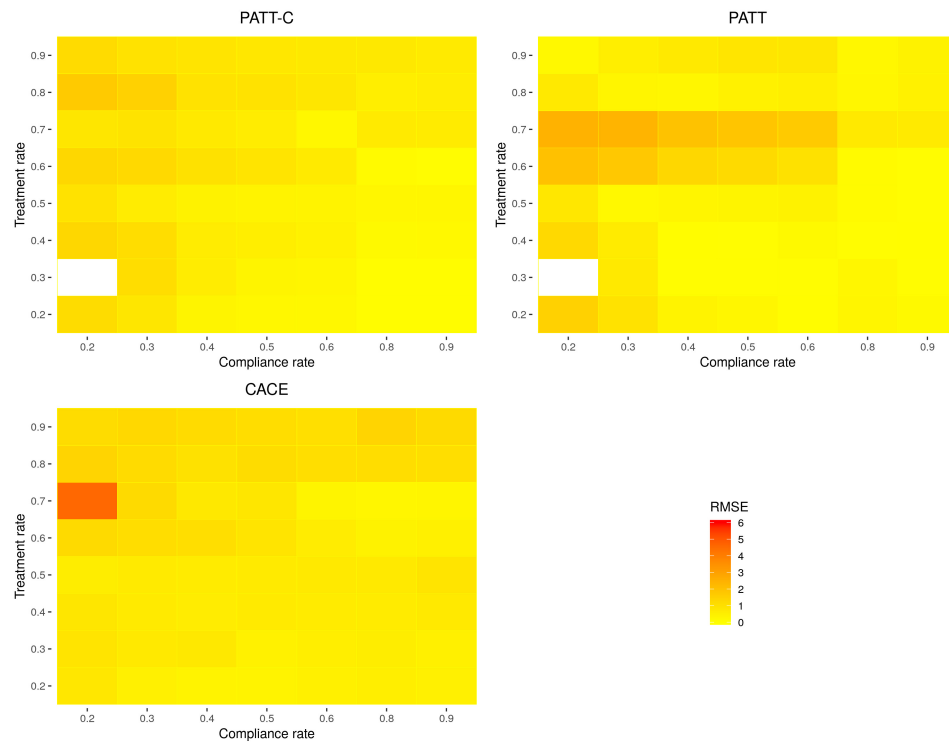


Figure 3: Average RMSE binned by compliance rate and treatment rate.

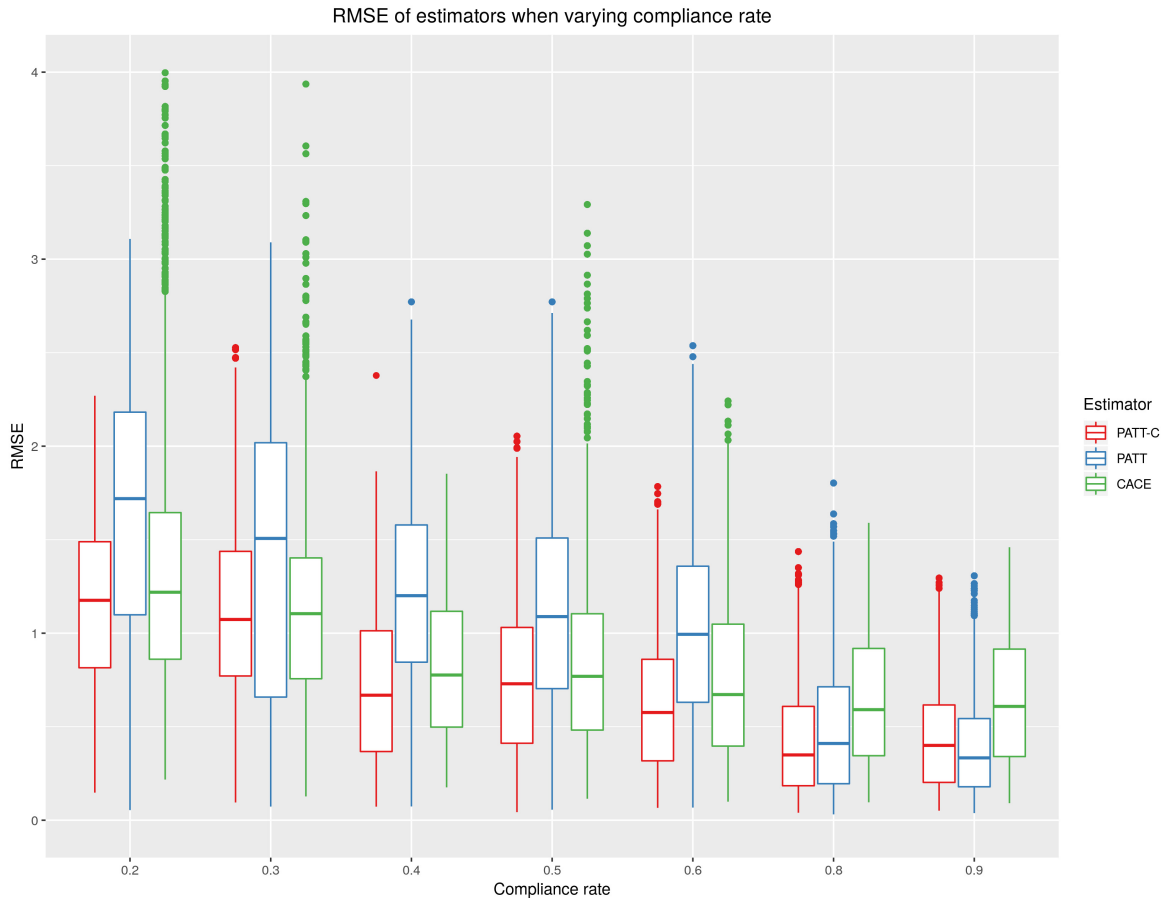


Figure 4: Average RMSE according to compliance rates in the population.

5 Application: Medicaid and health care use

We apply the PATT-C estimator to measure the effect of Medicaid coverage on health care use for a target population of adults who may benefit from expansions to the Medicaid program. In particular, we examine the population of nonelderly adults in the U.S. with household incomes at or below 138% of the Federal Poverty Level (FPL) – which amounts to \$32,913 for a four-person household in 2014 – who may be eligible for Medicaid following the Affordable Care Act (ACA) Medicaid expansion.

We draw RCT data from the Oregon Health Insurance Experiment (OHIE) [24–27], which randomly assigned Medicaid coverage to the uninsured and examined their subsequent health care use. Subsequent research calls in to question the external validity of the OHIE, which resulted in the counterintuitive finding that Medicaid increased ER use among RCT participants. Quasi-experimental studies on the impact of the 2006 Massachusetts health reform – which served as a model for the ACA – show that ER use decreased or remained constant following the reform [28, 29]. Kowalski [30] re-weights the LATE estimated on the OHIE sample to a broader population in Massachusetts using observational data from the Behavioral Risk Factor Surveillance System (BRFSS) [31] and finds Medicaid significantly decreased the probability of visiting the ER, and had no significant effect on the number of ER visits.

5.1 RCT sample

In 2008, a group of uninsured low-income adults participated in the OHIE for the chance to apply to receive health insurance through a state Medicaid program. In line with program eligibility requirements, participants were restricted to Oregon residents aged 19 to 64 who were not otherwise eligible for public insurance, who had been without insurance for six months, had income below the FPL, and held assets below \$2,000. Treatment assignment occurred at the household level: participants selected by the lottery won the opportunity for themselves and any household member to apply for Medicaid. Within a sample of 74,922 individuals representing 66,385 households, 29,834 participants were selected by the lottery; the remaining 45,008 participants served as controls in the experiment.

Participants in selected households were enrolled in Medicaid if they returned an enrollment application within 45 days of receipt. Only 30% of participants in selected households successfully enrolled. The low compliance rate is primarily due to failure to return an application or demonstrate income below the FPL. Compliance is measured using a binary variable indicating whether the participant was enrolled in any Medicaid program during the study period.

We include as covariates in our response and complier models (S.1 and S.3, respectively) pretreatment information on participant age, race, gender, education, marital status, number of children in the household, employment status, health status, and household income. We also include indicator variables on household size because lottery selection was random conditional on the number of household members. All analyses in the current application cluster-adjust standard errors at the household level because treatment occurs at the household level. The analyses also use survey weights to adjust for the probability of being sampled and non-response.

The response data originate from a mail survey containing questions about health insurance and health care use. The response variables measure health care use in terms of the number of ER and primary care (i.e., outpatient) visits in the past six months. Following Finkelstein *et al.* [24], indicator variables for survey wave and interactions with household size indicators are also included as predictors in the response and complier models because the proportion of treated participants varies across the survey waves.

5.2 Observational data

We acquire data on the target population from the National Health Interview Study (NHIS) [32] for the period 2008 to 2017.⁶ We restrict the sample to respondents with income below 138% of the FPL and who are uninsured or on Medicaid and select covariates on respondent characteristics that match the OHIE pretreatment covariates. We use a recoded variable that indicates whether respondents are on Medicaid as an analogue to the OHIE compliance measure. The outcomes of interest from the NHIS are based on questions that are virtually identical to the OHIE mail survey questions, except that the utilization questions in the NHIS are asked with a 12 month rather than a 6 month look-back period. Following Finkelstein *et al.* [24], we resolve this discrepancy by halving the NHIS responses in order to make them comparable to the OHIE outcomes.

5.3 Verifying assumptions

We verify the identification assumptions needed to identify $\tau_{\text{PATT-C}}$ prior to conducting placebo tests in Section 5.4 and reporting the results in Section 5.5.

⁶ A possible limitation of this application is that it ignores the complex sampling techniques of the NHIS sample design such as differential sampling, which is discussed in detail in Parsons *et al.* [33].

5.3.1 Consistency

Assumption (1) ensures that potential outcomes for participants in the target population would be identical to their outcomes in the RCT if they had been randomly assigned their observed treatment. In the empirical application, Medicaid coverage for uninsured individuals was applied in the same manner in the RCT as it is in the population. Differences in potential outcomes due to sample selection might arise, however, if there are differences in the mail surveys used to elicit health care use responses between the RCT and the nonrandomized study.

5.3.2 Conditional independence

Assumption (2) is violated if assignment to treatment influences the compliance status of individuals with the same covariates. The compliance ensemble can accurately classify compliance status for 78% of treated RCT participants with only the number of household members, survey wave (and the interaction between these indicators and household size indicators), and pretreatment covariates — and not treatment assignment — as predictors.⁷ This gives evidence in favor of the conditional independence assumption.

5.3.3 Strong ignorability

We cannot directly test Assumptions (3) and (4), which state that potential outcomes for treatment and control are independent of sample assignment for individuals with the same covariates and assignment to treatment. The assumptions are only met if every possible confounder associated with the response and the sample assignment is accounted for. In estimating the response surface, we use all demographic, socioeconomic, and pre-existing health condition data that were common in the OHIE and NHIS data. Potentially important unobserved confounders include the number of hospital and outpatient visits in the previous year, proximity to health services, and enrollment in other federal programs.

The final two columns of Table B2 compares RCT participants selected for Medicaid with population members on Medicaid. Compared to the RCT compliers, the population members who received treatment are younger, female, and more racially and ethnically diverse. Diagnoses of diabetes, asthma, high blood pressure, and heart disease are more common among the population on Medicaid than the RCT treated.

Strong ignorability assumptions may also be violated due to the fact that the OHIE applied a more stringent exclusion criteria compared to the NHIS sample. While the RCT and population sample both screened for individuals below the FPL, only the RCT required those enrolled to recertify their household income eligibility during the study period. Strong ignorability would not hold if the failure to recertify is correlated with unobserved variables.

5.3.4 No defiers

Angrist et al. [12] show that the bias due to violations of Assumption (5) is equivalent to the difference of average causal effects of treatment received for compliers and defiers, multiplied by the relative proportion of defiers, $\mathbb{P}(i \text{ is a defier}) / (\mathbb{P}(i \text{ is a complier}) - \mathbb{P}(i \text{ is a defier}))$.

Table B4 reports the distribution of participants in the OHIE by status of treatment assignment and treatment received. Assumption (5) does not hold due to the presence of defiers; i.e., participants who were assigned to control and enrolled in Medicaid during the study period. About 7% of the RCT sample were as-

⁷ The compliance ensemble is evaluated in terms of 10-fold cross-validated MSE. The distribution of MSE for the ensemble and its candidate algorithms are provided in Table B1.

signed to control but were enrolled in Medicaid ($T_i < D_i$) and 66% of the sample complied with treatment assignment ($D_i = T_i$), which results in a bias multiplier of 0.1. Suppose that the difference of average causal effects of Medicaid received on ER use for compliers and defiers is 1.2%. The resulting bias is only 0.1%, which would not meaningfully alter the interpretation of the PATT-C or CACE estimates.

5.3.5 Implied assumptions

We implicitly assume no interference between households in the OHIE because treatment assignment occurred at the household level. Within-household interference is not possible in this RCT because household members share the same treatment status. Interference between households would threaten the no-interference assumption in the unlikely case that the Medicaid coverage of individuals in treated households affects the health care use of individuals in households assigned to control.

The implied exclusion restriction assumption ensures treatment assignment affects the response only through enrollment in Medicaid. It is reasonable that a person's enrollment in Medicaid, not just their eligibility to enroll, would affect their hospital use. For private health insurance one might argue that eligibility may be negatively correlated with hospital use, as people with pre-existing conditions are less often eligible yet go to the hospital more frequently. This should not be the case with a federally funded program such as Medicaid.

5.4 Placebo tests

We conduct placebo tests to check whether the average outcomes differ between the RCT compliers on Medicaid and the adjusted population members who received Medicaid. If the placebo tests detect a significant difference between the mean outcomes of these groups, it would indicate bias arising from violations of the identification and modeling assumptions.

Table B3 reports the results of placebo tests for the weighted difference-in-mean outcomes of RCT compliers and adjusted population members who received treatment. The mean outcome of RCT compliers is calculated from the observed RCT sample and is weighted by OHIE survey weights. The adjusted population mean is the counterfactual Y_{i11} estimated from S.4 of the estimation procedure, and weighted by NHIS survey weights.

We first perform a Test of One-Sided Significance (TOST) [34] that evaluates equivalence between the weighted distributions, with a failure to reject the null of a substantively large difference. The TOST p -value is below the conventional level of significance ($p \leq 0.05$), indicating rejection of the null of a substantively large difference. Secondly, we conduct standard tests-of-difference and fail to reject the null of no difference. These results imply that the PATT-C estimator is not biased by differences in how Medicaid is delivered or health outcomes are measured between the RCT and population, or by differences in the unobserved characteristics of individuals in the sample or population.

5.5 Empirical results

Table 1 compares population and sample estimates of the effect of Medicaid on health care use. The PATT-C estimates indicate that Medicaid coverage has a statistically significant and negative effect on the number of ER and outpatient visits. The PATT estimates, which do not account for noncompliance in the RCT, indicate a similarly sized negative effect on the number of ER visits and no effect on the number of outpatient visits attributable to Medicaid coverage. Our population estimates are consistent with the study of Kowalski [30], which extrapolates LATE estimates to the Massachusetts population and find negative LATEs on ER utilization.

Table 1: Comparison of population and sample estimates.

Estimator	Outcome	
	# ER visits	# outpatient visits
PATT-C	-0.002 (0.0006)	-0.017 (0.0008)
PATT	-0.001 (0.0004)	0.003 (0.0027)
CACE	0.023 (0.0327)	0.225 (0.0920)

Notes: survey-weighted difference-in-means; bootstrapped standard errors (in parentheses) are clustered on the household. Response models include binary indicators for treatment received, household size, survey wave (and interactions), and pretreatment covariates.

The PATT-C estimates differ in direction and magnitude relative the CACE estimates on the OHIE sample, which can be explained by differences in the covariate distributions of the RCT sample and population. The CACE estimates indicate a positive and significant effect on primary care visits and no effect on ER use. The direction and magnitude of the CACE estimated treatment effect on ER use is almost identical to the corresponding LATE estimates on the RCT sample reported by Finkelstein et al. [24], although the authors uncover a significant (positive) effect of Medicaid on ER use.

Treatment effect heterogeneity in the population helps to explain the differences between the complier-adjusted sample and population treatment effect estimates. Figure 5 plots PATT-C estimates of heterogeneous treatment effects on ER and outpatient use in the population. Heterogeneous treatment effects are estimated by taking differences across response surfaces for a given binary covariate, and response surfaces are estimated with the ensemble mean predictions. We find significant decreases in ER use due to Medicaid coverage for adult population members who are female, black, below the age of 50, have a vocational or two-year degree, and have household income of over \$10,000. Population members with diabetes, asthma, and heart

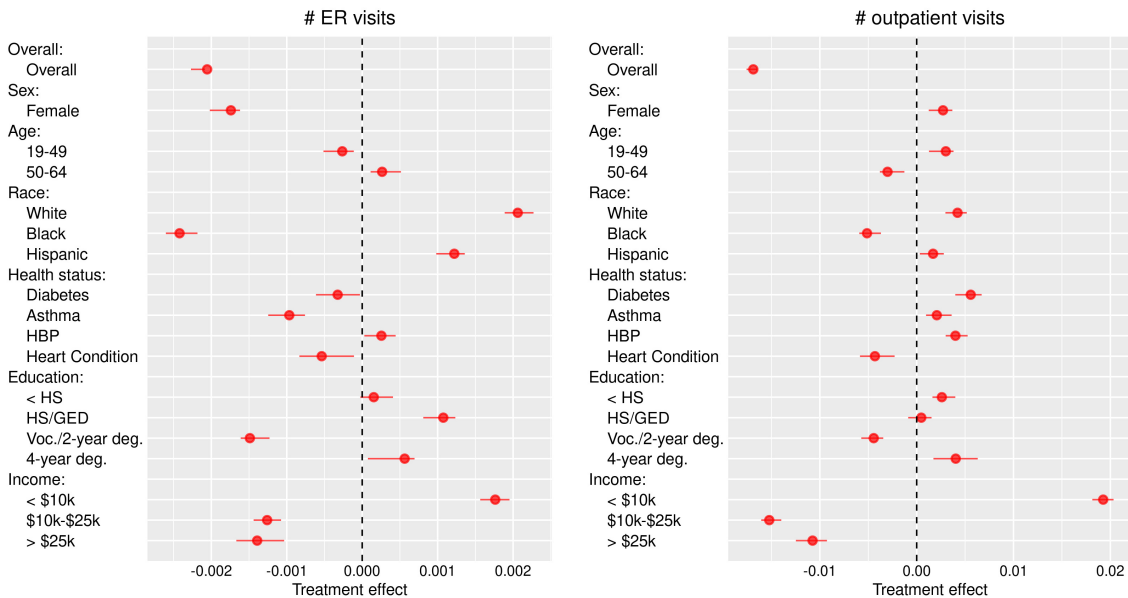


Figure 5: Heterogeneity in PATT-C treatment effect estimates. *Notes:* Horizontal lines represent 95% percentile bootstrap confidence intervals.

conditions are also less likely to visit the ER due to Medicaid coverage. The estimated effect of Medicaid on the number of outpatient visits is negative for population members who are over the age of 50, black, diagnosed with a heart condition, have a vocational or two-year degree, and have household income of over \$10,000.

6 Discussion

The simulation results presented in Section 4 show that the PATT-C estimator outperforms its unadjusted counterpart when the population compliance rate is low. Of course, the simulation results depend on the particular way we parameterized the compliance, selection, treatment assignment, and response schedules.

In particular, the strength of correlation between the covariates and compliance governs how well the estimator will perform, since S.1 of the estimation procedure predicts who *would be* a complier in the RCT control group, had they been assigned to treatment. If it is difficult to predict compliance using the observed covariates, then the estimator will perform badly because of noise introduced by incorrectly treating non-compliers as compliers. Further research should be done into ways to test how well the model of compliance works in the population or explore models to more accurately predict compliance in RCTs. Accurately predicting compliance is not only essential for yielding unbiased estimates of the average causal effects for target populations, it is also useful for researchers and policymakers to know which groups of individuals are unlikely to comply with treatment.

In the OHIE trial, less than one-third of those selected to receive Medicaid benefits actually enrolled. We accurately classified compliance status for 78% of treated RCT participants using only the pretreatment covariates as features. While we do not know how well the compliance ensemble predicts for the control group, the control group should be similar to the treatment group on pretreatment covariates because of the RCT randomization. The model's performance on the training set suggests that compliance is not purely random and depends on observed covariates, which gives evidence in favor of using PATT-C.

In the empirical application, the sample population differs in several dimensions from the target population of individuals who would be covered by other Medicaid expansions, such as the ACA expansion. For instance, the RCT participants are disproportionately white and over the age of 50. The RCT participants volunteered for the study and may be in poorer health compared to the target population. These differences in baseline covariates make reweighting or response surface methods necessary to extend the RCT results to the population.

Explicitly modeling compliance allows us to decompose population estimates by subgroup according to pretreatment covariates common to both RCT and observational datasets; e.g, demographic variables, pre-existing conditions, and insurance coverage. We find substantial treatment effect heterogeneity in terms of race, education, and health status subgroups. This pattern is expected because RCT participants volunteered for the study.

The PATT-C estimates indicate that Medicaid coverage significantly decreases the number of ER and outpatient visits. The estimated effect on ER utilization is consistent with quasi-experimental studies on the impact of the 2006 Massachusetts health reform and a previous study that extrapolates LATE estimates on the OHIE sample to the Massachusetts population. The PATT-C estimates differ in direction and magnitude relative the complier-adjusted sample estimates, which can be explained by treatment effect heterogeneity in the population and the differences in covariate distributions of the RCT sample and population.

Acknowledgement: Poulos acknowledges support from the National Science Foundation Graduate Research Fellowship [grant number DGE 1106400]. Code for reproducing the results in this paper is available in the public repository <https://github.com/jvpoulos/patt-c>.

References

- [1] K. Imai, G. King, and E. Stuart, Misunderstandings among experimentalists and observationalists about causal inference, *J. R. Stat. Soc. Ser. A. Stat. Soc.* 171 (2008), no. 2, 481–502.
- [2] E. A. Stuart, S. R. Cole, C. P. Bradshaw, and P. J. Leaf, The use of propensity scores to assess the generalizability of results from randomized trials, *J. R. Stat. Soc. Ser. A. Stat. Soc.* 174 (2011), no. 2, 369–386.
- [3] E. Hartman, R. Grieve, R. Ramsahai, and J. S. Sekhon, From SATE to PATT: Combining experimental with observational studies to estimate population treatment effects, *J. R. Stat. Soc. Ser. A. Stat. Soc.* 178 (2015), no. 3, 757–778.
- [4] A. Balke and J. Pearl, Bounds on treatment effects from studies with imperfect compliance, *J. Am. Stat. Assoc.* 92 (1997), no. 439, 1171–1176.
- [5] K. Imai, D. Tingley, and T. Yamamoto, Experimental designs for identifying causal mechanisms, *J. R. Stat. Soc. Ser. A. Stat. Soc.* 176 (2013), no. 1, 5–51.
- [6] C. E. Frangakis and D. B. Rubin, Principal stratification in causal inference, *Biometrics* 58 (2002), no. 1, 21–29.
- [7] P. Frumento, F. Mealli, B. Pacini, and D. B. Rubin, Evaluating the effect of training on wages in the presence of noncompliance, nonemployment, and missing outcome data, *J. Am. Stat. Assoc.* 107 (2012), no. 498, 450–466.
- [8] L. H. Yau and R. J. Little, Inference for the complier-average causal effect from longitudinal data subject to noncompliance and missing data, with application to a job training assessment for the unemployed, *J. Am. Stat. Assoc.* 96 (2001), no. 456, 1232–1244.
- [9] G. W. Imbens and D. B. Rubin, *Causal Inference in Statistics, Social, and Biomedical Sciences* (2015), Cambridge: Cambridge University Press.
- [10] D. B. Rubin, Comment: Neyman (1923) and causal inference in experiments and observational studies, *Stat. Sci.* 5 (1990), no. 4, 472–480.
- [11] P. M. Rothwell, External validity of randomised controlled trials: ‘To whom do the results of this trial apply?’, *The Lancet* 365 (2005), no. 9453, 82–93.
- [12] J. D. Angrist, G. W. Imbens, and D. B. Rubin, Identification of causal effects using instrumental variables, *J. Am. Stat. Assoc.* 91 (1996), no. 434, 444–455.
- [13] J. Pearl, Causal diagrams for empirical research, *Biometrika* 82 (1995), no. 4, 669–688.
- [14] N. J. Perkins and E. F. Schisterman, The inconsistency of ‘optimal’ cut-points using two ROC based criteria, *Am. J. Epidemiol.* 163 (2006), no. 7, 670–675.
- [15] M. J. Van der Laan, E. C. Polley, and A. E. Hubbard, Super learner, *Stat. Appl. Genet. Mol. Biol.* 6 (2007), no. 1, 1–21.
- [16] E. C. Polley, S. Rose, and M. J. Van der Laan, Super learning, in *Targeted learning* (2011), 43–66, New York: Springer.
- [17] A. Buja, T. Hastie, and R. Tibshirani, Linear smoothers and additive models, *Ann. Stat.* 17 (1989), no. 2, 453–510.
- [18] J. H. Friedman, Greedy function approximation: A gradient boosting machine, *Ann. Stat.* 29 (2001), no. 5, 1189–1232.
- [19] R. Tibshirani, J. Bien, J. Friedman, T. Hastie, N. Simon, J. Taylor, and R. J. Tibshirani, Strong rules for discarding predictors in lasso-type problems, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 74 (2012), no. 2, 245–266.
- [20] L. Breiman, Random forests, *Mach. Learn.* 45 (2001), no. 1, 5–32.
- [21] G. W. Imbens and D. B. Rubin, Bayesian inference for causal effects in randomized experiments with noncompliance, *Ann. Stat.* 25 (1997), no. 1, 305–327.
- [22] G. W. Imbens and J. D. Angrist, Identification and estimation of local average treatment effects, *Econometrica* 62 (1994), no. 2, 467–475.
- [23] D. A. Freedman, Statistical models for causation: What inferential leverage do they provide?, *Eval. Rev.* 30 (2006), no. 6, 691–713.
- [24] A. Finkelstein, S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, and H. Allen, The Oregon health insurance experiment: Evidence from the first year, *Q. J. Econ.* 127 (2012), no. 3, 1057–1106.
- [25] K. Baicker, S. L. Taubman, H. L. Allen, M. Bernstein, J. H. Gruber, J. P. Newhouse, E. C. Schneider, B. J. Wright, A. M. Zaslavsky, and A. N. Finkelstein, The Oregon experiment—effects of Medicaid on clinical outcomes, *N. Engl. J. Med.* 368 (2013), no. 18, 1713–1722.
- [26] K. Baicker, A. Finkelstein, J. Song, and S. Taubman, The impact of Medicaid on labor market activity and program participation: Evidence from the Oregon health insurance experiment, *Am. Econ. Rev.* 104 (2014), no. 5, p. 322–328.
- [27] S. Taubman, H. Allen, B. Wright, K. Baicker, and A. Finkelstein, Medicaid increases emergency department use: Evidence from Oregon’s health insurance experiment, *Science* 343 (2014), no. 6168, 263–268.
- [28] S. Miller, The effect of insurance on emergency room visits: An analysis of the 2006 Massachusetts health reform, *J. Public Econ.* 96 (2012), no. 11–12, 893–908.
- [29] J. T. Kolstad and A. E. Kowalski, The impact of health care reform on hospital and preventive care: Evidence from Massachusetts, *J. Public Econ.* 96 (2012), no. 11–12, 909–929.
- [30] A. E. Kowalski, “Doing more when you’re running LATE: Applying marginal treatment effect methods to examine treatment effect heterogeneity in experiments,” NBER Work. Pap. Ser. 22363 (2016).
- [31] Centers for Disease Control and Prevention. Behavioral Risk Factor Surveillance System Survey Data (2016). Atlanta, Georgia: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention.

- [32] Centers for Disease Control and Prevention. National Health Interview Survey (2019). Atlanta, Georgia: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention.
- [33] V. L. Parsons, C. L. Moriarity, K. Jonas, T. F. Moore, K. E. Davis, and L. Tompkins, Design and estimation for the National Health Interview Survey, 2006–2015 (2014). Atlanta, Georgia: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention.
- [34] R. L. Berger and J. C. Hsu, Bioequivalence trials, intersection-union tests and equivalence confidence sets, *Stat. Sci.* 11 (1996), no. 4, 283–319.

A Proof of Theorem (1)

Proof. We separate the expectation linearly into two terms and consider each individually.

$$\begin{aligned}
 \mathbb{E}(Y_{i01} | S_i = 0, D_i = 1) &= \mathbb{E}(Y_{i11} | S_i = 0, D_i = 1) && \text{by Assumption (1)} \\
 &= \mathbb{E}(Y_{i11} | S_i = 0, T_i = 1, C_i = 1) && \text{by Assumption (5)} \\
 &= \mathbb{E}_{01} [\mathbb{E}(Y_{i11} | S_i = 0, T_i = 1, C_i = 1, W_i)] \\
 &= \mathbb{E}_{01} [\mathbb{E}(Y_{i11} | S_i = 1, T_i = 1, C_i = 1, W_i)] && \text{by Assumption (3)} \\
 &= \mathbb{E}_{01} [\mathbb{E}(Y_{i11} | S_i = 1, D_i = 1, W_i)].
 \end{aligned}$$

Intuitively, conditioning on W_i makes sample selection ignorable under Assumption (3). This is the critical connector between the third and fourth lines of the first expectation derivation.

$$\begin{aligned}
 \mathbb{E}(Y_{i00} | S_i = 0, D_i = 1) &= \mathbb{E}(Y_{i10} | S_i = 0, D_i = 1) && \text{by Assumption (1)} \\
 &= \mathbb{E}(Y_{i10} | S_i = 0, T_i = 1, C_i = 1) && \text{by Assumption (5)} \\
 &= \mathbb{E}_{01} [\mathbb{E}(Y_{i10} | S_i = 1, T_i = 1, C_i = 1, W_i)] && \text{by Assumption (4)} \\
 &= \mathbb{E}_{01} [\mathbb{E}(Y_{i10} | S_i = 1, T_i = 0, C_i = 1, W_i)]. && \text{by Assumption (2)}
 \end{aligned}$$

The last line follows because Assumption (2) allows us to use RCT controls who would have complied had they been assigned to treatment. Finally, the result follows by plugging these two expressions into Eq. (1). \square

B Tables & Figures

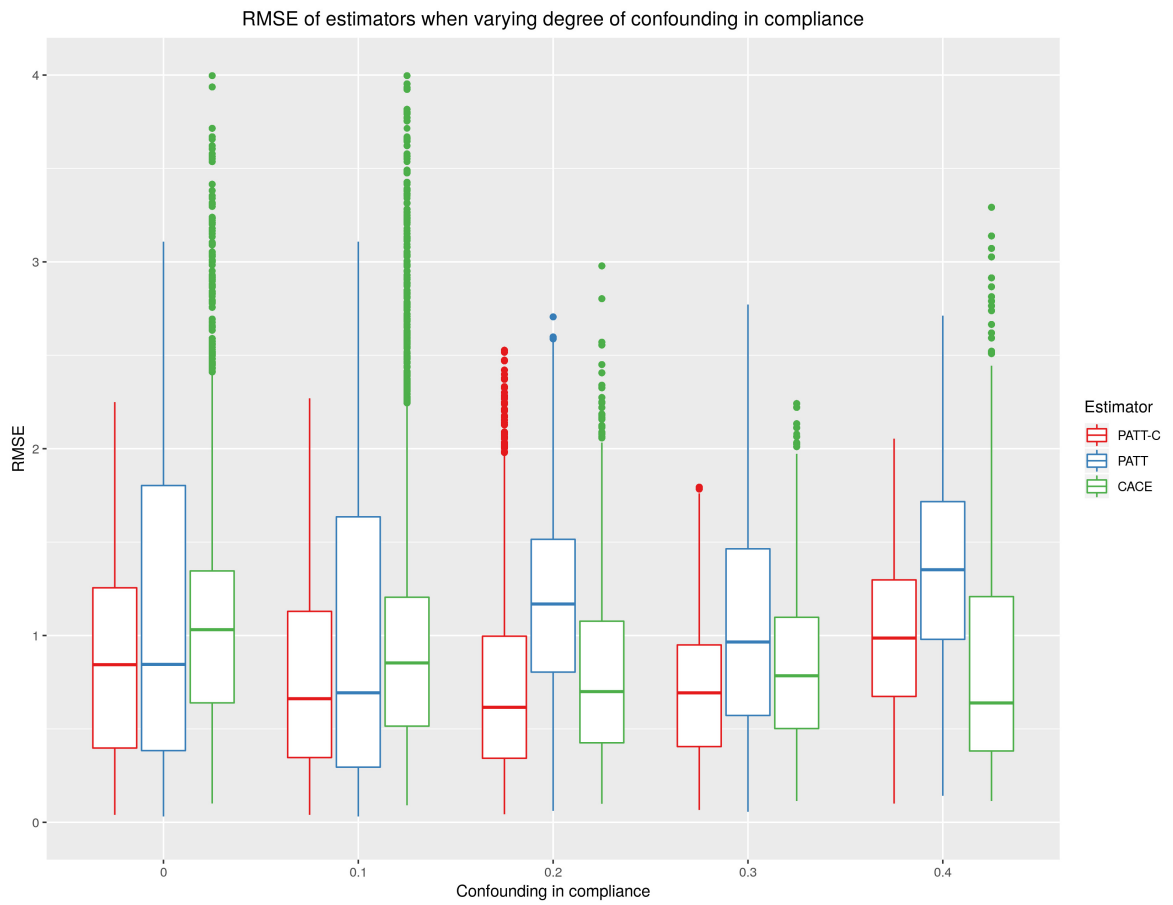


Figure B1: Average RMSE according to degree of confounding in compliance.

Table B1: Distribution of MSE for compliance ensemble.

Algorithm	Mean	SE	Min.	Max.
Super learner (<i>SuperLearner</i>)	0.22	0.001	0.21	0.23
Lasso regression (<i>glmnet</i>)	0.22	0.001	0.21	0.23
Random forests (<i>randomForest</i>)	0.27	0.002	0.25	0.29
Ridge regression (<i>glmnet</i>)	0.22	0.001	0.21	0.23

MSE is 10-fold cross-validated error for super learner ensemble and candidate algorithms. R package used for implementing each algorithm in parentheses.

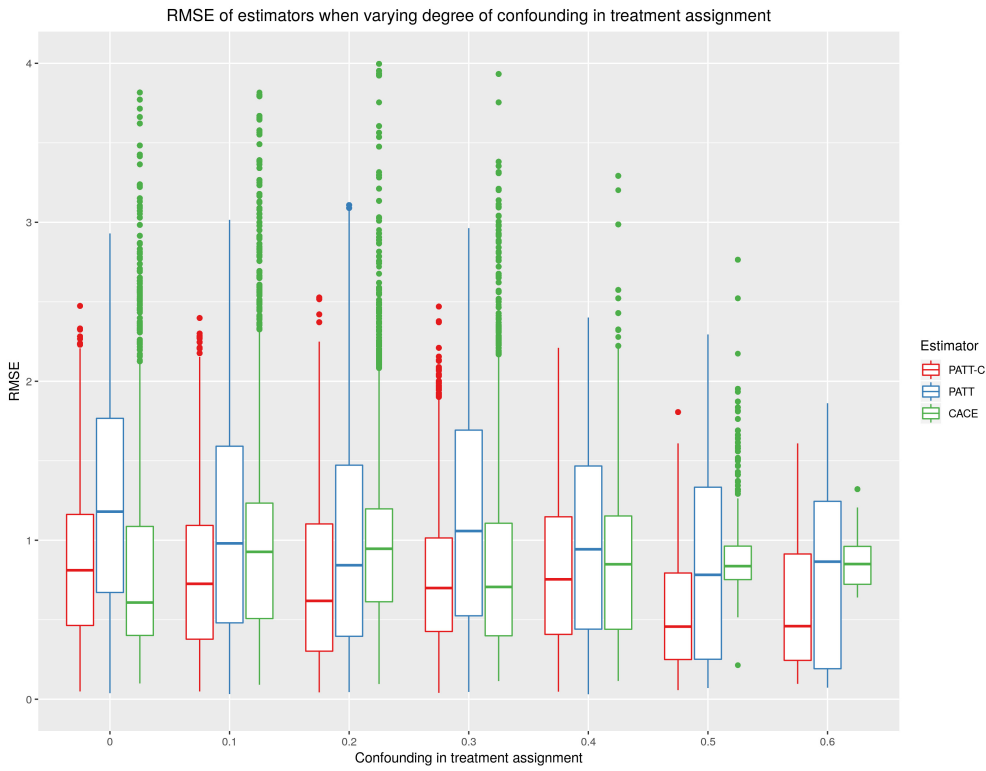


Figure B2: Average RMSE according to degree of confounding in treatment assignment.

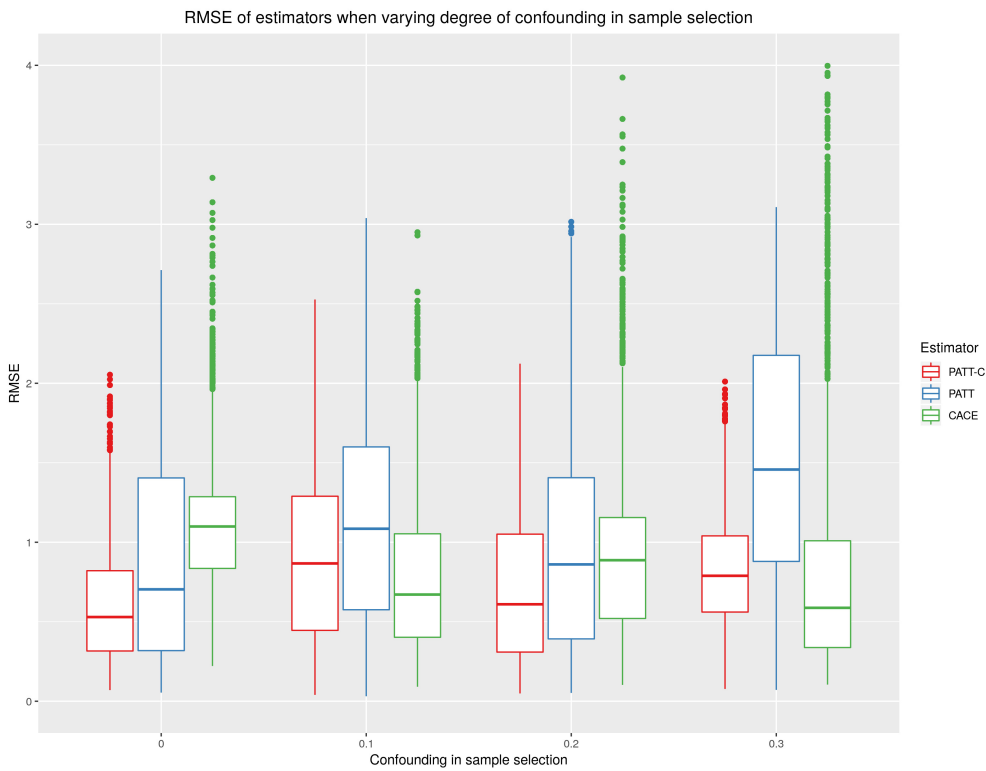


Figure B3: Average RMSE according to degree of confounding in sample selection.

Table B2: Pretreatment covariates and responses for OHIE and NHIS respondents by Medicaid coverage status.

Covariate	OHIE no Medicaid <i>n</i> = 13, 573		OHIE Medicaid <i>n</i> = 5, 547		NHIS Medicaid <i>n</i> = 6, 256	
	N	%	N	%	N	%
<i>Age:</i>						
19-49	4,166	31.0	1,473	26.9	4,322	69.2
50-64	9,407	69.0	4,074	73.1	1,934	30.8
<i>Education:</i>						
Less than high school	2,181	15.8	1,028	18.3	1,941	30.8
High school diploma or GED	6,948	50.9	2,948	52.7	2,074	33.3
Voc. training / 2-year degree	2,889	21.5	1,118	20.8	1,809	29.1
4-year college degree or more	1,555	11.8	453	8.2	432	6.8
<i>Ethnicity:</i>						
Black	193	3.1	247	3.5	1,723	27.0
Hispanic	1,242	9.3	509	9.0	1,570	24.9
White	11,453	84.2	4,675	84.7	3,899	63.0
<i>Gender:</i>						
Female	7,807	57.8	3,280	59.3	4,285	68.4
<i>Health status:</i>						
Asthma	2,114	15.6	1,004	18.3	1,269	20.1
Diabetes	1,484	10.7	595	11.0	864	14.0
Heart condition	302	2.1	178	3.1	529	8.5
High blood pressure (HBP)	3,740	27.4	1,533	28.1	2,162	34.9
<i>Income:</i>						
< \$10k	6,003	43.7	3,488	62.6	2,587	41.8
\$10k-\$25k	5,583	41.7	1,678	30.5	3,095	49.2
> \$25k	1,987	14.6	381	6.9	574	9.0
Responses						
	Mean	S.d.	Mean	S.d.	Mean	S.d.
# ER visits	0.45	1.00	0.40	0.85	0.25	0.54
# outpatient visits	1.98	2.97	1.98	2.92	1.07	1.21

Notes: percentages calculated using OHIE or NHIS survey weights.

Table B3: Placebo tests.

Outcome	RCT complier weighted mean	Adjusted population weighted mean	Weighted difference-in-means	TOST <i>p</i> -value	Two-sided <i>p</i> -value
# ER visits	0.449	0.466	-0.017	< 0.001	0.202
# outpatient visits	1.966	2.029	-0.063	< 0.001	0.090

Notes: *p*-values correspond to survey-weighted difference-in-means estimated from bootstrapped standard errors, which are clustered on the household. TOST *p*-value evaluates equivalence between the weighted distributions, and the two-sided *p*-value evaluates difference between the weighted distribution.

Table B4: Distribution of OHIE participants by status of treatment assignment (T_i) and treatment received (D_i).

	$D_i = 0$	$D_i = 1$	N
$T_i = 0$	8,343	1,265	9,608
$T_i = 1$	5,230	4,282	9,512
N	13,573	5,547	19,120

Table B5: Error and weights for candidate algorithms in response ensemble for RCT compliers.

# ER visits		
Algorithm	MSE	Weight
Additive regression, degree = 2 (gam)	0.81	0
Additive regression, degree = 3 (gam)	0.81	0
Additive regression, degree = 4 (gam)	0.81	0
Boosted regression (gbm)	0.79	0
Lasso regression (glmnet)	0.79	0.76
Random forests, #preds. = 1 (randomForest)	0.79	0.18
Random forests, #preds. = 10 (randomForest)	0.83	0.06
Regularized linear regression, $\alpha = 0.25$ (glmnet)	0.79	0
Regularized linear regression, $\alpha = 0.5$ (glmnet)	0.79	0
Regularized linear regression, $\alpha = 0.75$ (glmnet)	0.79	0
Ridge regression (glmnet)	0.79	0
# outpatient visits		
Algorithm	MSE	Weight
Additive regression, degree = 2 (gam)	6.60	0
Additive regression, degree = 3 (gam)	6.60	0
Additive regression, degree = 4 (gam)	6.60	0
Boosted regression (gbm)	6.48	0
Lasso regression (glmnet)	6.47	0.87
Random forests, #preds. = 1 (randomForest)	6.47	0
Random forests, #preds. = 10 (randomForest)	6.73	0.13
Regularized linear regression, $\alpha = 0.25$ (glmnet)	6.47	0
Regularized linear regression, $\alpha = 0.5$ (glmnet)	6.47	0
Regularized linear regression, $\alpha = 0.75$ (glmnet)	6.47	0
Ridge regression (glmnet)	6.47	0

Notes: cross-validated error and weights used for each algorithm in super learner ensemble. *MSE* is the ten-fold cross-validated mean squared error for each algorithm. *Weight* is the coefficient for super learner, which is estimated using non-negative least squares based on the Lawson-Hanson algorithm. R package used for implementing each algorithm in parentheses. #preds. is the number of predictors randomly sampled as candidates in each decision tree in random forests algorithm. α is a parameter that mixes L1 and L2 norms. degree is the smoothing term for smoothing splines.

Table B6: Error and weights for candidate algorithms in response ensemble for all RCT participants.

# ER visits		
Algorithm	MSE	Weight
Additive regression, degree = 2 (gam)	0.76	0
Additive regression, degree = 3 (gam)	0.76	0
Additive regression, degree = 4 (gam)	0.76	0
Boosted regression (gbm)	0.75	0
Lasso regression (glmnet)	0.75	0
Random forests, #preds. = 1 (randomForest)	0.75	1
Random forests, #preds. = 10 (randomForest)	0.79	0
Regularized linear regression, $\alpha = 0.25$ (glmnet)	0.75	0
Regularized linear regression, $\alpha = 0.5$ (glmnet)	0.75	0
Regularized linear regression, $\alpha = 0.75$ (glmnet)	0.75	0
Ridge regression (glmnet)	0.75	0
# outpatient visits		
Algorithm	MSE	Weight
Additive regression, degree = 2 (gam)	6.44	0
Additive regression, degree = 3 (gam)	6.44	0
Additive regression, degree = 4 (gam)	6.44	0
Boosted regression (gbm)	6.39	0
Lasso regression (glmnet)	6.39	0.87
Random forests, #preds. = 1 (randomForest)	6.39	0
Random forests, #preds. = 10 (randomForest)	6.61	0.13
Regularized linear regression, $\alpha = 0.25$ (glmnet)	6.39	0
Regularized linear regression, $\alpha = 0.5$ (glmnet)	6.39	0
Regularized linear regression, $\alpha = 0.75$ (glmnet)	6.39	0
Ridge regression (glmnet)	6.39	0

See notes to Table B5.