Political Analysis (2026), **34**, 38–52
doi:10.1017/pan.2025.10005

PA

ARTICLE

Estimating the Local Average Treatment Effect Without the Exclusion Restriction

Zachary Markovich

Scientist, Uber Technologies, Seattle, WA, USA; Email: zachmarkovich@gmail.com

(Received 1 June 2023; revised 4 January 2025; accepted 18 January 2025; published online 21 July 2025)

Abstract

Existing approaches to conducting inference about the *Local Average Treatment Effect* or LATE require assumptions that are considered tenuous in many applied settings. In particular, Instrumental Variable techniques require monotonicity and the exclusion restriction while principal score methods rest on some form of the principal ignorability assumption. This paper provides new results showing that an estimator within the class of principal score methods allows conservative inference about the LATE without invoking such assumptions. I term this estimator the *Compliance Probability Weighting* estimator and show that, under very mild assumptions, it provides an asymptotically conservative estimator for the LATE. I apply this estimator to a recent survey experiment and provide evidence of a stronger effect for the subset of compliers than the original authors had uncovered.

Keywords: experimental design; instrumental variables; causal inference

Edited by: Jeff Gill

1. Introduction

Compliance is often a concern when analyzing randomized experiments. Treatment effects cannot be estimated among subjects who did not receive their assigned treatment (non-compliers) without strong assumptions, so researchers instead typically only seek to estimate the treatment effect for the subset of subjects who always receive the treatment that they are assigned (the *Local Average Treatment Effect* (LATE)). Instrumental variables (IV) methods are the most common approach taken when estimating the LATE in political science research, but rest on assumptions (namely, the exclusion restriction and monotonicity) that will fail in many applied settings (Angrist, Imbens, and Rubin 1996).

Principal score methods represent an alternate approach to conducting inference about the LATE. While the class of estimators contained within principal score methods is quite broad, they are unified by the general approach of conditioning causal effects on some post-treatment variable using estimated probabilities that units fall within a particular strata of that variable (Ding and Lu 2017; Feller, Mealli, and Miratrix 2017; Hill, Waldfogel, and Brooks-Gunn 2002; Jo and Stuart 2009; Wang *et al.* 2023). Weighting by those probabilities is a straightforward approach to incorporating principal scores into the analysis, but matching and regression methods have also been proposed (Jo and Stuart 2009). The weighted per-protocol estimator proposed in the context of medical trials can also be understood as an example of principal score methods (Hernán and Robins 2017; Robins and Finkelstein 2000). All current proposals for using principal score methods; however, rely on some form the principal ignorability assumption. In its weakest form, this assumption requires that the outcomes for units assigned to the control condition are unaffected by compliance type, conditional on the observed covariates. Although

© The Author(s), 2025. Published by Cambridge University Press on behalf of The Society for Political Methodology.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

sometimes plausible, this assumption has been labeled as "quite strong" by previous researchers (Feller *et al.* 2017).

In this paper, I show that a particular principal score estimator can facilitate asymptotically conservative inference about the LATE. This estimator follows a two step process. Specifically, it requires first, directly estimating the probability that each unit is a complier and then weighting the difference in means between the treatment and control groups using these probabilities. To avoid confusion between this estimator and the broader class of principal score methods that it is contained within, I refer to it as the *Compliance Probability Weighting* (CPW) estimator.[1] Intuitively, this weighting leads to an asymptotically conservative estimator for the LATE for two reasons. First, it adjusts for the imbalance in observed covariates between the set of compliers and non-compliers. Second, it upweights covariate strata where there are more compliers, effectively upweighting covariate strata where the gap between the *Intention to Treat* (ITT) estimand and the LATE will be smallest. The remainder of the paper focuses on formalizing this intuition and shows that under mild assumptions, the CPW estimator is asymptotically conservative for the LATE.[2] While previous authors have shown that similar estimators can point identify the LATE under the assumption of principal ignorability, this paper is the first to establish its conservatism without invoking such an assumption.

To demonstrate the utility of the CPW estimator, I use it to re-analyze the Kalla and Broockman (2022) experiment on the persuasive effects of political advertising. Specifically, Kalla and Broockman (2022) randomly aired different political ads in different parts of the country. They then measured the effect of those ads on political attitudes using a survey. Since many respondents did not watch or pay attention to the ads airing in their region, non-compliance was a major concern for their analysis. Kalla and Broockman (2022) measured adherence by including a question focused on a specific policy fact embedded in an ad, but did not condition causal effects on it because the exclusion restriction did not seem plausible. Using the CPW estimator, I provide evidence that, while no persuasive effect is present in the full sample, an effect on attitudes can be observed among the subset of subjects that paid close attention to the ads.

In the next section, I provide some initial details on this illustrative example. I then introduce some basic notation and define the target estimand in Section 3 and turn to reviewing existing approaches for inference in the presence of non-compliance in Section 4. Section 5 introduces the CPW estimator and Section 6 presents some simulation results exploring the performance of this estimator. Section 7 revisits the empirical example and Section 8 concludes.

2. An Illustrative Example

I use an experiment conducted by Kalla and Broockman (2022) as a running example. In this experiment, Kalla and Broockman (2022) partnered with a political advocacy organization to randomly air tv ads in different parts of the country. The ads focused on either LGBT rights or immigration policy. Kalla and Broockman (2022) estimated the effect of ads on political attitudes by conducting a survey shortly after the ads aired.

A major concern for this design is non-compliance. Some respondents might not regularly watch or even own a TV, negating the persuasive effect of television advertising. Similarly, many respondents might not pay much attention to the ads they do see, similarly limiting their persuasive effect. Consequently, it is hard to take the ITT estimand as a reasonable approximation to what the persuasive effect would have been if all respondents closely watched the ad in this setting.[3] To hedge against this concern, Kalla and Broockman (2022) included a question asking whether respondents knew a specific fact regarding immigration policy that was embedded in some of the ads. They observed that the ads had

[1] In particular, this estimator is very similar to those proposed by Stuart and Jo (2015) and Ding and Lu (2017).

[2] That is, the CPW estimator converges in probability to a value which is less than or equal to the LATE.

[3] Note, the ITT estimand represents the average effect of *assignment* to treatment on the outcome, rather than the effect of actual the receipt of the treatment. This is discussed more formally in Section 3.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

a significant effect on whether respondents knew this fact about immigration policy, but they observed no significant effect on political attitudes. Consequently, Kalla and Broockman (2022) concluded that many respondents had paid close attention to the ads, but were not persuaded by them.

Kalla and Broockman 2022's design can be understood as an experiment with non-compliance. In this case, compliers can be defined as respondents that paid careful attention to the ad and learned the policy fact embedded in it. However, IV methods are inappropriate in this setting because is easy to imagine a respondent being persuaded by the ad without noticing the rather minor detail about immigration policy embedded in it, which would represent a straightforward exclusion restriction violation.

3. Notation

I assume that the researcher observes a simple random sample of size N from a population where for each unit $i = 1 \dots N$ they observe (Z_i, T_i, X_i, Y_i). I use Z_i to denote the binary treatment assignment indicator (i.e., the encouragement or instrument) where $Z_i = 1$ indicates the unit was assigned to treatment and $Z_i = 0$ indicates that it was assigned to the control. T_i is a binary indicator for the treatment received. Similar to the assignment indicator, $T_i = 1$ indicates the unit received the treatment and $T_i = 0$ indicates the unit received the control. Finally, $Y_i \in \mathbb{R}$ denotes the outcome and $X_i \in \mathbb{R}^K$ for some integer K denotes a set of pre-treatment covariates. I denote the domain of X_i as $\mathcal{X}$ and its probability law as $f_{X_i}(\cdot)$. As a basic regularity condition, I also require that Y_i have finite first and second moments. I further denote the set of potential outcomes for T_i as $\{T_i(z) : z \in \{0,1\}\}$ and the set of potential outcomes for Y_i as $\{Y_i(z,t) : z \in \{0,1\}, t \in \{0,1\}\}$. Note that the Stable Unit Value Treatment Assumption (SUTVA) for both Z_i and Y_i is implicit in this notation (Rubin 1974). I denote the effect of assignment to treatment on the treatment received for unit i as $\gamma_i = T_i(1) - T_i(0)$ and the effect of assignment to treatment on the outcome as $\tau_i = Y_i(1, T_i(1)) - Y_i(0, T_i(0))$.

I note that because T_i and Z_i are binary, γ_i can only take on 3 values: -1, 0, or 1. I consider a complier to be a unit that receives the treatment when assigned the treatment and the control when assigned the control; that is, a unit for which $\gamma_i = 1$. Units with any other value of γ_i are either non-compliers or defiers. The target of inference will be the treatment effect among this subset of compliers:

$$\begin{aligned} \text{LATE} &= \mathbb{E}\left[Y_i(1, T_i(1)) - Y_i(0, T_i(0)) | T_i(1) = 1, T_i(0) = 0\right] \\ &= \mathbb{E}\left(\tau_i | \gamma_i = 1\right) \end{aligned}$$

The ITT estimand on the other hand is

$$\text{ITT} = E\left[Y_i(1, T_i(1)) - Y_i(0, T_i(0))\right] = \mathbb{E}\left(\tau_i\right).$$

At points I will also refer to the analogs ITT estimator which simply replaces sample quantities for the expectations above:

$$\widehat{ITT} = \hat{\mathbb{E}}\left(Y_i | Z_i = 1\right) - \hat{\mathbb{E}}\left(Y_i | Z_i = 0\right),$$

where $\hat{\mathbb{E}}(\cdot)$ indicates the sample mean for the specified quantity

4. Current Approaches to Estimating the LATE

This section briefly outlines two common approaches to estimating the LATE. Section 4.1 discusses the IV estimator while Section 4.2 reviews principal score methods.

4.1. Instrumental Variables

Angrist *et al.* (1996) introduced the IV estimator for the LATE. Specifically, the IV estimator takes the following form:

$$\frac{\widehat{ITT}}{\hat{\mathbb{E}}\left(T_i | Z_i = 1\right) - \hat{\mathbb{E}}\left(T_i | Z_i = 0\right)},$$

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

where $\hat{\mathbb{E}}(\cdot)$ again indicates the sample mean for the specified quantity. Intuitively, this estimator works by taking the estimated effect of assignment to treatment on the outcome and then dividing it by the estimated effect of assignment to treatment on the treatment received. Angrist *et al.* (1996) show that, under the assumptions discussed in the remainder of this section, the the numerator will converge to the LATE multiplied by the fraction of compliers while the denominator will converge to the fraction of compliers. Consequently, the IV estimator will be a consistent for the LATE if several assumptions are deemed plausible.

The first assumption needed for the consistency of the IV estimator is *random assignment independent of co-variates*.

Assumption 1. *Random assignment independent of covariates: Fraction* $p \in (0,1)$ *of the sample is randomly assigned to treatment* ($Z_i = 1$) *such that* $\forall i \in \{1 \dots N\}$,

$$Z_i \perp\!\!\!\perp \{Y_i(z,t), T_i(z) : z \in \{0,1\}, t \in \{0,1\}\}.$$

This assumption requires that the assignment to treatment be independent of the potential outcomes of the treatment received variable and the outcomes. It is standard in the context of completely randomized experiments and will be maintained for the CPW estimator as well.

The next assumption requires that the average effect of assignment to treatment on the treatment received is non-zero.

Assumption 2. *Relevance:* $\forall i \in \{1 \dots N\}$,

$$\mathbb{E}[T_i(1) - T_i(0)] = \mathbb{E}(\gamma_i) \neq 0.$$

Because the treatment received and treatment assigned are both fully observed by the researcher, Assumption 2 is empirically testable under Assumption 1 (Andrews, Stock, and Sun 2019).

The next assumption required for the IV estimator is monotonicity. This assumption requires that $T_i(z)$ is increasing with z.

Assumption 3. *Monotonicity:* $\forall i \in \{1 \dots N\}$,

$$T_i(1) - T_i(0) = \gamma_i \geq 0.$$

This assumption is also referred to as the no defiers assumption, since it requires that there be no subjects who receive the opposite of the treatment they are assigned to.

The final assumption needed for the IV estimator is the exclusion restriction, which requires that assignment to treatment only impact the outcome through the treatment actually received.

Assumption 4. *The Exclusion Restriction:* $\forall i \in \{1 \dots N\}$ *and* $t \in \{0,1\}$,

$$Y_i(1,t) = Y_i(0,t).$$

Although frequently utilized for analyzing both experiments and quasi-experiments with non-compliance, the assumptions required for the IV estimator are often questioned (Lal *et al.* 2023), suggesting there is a need for methods that can estimate the LATE without invoking these assumptions.

4.2. Principal Score Methods

Principal score methods in essence swap the exclusion restriction and monotonicity for a selection on observables strategy. Specifically, they posit that the differences in average outcomes between the different compliance strata can be explained by a set of observed covariates. Following the lead of

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

Feller *et al.* (2017), I present two forms of this principal ignorability assumption. The stronger version requires that, conditional on X_i, average potential outcomes be the same across covariate strata.

Assumption 5. *Strong Principal Ignorability:* $\forall i \in \{1 \dots N\}$,

$$\mathbb{E}\left[Y_i(1,T_i)|X_i = x, \gamma_i = 1\right] = \mathbb{E}\left[Y_i(1,T_i)|X_i = x, \gamma_i = 0\right] = \mathbb{E}\left[Y_i(1,T_i)|X_i = x, \gamma_i = -1\right]$$

and

$$\mathbb{E}\left[Y_i(0,T_i)|X_i = x, \gamma_i = 1\right] = \mathbb{E}\left[Y_i(0,T_i)|X_i = x, \gamma_i = 0\right] = \mathbb{E}\left[Y_i(0,T_i)|X_i = x, \gamma_i = -1\right].$$

In plain English, this assumption requires that whether a subject complies with their assigned treatment is unrelated to the effect that treatment would have on them after adjusting for covariates. This assumptions is quite strong and likely fails in many applied settings, but it can facilitate the point identification of the LATE without invoking the exclusion restriction or monotonicity.

Consequently, researchers have also considered a weaker form of principal ignorability.

Assumption 6. *Weak Principal Ignorability:* $\forall i \in \{1 \dots N\}$ *and* $t \in \{0,1\}$,

$$\mathbb{E}\left[Y_i(1,1)|X_i = x, T_i(1) = 1, T_i(0) = t\right] = \mathbb{E}\left[Y_i(1,1)|X_i = x, T_i(1) = 1\right]$$

and

$$\mathbb{E}\left[Y_i(0,0)|X_i = x, T_i(0) = 0, T_i(1) = t\right] = \mathbb{E}\left[Y_i(0,0)|X_i = x, T_i(0) = 0\right].$$

The first equation in Assumption 6 requires that, given X_i, the average outcome for units assigned to treatment is not dependent on the treatment they would have received if they were assigned to the control. Similarly, the second equation requires that, given X_i, the average outcome for units assigned to the treatment is not dependent on the treatment they would have received if they were assigned to the control. On its own, weak principal ignorability is not sufficient to point identify the LATE, and identification can only be achieved along with monotonicity (Feller *et al.* 2017).[4] Although more readily plausible than principal ignorability, weak principal ignorability may fail in a number of applied settings.

There are a large number of estimators that identify the LATE under principal ignorability. The CPW estimator is similar to the weighting approach proposed by Stuart and Jo (2015) and further developed by Ding and Lu (2017). Regression (Bein 2015; Joffe, Small, and Hsu 2007) and matching (Hill *et al.* 2002; Jo and Stuart 2009) estimators also exist. Porcher *et al.* (2016) provides a more thorough review of the various principal scores methods.

5. Estimation

This section outlines the CPW estimator. Section 5.1 presents the assumptions needed by the CPW estimator and Section 5.2 compares those assumptions with the IV assumptions and principal ignorability. Section 5.3 defines the CPW estimator and presents the main result about its asymptotic conservatism.

5.1. Assumptions

Along with the standard assumptions of SUTVA, random sampling, positivity, and random assignment to treatment discussed in Sections 3 and 4.1, the CPW estimator will require two additional assumptions.

The first is that the researcher has specified a consistent estimator for the probability that each unit is a complier.

[4] Note that these assumptions can be further weakened in the case of one sided non-compliance or if the exclusion restriction is believed to hold among some compliance strata. See Feller *et al.* (2017) Sections 2 and 3 for a more thorough presentation.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

Assumption 7.

$$\hat{P}(\gamma_i = 1|X_i) \xrightarrow[N\to\infty]{p} P(\gamma_i = 1|X_i)$$

and $\forall x \in \mathcal{X}$,

$$0 < P(\gamma_i = 1|X_i = x) \leq 1.$$

There are many potential choices for estimators that will satisfy this condition. For example, a properly specified parametric model will be consistent and can be used if the analyst deems the functional form assumptions reasonable. Non-parametric options are also available. For example, if monotonicity is assumed to hold, γ_i can only equal 0 or 1, so any estimator for heterogeneity in the effect of assignment to treatment on the treatment received (i.e., $\mathbb{E}(\gamma_i|X_i)$) will also provide an estimate for the probability that $\gamma_i = 1$. The simulations and applications use this approach, relying on the estimator for treatment effect heterogeneity provided by Nie and Wager (2021). In the case of two-sided non-compliance, the estimator provided by Aronow and Carnegie (2013) could be used instead. Note that the identification of the conditional probability required by Assumption 7 will almost certainly require additional assumptions that will depend on the precise estimator that is being used and that may be deemed strong in a particular application. For example, the estimator provided by Aronow and Carnegie (2013) rests on parametric functional form assumptions that may not be plausible in all settings. Even non-parametric estimators, like that provided by Nie and Wager (2021), still typically require that the distribution of the covariates and treatment effects obey some basic regularity conditions. Identification of $P(\gamma_i = 1|X_i)$ will also almost certainly require restrictions on the dimensionality of X_i (typically that it grow at a rate significantly slower than N) and researchers may find LASSO based methods to be particularly useful in settings where X_i is high dimensional (Belloni, Chernozhukov, and Hansen 2014).

This assumption also implicitly requires that some measure of the treatment received, T_i, be observed, as the probability that each unit is a complier is otherwise inestimable. This requirement is, of course, shared by IV and other principal score methods as well. Indeed, in practice, this requirement is actually likely weaker for the CPW estimator than it is for IV methods, as the exclusion restriction is unlikely to be plausible in cases of partial compliance (say where non-compliers receiver a weaker dose of the treatment rather than no treatment) or if the treatment received is measured with error, as is likely the case if compliance is self-reported in an exit survey. Indeed, such settings would represent excellent use cases for the CPW estimator as they are not inconsistent with any of the assumptions discussed in this section holding.

It is worth emphasizing that Assumption 7 requires that the chosen estimator converge to the conditional probability that a unit is a complier given the chosen covariates, it does not require that those covariates themselves be strong predictors of compliance type. Indeed, if X_i is completely independent of γ_i, such that $P(\gamma_i = 1|X_i = x)$ is the same for all values of x, the main result about the asymptotic conservatism of the CPW estimator will still hold as long as the chosen estimator for $\hat{P}(\gamma_i = 1|X_i)$ converges to this constant value as well and the other relevant assumptions are satisfied.

The requirement that $P(\gamma_i = 1|X_i)$ always be greater than zero is worth some discussion as well. In essence, this can be interpreted as something like a positivity requirement, as the effect of the receipt of treatment on the outcome is fundamentally unknowable for covariate strata that contain no compliers. Note, that Assumption 3 implies that this condition holds.

The second assumption needed for the main result to hold is that τ_i be larger for compliers than non-compliers, averaging across covariate strata.

Assumption 8.

$$\int_{\mathcal{X}} f_{X_i|\gamma_i}(X_i = x|\gamma_i = 1)\left(\mathbb{E}[\tau_i|X_i = x, \gamma_i = 1] - \mathbb{E}[\tau_i|X_i = x]\right) dx \geq 0.$$

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

In essence, this assumption requires that the effect of assignment to treatment be larger for compliers than in the full sample when the distribution of pre-treatment covariates is fixed at that seen for just compliers. A more interpretable condition that is sufficient for fulfilling this assumption is that the effect of assignment to treatment is, on average, larger for compliers than non-compliers in all covariate strata. I consider it likely that this assumption holds in many applied settings, as it is reasonable to assume that, even if the exclusion restriction fails, the actual receipt of the treatment should be associated with a larger effect than just the assignment to treatment alone.

I also consider a stronger version of Assumption 8 which will require that compliers have a larger magnitude treatment effect across all covariate strata and that average treatment effects are consistently signed for all covariate strata as well.

Assumption 9. *Either:*

a. $\forall x \in \mathcal{X}, \mathbb{E}(\tau_i|X_i = x, \gamma_i = 1) \geq \mathbb{E}(\tau_i|X_i = x) \geq 0$
b. $\forall x \in \mathcal{X}, \mathbb{E}(\tau_i|X_i = x, \gamma_i = 1) \leq \mathbb{E}(\tau_i|X_i = x) \leq 0.$

Note, a. implies Assumption 8 while b. would imply a similar relationship, but when treatment effects are negative.

5.2. Comparison with IV Assumptions and Principal Ignorability

A natural question is how the assumptions presented above compare with the IV assumptions and principal ignorability. This subsection provides such a discussion, first, considering the IV assumptions and then proceeding to principal ignorability.

5.2.1. IV Assumptions

While the exclusion restriction does not naturally imply the CPW assumptions, it can be shown that if treatment effects are positive for all covariate strata, Assumption 4 does imply Assumption 8, as stated in the following proposition.

Proposition 1. *If for any* $x \in \mathcal{X}$,

$$\mathbb{E}(\tau_i|\gamma_i = 1, X_i = x) > 0.$$

then Assumption 4 implies Assumption 8.

Proof in Section A.1 of the Supplementary Material.

So in situations where the treatment effect is positive for all covariate strata, Assumption 8 will be weaker than the exclusion restriction typically invoked for IV models. Of course, researchers are often unwilling to assume away such backfire effects for all covariate strata, and in such cases, the assumptions needed for the IV estimator will not naturally nest those used for the CPW estimator.[5]

Nonetheless, there are still many applied settings where the assumptions needed for the CPW estimator will be more plausible than the IV ones. For example, if compliance is measured with error, then the exclusion restriction will almost certainly fail, but the CPW estimator will still be useful. Similarly, if non-compliers receive a weaker version of the treatment, then the exclusion restriction will not hold, but the CPW assumptions may be plausible. This is effectively the case in the Kalla and Broockman (2022) experiment where non-compliers are defined as respondents who were likely exposed the ad, but did not pay careful attention. The question of what conditions will lead the CPW estimator to outperform the IV estimator is discussed in more depth in Section 6, which uses simulations

[5] Although see Wood and Porter (2019) and Guess and Coppock (2020) for cases where evidence suggests that such backfire effects do *not* exist.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

to explore the performance of both estimators under different assumptions about the underlying data generating process.

5.2.2. Principal Ignorability

I consider the CPW assumptions to also be more frequently plausible than both versions of principal ignorability. Even in its weaker form, principal ignorability requires that the outcomes not be related to the treatment units would have received if they were assigned to the other treatment group. Such an assumption is likely to be questionable in many social science applications where compliance is inherently related to the level of interest that a subject has in the treatment. For example, in the context of the Kalla and Broockman (2022) experiment, principal ignorability requires that subjects in the control group who would have closely watched the ad and learned the fact about immigration from it not have different attitudes toward immigrants than those who would not. On its face, such an assumption is difficult to sustain as subjects that would closely watch the ad likely have a greater level of interest in immigration policy and could easily hold different opinions. In contrast, the assumptions required for asymptotic conservatism place no restriction on this relationship.

5.3. The CPW Estimator

Intuition about the approach the CPW estimator takes can be developed using the following decomposition of the difference between the LATE and ITT estimands.

Lemma 1.

$$\begin{aligned} LATE - ITT = &\int_{\mathcal{X}} f_{X_i|\gamma_i}(x|\gamma_i = 1)\left[\mathbb{E}\left(\tau_i|X_i = x, \gamma_i = 1\right) - \mathbb{E}\left(\tau_i|X_i = x\right)\right] dx \\ &+ \int_{\mathcal{X}} \mathbb{E}\left[\tau_i|X_i = x\right]\left[f_{X_i|\gamma_i}(x|\gamma_i = 1) - f_{X_i}(x)\right] dx. \end{aligned}$$

Proof in Section A.2 of the Supplementary Material.

Specifically, this decomposition splits the difference between the ITT estimand and the LATE into two portions: the first represents the difference in the effect of assignment to treatment on compliers versus non-compliers while the second represents treatment effect heterogeneity due to the difference in the balance of covariates between compliers and non-compliers. Note that while this difference is helpful for building intuition about the CPW estimator, the ITT estimand is often a major quantity of interest in its own right and this decomposition should not be understood as a representing a form of bias for the ITT estimator.

The CPW estimator works by choosing weights that will eliminate the second term of the difference in Lemma 1. Specifically, I propose weighting by the estimated probability that unit i is a complier given X_i: $\hat{P}(\gamma_i = 1|X_i)$, which will help eliminate the bias due to covariate mismatch and upweight covariate strata that are more likely to include compliers. Formally, this estimator will take the following form:

$$\widehat{\text{CPW}} = \sum_{i=1}^{N} w_i \left(Z_i Y_i - (1 - Z_i) Y_i\right),$$

where

$$w_i = \frac{Z_i \hat{P}(\gamma_i = 1|X_i)}{\sum_{j=1}^{N} Z_i \hat{P}(\gamma_j = 1|X_j)} + \frac{(1 - Z_i)\hat{P}(\gamma_i = 1|X_i)}{\sum_{j=1}^{N}(1 - Z_i)\hat{P}(\gamma_j = 1|X_j)}.$$

Interval and standard error estimates for the CPW estimator can be generated using bootstrap methods, as is standard in the principal score methods literature (Feller *et al.* 2017).

A useful precursor to my main results is to provide a closed form representation for the asymptote of the CPW estimator and the asymptotic difference between the CPW estimator and the LATE.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

Lemma 2. *Under Assumption 7,*

$$\widehat{CPW} \xrightarrow[N\to\infty]{p} \int_{\mathcal{X}} \mathbb{E}\left(\tau_i|X_i = x\right) f_{X_i|\gamma_i}(X_i = x|\gamma_i = 1)dx.$$

And,

$$LATE - \widehat{CPW} \xrightarrow[N\to\infty]{p} \int_{\mathcal{X}} f_{X_i|\gamma_i}(X_i = x|\gamma_i = 1)\left[E\left(\tau_i|X_i = x, \gamma_i = 1\right) - \mathbb{E}\left(\tau_i|X_i = x\right)\right]dx.$$

Proof in Section A.3 of the Supplementary Material.

The first result in this lemma shows that the CPW estimator will converge to the average of the conditional average effect of assignment to treatment within covariate stratum, weighted by the probability that the member of each stratum is a complier. The second result represents the asymptotic difference between the CPW estimator and the LATE as the average of the difference in the effect of assignment to treatment on the outcome between compliers and the full sample among covariate strata, weighted by the probability that a unit in each stratum is a complier. Note that, in the likely case that covariate strata with more compliers have a smaller gap between $\mathbb{E}\left(\tau_i|X_i = x, \gamma_i = 1\right)$ and $\mathbb{E}\left(\tau_i|X_i = x\right)$, the asymptotic bias of the CPW estimator will be falling with the strength of the relationship between the pre-treatment covariates and compliance type.

While there is not an immediate statistic that can be estimated and used to summarize this bias ($\mathbb{E}\left(\tau_i|X_i = x, \gamma_i = 1\right) - \mathbb{E}\left(\tau_i|X_i = x\right)$ cannot be identified without stronger assumptions), researchers will likely still want to consider how well the available covariates predict compliance type. In many cases, theoretical grounds can be used to justify this belief. For example, in the case of the Kalla and Broockman (2022) application, it is quite reasonable to think that general interest in politics will predict how likely a respondent was to be attentive to the ad being aired. It also worth emphasizing that the strength of the relationship between the covariates and compliance type is an empirical question and that the performance of the estimator chosen for $\hat{P}\left(\gamma_i = 1|X_i\right)$ can often be summarized with more general measures of goodness of fit.

The representation of the asymptotic bias of the CPW estimator provided in Lemma 2 also immediately leads to my main result when paired with Assumption 8.

Proposition 2. *Under Assumptions 1, 7, and 8,*

$$\lim_{N\to\infty} P\left(LATE \geq \widehat{CPW}\right) = 1.$$

The proof follows immediately from Assumption 8 and Lemma 2

The interpretation of this proposition will depend on the sign that is assumed for the LATE. If the LATE is negative, then the CPW estimator will still be more negative than the LATE. Since the coding of treatment and control groups as 1 and 0 is arbitrary, it can be assumed, without loss of generality, that the LATE is positive. However, this may be an unsatisfying result as it does not eliminate the possibility that the CPW estimator converge to a large negative value while the LATE remains positive.

A stronger conclusion then is presented below, which uses Assumption 9 to deliver a result about the magnitude of the LATE and CPW estimator.

Proposition 3. *Under Assumptions 1, 7, and 9,*

$$\lim_{N\to\infty} P\left(|LATE| \geq |\widehat{CPW}|\right) = 1.$$

Proof in Section A.4 of the Supplementary Material.

Note that Assumption 9 places stronger limitations on the heterogeneity in treatment effects between covariate strata than Assumption 8, so delivering this stronger conclusion does require stronger assumptions about the nature of the data generating process.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

6. Simulation Results

To explore the performance of the CPW estimator, I conducted a set of simulations aimed at examining how the ITT, IV, and CPW estimators compared.[6] Specifically, I generated a length N treatment assignment vector Z as a set of Bernouli draws with a probability of success of .5, and X as an $N \times K$ matrix of draws from the standard normal distribution. I then generated a length K vector β again as independent draws from a normal distribution with mean 0 and standard deviation σ. Then, for unit i, I let $\mu_i = \text{Logit}^{-1}(X_i'\beta)$ and generated γ_i as a bernouli draw with probability of success of μ_i. Finally, I let $T_i = \gamma_i Z_i$ and $Y_i = \gamma_i Z_i + (1-\gamma_i) Z_i \tau_{NC} + \epsilon_i$ where τ_{NC} is the treatment effect among non-compliers and ϵ_i is a standard normal random variable.[7] Note, implicit in this simulation setup is that non-compliance is one-sided. Section B.5 of the Supplementary Material presents additional simulation results which allow non-compliance to be two-sided instead.[8]

In the simulations, I allowed N to take on values between 100 and 1,000, but fixed K at 10. I also allowed σ to take on values of 0, 1, and 10 and τ_{NC} to take on values of 0, .25, and .5.[9] Note that values of τ_{NC} other than zero represent exclusion restriction violations, while σ controls the strength of the relationship between X_i and γ_i, with more extreme values of β suggesting that X_i is more predictive of γ_i. When $\tau_{NC} = 0$, there is no exclusion restriction violation, so the IV estimator should perform well. The performance of the CPW estimator instead will depend on σ, which controls how strong the relationship between X_i and γ_i is.[10] For each combination of simulation parameters, I generated 100 datasets and calculated the bias and root mean squared error (RMSE) of CPW, ITT, and IV estimators. Although the ITT estimator is typically not used to estimate the LATE (instead, it is unbiased for the ITT estimand, which is a separate quantity of interest), I believe its performance as an estimator for the LATE is a useful benchmark for the CPW estimator as the ITT estimator is also conservative for the LATE under Assumption 8.

Figure 1 visualizes the results of this analysis. As expected, the performance of the CPW estimator relative to the IV and ITT options depends on the two parameters discussed above. When $\sigma = 0$, X_i is not related to γ_i, so the weights used by the CPW estimator will not succeed in upweighting units that are likely to be compliers. Consequently, when $\sigma = 0$, the ITT and CPW estimators both perform similarly. As σ increases, so does the performance of the CPW estimator. The performance of the IV estimator on the other hand is determined by τ_{NC}. When $\tau_{NC} = 0$, the treatment effect for non-compliers is 0, so the exclusion restriction is satisfied. As τ_{NC} increases though, the IV estimator increasingly over-estimates the LATE, while the ITT and CPW estimators remain conservative. Note that the bias of both the CPW and ITT estimators decline with τ_{NC}, suggesting that they will be the most useful when the violation of the exclusion restriction is the most severe.

Overall, these simulation results suggest some concrete guidance about when the CPW estimator ought to be used. Specifically, it will be most useful when there is a strong relationship between the pre-treatment covariates and compliance and when the exclusion restriction is unlikely to hold. On the other hand, if the assumptions of the IV estimator are deemed plausible, it will likely yield the best performance. Finally, if the IV assumptions are unlikely to hold and pre-treatment covariates do not predict compliance, then the ITT estimator may remain the best option.

[6] These simulations were executed using the MIT Super Cloud Reuther *et al.* 2018. Replication materials for these simulations and the empirical example are available at Markovich 2025.

[7] Note, the estimation approach used here uses the non-parametric approach proposed by (Nie and Wager 2021). Section B.4 of the Supplementary Material implements the parametric estimator proposed by Aronow and Carnegie (2013) and considers both the case when that parametric model is correctly and incorrectly specified.

[8] This simulation setup also implicitly fixes the probability of compliance at .5. Section B.1 of the Supplementary Material presents additional results where this probability is allowed to vary.

[9] Section B.2 of the Supplementary Material, I allow K to take on additional values. Section B.3 of the Supplementary Material presents results when $\tau_{NC} = 2$, which implies a failure of Assumption 8.

[10] Note that this simulation set up implicitly assumes that half of the sample is compliers. Section B of the Supplementary Material provides additional simulation results that allow the fraction of compliers to vary.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

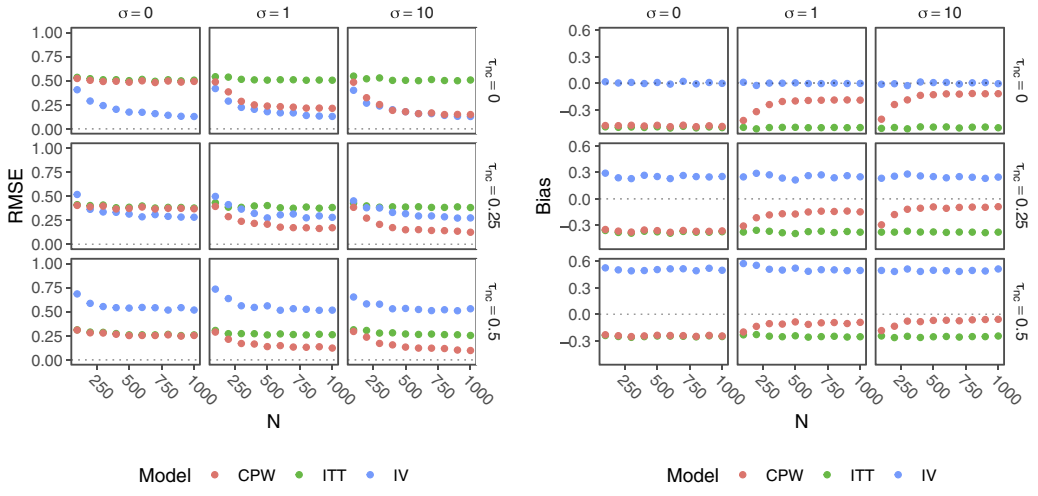


Figure 1. Simulation results.
Note: Figure presents results from simulations comparing efficacy of the ITT, IV, and CPW estimators. The vertical facets identify the value of σ while the horizontal facets identify the value of τ_{NC}. Each point represents the results from 100 simulations with those parameters.

7. Empirical Application

I used the CPW estimator to re-analyze the Kalla and Broockman (2022) study. Specifically, I used an indicator for whether the ad was aired in the respondent's area as the treatment assignment variable, an indicator for whether the respondent knew the fact about immigration policy as the treatment received variable, and the index representing attitudes toward immigration as the outcome. The set of pre-treatment covariates used to predict compliance were composed of a battery of demographic and ideological questions that Kalla and Broockman (2022) gathered in the first wave of the survey, which was conducted before the ads were aired. The LATE in this setting represents the effect of learning about immigration policy from the ad, while ITT estimand represents the effect of living in an area where the ad was aired.

I consider it likely the assumptions needed for the CPW estimator will hold in this setting. In particular, it seems eminently plausible that the effect of the ad is largest for respondents who watched it carefully and learned the fact about immigration policy than for those that did not, satisfying Assumption 8. I also use a non-parametric estimator for $\hat{P}(\gamma_i = 1|X_i)$ which will be consistent as long as non-compliance is one sided.[11]

This reanalysis of Kalla and Broockman (2022)'s experiment shows that there is a meaningful treatment effect for respondents that paid close enough attention to learn the embedded fact about immigration policy. Specifically, the CPW estimator suggests that the ads were able to increase very attentive respondents' score on Kalla and Broockman (2022)'s immigration policy battery by 0.092 (95% ci=(.034, .11)) – much larger than the negligible ITT estimate of 0.0089 (95% ci = (−0.025, 0.043)). These results add important color to Kalla and Broockman (2022)'s conclusion that "Despite the immigration advertisements being memorable and imparting new information, there is not clear evidence that these ads had persuasive effects on issue attitudes." Indeed, they show that there likely was

[11]Specifically, I used the r-learner proposed by Nie and Wager (2021). Specifically, I used the implementation of their algorithm which relies on xgboost for predictions. When implementing this estimator, I also tuned the number of search rounds and number of trees used by xgboost, which I found improved the model's predictive accuracy. Finally, because tree based estimators rely on a random sampling process when generating estimates, I present the average of 10 runs of the algorithm to eliminate this additional variability.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

a persuasive effect on issue attitudes, it just only emerges among that subset of respondents that learned new information from the ad.

These results have direct relevance for scholars of American public opinion. In contrast to evidence suggesting political ads are rarely persuasive (e.g., Coppock, Hill, and Vavreck (2020); Kalla and Broockman (2018)), these results suggest that political advertising can be persuasive when respondents pay close attention to the message being conveyed.[12] Indeed, these results suggest that political advertising might be particularly effective when it is meant to persuade a smaller, more engaged electorate (as is the case in primary campaigns) or when they are aired on more overtly political networks where users might be more inclined to carefully watch a political ad.

Nonetheless, some care must be taken in interpreting these results, as some organizations might prefer to calibrate their advertising strategies based on their performance in the full population (including those who do not own TVs). For example, an organization trying to win elections or enact broad political change might be most interested in the persuasive effect on the full electorate and may not want to base their strategy on the effect seen in the subset of most attentive respondents.

8. Conclusion

Researchers analyzing randomized experiments with non-compliance have few tools currently available. Because the true treatment effect cannot be known for non-compliers, researchers typically focus on the treatment effect just among the subset of units that receive the treatment they were assigned. However, estimating this treatment effect can be difficult. IV estimators represent the dominant approach, but require assumptions (namely, monotonicity and the exclusion restriction) that will not be plausible in all settings. This paper proposes an alternate approach to conducting inference about the LATE. Specifically, it suggests directly estimating the probability that each unit is a complier and then estimating the LATE as the difference in sample means, weighted by these estimated probabilities. This estimator is asymptotically conservative under assumptions that may hold even when the IV assumptions fail.

Note that the CPW estimator is designed to be conservative for the LATE, but it is possible other estimands will be more relevant for answering a particular substantive question. In particular, the LATE provides the effect of treatment on the compliers, but if selection into treatment is strongly determined by the benefit that units receive from that treatment, it may be not representative of the average treatment effect in the full sample.[13] Ultimately, the relevance of the LATE will be determined by the research question being considered. In the context of the Kalla and Broockman (2022) replication, a positive treatment effect for compliers suggests that political advertising effective for a highly attentive subset of the full population. Although this contrasts with past work suggesting that persuasive effects from political advertising are consistently small, with limited heterogeneity (e.g., Coppock *et al.* 2020; Kalla and Broockman 2018), they do not speak to the persuasive effect that might be realized if the full sample could be induced to carefully watch the ads. Indeed, the composition of the population of compliers is always important for the interpretation of the LATE and large amount of recent methodological work considers how this might impact the interpretation of causal effects (Aronow and Carnegie 2013; Marbach and Hangartner 2020).

A major limitation of the CPW estimator is its dependence on estimating the probability that each unit is a complier. When this model can accurately identify which units are likely compliers, the CPW estimator will perform well. If X_i is only a weak predictor of compliance though, this approach will produce weights that are unrelated to a unit's likelihood of complying and might inflate the variance of the estimates without producing any reduction in bias. The CPW estimator also assumes that the researcher has provided a consistent estimator for the conditional probability of compliance given

[12] Although see de Benedictis-Kessner *et al.* (2019); Gerber *et al.* (2011); Wittenberg *et al.* (2021) for examples of past studies that support a larger persuasive effect from television and video advertising.

[13] See Aronow and Carnegie (2013) and Sasaki (n.d.).

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

the covariates, and asymptotic conservatism will not be guaranteed if an incorrectly specified model is used. Since even non-parametric estimators for this probability generally require that the data obey some basic regularity conditions, estimating the probability of compliance will almost certainly require invoking untestable assumptions about the data generating process.

Consequently, the choice of whether to rely on the IV estimator, the ITT estimator, or the proposed CPW estimator will depend on the quality of the data available and the assumptions the analyst is willing to make. If the IV assumptions are considered plausible, then it remains the best choice for estimating the LATE. Indeed, there are certainly many cases when the IV assumptions will be plausible, but the researcher lacks a rich set of covariates that can be used to estimate the probability of compliance or does not feel comfortable specifying a consistent model (either parametric or non-parametric) to estimate those probabilities. In such cases, the IV estimator will remain superior to the CPW estimator. On the other hand, if the IV assumptions fail, but a set of pre-treatment covariates that strongly predict compliance are available, the CPW estimator will likely be an appealing choice. Ultimately, researchers will have to rely on substantive knowledge of the data generating process to assess which set of assumptions are more plausible.

This suggests that, when making a decision about the usefulness of the CPW estimator, a researcher will want to consider the likelihood that the available covariates strongly predict compliance type. Substantive knowledge will be helpful for making such an assessment. In particular, the CPW estimator will be particularly relevant for settings where an observed covariate predicts the utility that a particular unit receives from the treatment. For example, if the treatment involves exposure to political messaging of some form, general interest in politics will likely determine how attentive a subject is to that messaging. This dynamic is present in the Kalla and Broockman (2022) replication, but emerges in many other studies as well (de Benedictis-Kessner *et al.* 2019; Guess *et al.* 2021). A similar dynamic might emerge when studying health related interventions as well, where compliance will be predicted by the perceived health benefit that a subject will receive. It is worth emphasizing that how well the available covariates predict compliance is an empirical question, and that it may be useful to consider various measures of model fit such as the test set loss or model deviance when assessing the utility of the CPW estimator for a particular setting.

Nonetheless, many applied settings seem likely to satisfy these requirements. Survey experiments in particular represent a useful area for using the CPW estimator. It is common for researchers designing survey experiments to include questions after the treatment to make sure it was correctly understood and interpreted by respondents (i.e., factual manipulation checks), but researchers rarely condition causal effects on correctly answering these questions because they are post-treatment and the exclusion restriction is rarely plausible Kane and Barabas (2019); Montgomery, Nyhan, and Torres (2018). The CPW estimator; however, is well suited for such settings as researchers frequently include pre-treatment questions aimed at measured attentiveness that are likely to predict performance on the factual manipulation check. Indeed, these conditions are similar to those present in the Kalla and Broockman (2018) experiment, which also focused on non-compliance among survey respondents. Such conditions can also arise outside of the social science. For example, similar problems are common in experiments that assign subjects to a diet or exercise regimen. In such settings, many subjects will fail to fully comply with the assigned intervention; however, some change in lifestyle is still likely suggesting the exclusion restriction is unlikely to hold. This estimator will also be helpful in cases of partial compliance, where some subjects receive a weaker dose of the treatment. This setting implies an inherent exclusion restriction violation, as a weak treatment likely still has some effect, but the CPW assumptions may be plausible. While the IV estimator will remain invaluable when analyzing experiments with non-compliance when its assumptions are satisfied, I believe the CPW estimator will be a useful tool for analyzing experiments with non-compliance when its assumptions fails.

Acknowledgments. I thank Adam Berinsky, Shiyao Liu, Chloe Wittenberg, Teppei Yamamoto, and all the members of the Kim Research Group for their helpful comments. I also acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing (HPC, database, consultation) resources that have contributed to the research results reported within this paper.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

Data Availability Statement. Replication code for this article is available at Markovich (2025). A preservation copy of the same code and data can also be accessed via Dataverse at https://doi.org/10.7910/DVN/Y6JJ0A.

Supplementary Material. The supplementary material for this article can be found at https://doi.org/10.1017/pan.2025.10005.

References

Andrews, I., J. H. Stock, and L. Sun. 2019. "Weak Instruments in Instrumental Variables Regression: Theory and Practice." *Annual Review of Economics* 11: 727–753.

Angrist, J. D., G. W. Imbens, and D. B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American statistical Association* 91 (434): 444–455.

Aronow, P. M., and A. Carnegie. 2013. "Beyond LATE: Estimation of the Average Treatment Effect with an Instrumental Variable." *Political Analysis* 21 (4): 492–506.

Bein, E. 2015. "Proxy Variable Estimators for Principal Stratification Analyses." Technical report. Abt Associates Working Paper.

Belloni, A., V. Chernozhukov, and C. Hansen. 2014. "Inference on Treatment Effects After Selection Among High-Dimensional Controls." *The Review of Economic Studies* 81 (2): 608–650.

de Benedictis-Kessner, J., M. A. Baum, A. J. Berinsky, and T. Yamamoto. 2019. "Persuading the Enemy: Estimating the Persuasive Effects of Partisan Media with the Preference-Incorporating Choice and Assignment Design." *American Political Science Review* 113 (4): 902–916.

Coppock, A., S. J. Hill, and L. Vavreck. 2020. "The Small Effects of Political Advertising are Small Regardless of Context, Message, Sender, or Receiver: Evidence From 59 Real-Time Randomized Experiments." *Science Advances* 6 (36): eabc4046.

Ding, P., and J. Lu. 2017. "Principal Stratification Analysis Using Principal Scores." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 79 (3): 757–777.

Feller, A., F. Mealli, and L. Miratrix. 2017. "Principal Score Methods: Assumptions, Extensions, and Practical Considerations." *Journal of Educational and Behavioral Statistics* 42 (6): 726–758.

Gerber, A. S., J. G. Gimpel, D. P. Green, and D. R. Shaw. 2011. "How Large and Long-Lasting are the Persuasive Effects of Televised Campaign Ads? Results From a Randomized Field Experiment." *American Political Science Review* 105 (1): 135–150.

Guess, A., and A. Coppock. 2020. "Does Counter-Attitudinal Information Cause Backlash? Results From Three Large Survey Experiments." *British Journal of Political Science* 50 (4): 1497–1515.

Guess, A. M., P. Barberá, S. Munzert, and J. Yang. 2021. "The Consequences of Online Partisan Media." *Proceedings of the National Academy of Sciences* 118 (14): e2013464118.

Hernán, M. A., and J. M. Robins. 2017. "Per-Protocol Analyses of Pragmatic Trials." *The New England Journal of Medicine* 377 (14): 1391–1398.

Hill, J., J. Waldfogel, and J. Brooks-Gunn. 2002. "Differential Effects of High-Quality Child Care." *Journal of Policy Analysis and Management: The Journal of the Association for Public Policy Analysis and Management* 21 (4): 601–627.

Jo, B., and E. A. Stuart. 2009. "On the Use of Propensity Scores in Principal Causal Effect Estimation." *Statistics in Medicine* 28 (23): 2857–2875.

Joffe, M. M., D. Small, and C.-Y. Hsu. 2007. "Defining and Estimating Intervention Effects for Groups that will Develop an Auxiliary Outcome." *Statistical Science* 22 (1): 74–97.

Kalla, J. L., and D. E. Broockman. 2018. "The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments." *American Political Science Review* 112 (1): 148–166.

Kalla, J. L., and D. E. Broockman. 2022. ""outside Lobbying" Over the Airwaves: A Randomized Field Experiment on Televised Issue Ads." *American Political Science Review* 116 (3): 1126–1132.

Kane, J. V., and J. Barabas. 2019. "No Harm in Checking: Using Factual Manipulation Checks to Assess Attentiveness in Experiments." *American Journal of Political Science* 63 (1): 234–249.

Lal, A., M. Lockhart, Y. Xu, and Z. Zu. 2023. "How Much Should We Trust Instrumental Variable Estimates in Political Science? Practical Advice Based on Over 60 Replicated Studies." arXiv preprint arXiv:2303.11399.

Marbach, M., and D. Hangartner. 2020. "Profiling Compliers and Noncompliers for Instrumental-Variable Analysis." *Political Analysis* 28 (3): 435–444.

Markovich, Z. 2025. "Estimating the Local Average Treatment Effect Without the Exclusion Restriction." https://doi.org/10.7910/DVN/Y6JJ0A.

Montgomery, J. M., B. Nyhan, and M. Torres. 2018. "How Conditioning on Posttreatment Variables Can Ruin Your Experiment and What to Do About It." *American Journal of Political Science* 62 (3): 760–775.

Nie, X., and S. Wager. 2021. "Quasi-Oracle Estimation of Heterogeneous Treatment Effects." *Biometrika* 108 (2): 299–319.

Porcher, R., C. Leyrat, G. Baron, B. Giraudeau, and I. Boutron. 2016. "Performance of Principal Scores to Estimate the Marginal Compliers Causal Effect of an Intervention." *Statistics in Medicine* 35 (5): 752–767.

Reuther, A., et al. 2018. "Interactive Supercomputing on 40,000 Cores for Machine Learning and Data Analysis." In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, 1–6. IEEE.

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press

Robins, J. M., and D. M. Finkelstein. 2000. "Correcting for Noncompliance and Dependent Censoring in an AIDS Clinical Trial With Inverse Probability of Censoring Weighted (IPCW) Log-Rank Tests." *Biometrics* 56 (3): 779–788.
Rubin, D. B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66 (5): 688.
Sasaki, T. n.d. "Bayesian Approach to Estimating the Average Treatment Effect in the Presence of Noncompliance." In *Annual Meeting of the Society of Political Methodology*.
Stuart, E. A., and B. Jo. 2015. "Assessing the Sensitivity of Methods for Estimating Principal Causal Effects." *Statistical Methods in Medical Research* 24 (6): 657–674.
Wang, C., Y. Zhang, F. Mealli, and B. Bornkamp. 2023. "Sensitivity Analyses for the Principal Ignorability Assumption Using Multiple Imputation." *Pharmaceutical Statistics* 22 (1): 64–78.
Wittenberg, C., B. M. Tappin, A. J. Berinsky, and D. G. Rand. 2021. "The (Minimal) Persuasive Advantage of Political Video Over Text." *Proceedings of the National Academy of Sciences* 118 (47): e2114388118.
Wood, T., and E. Porter. 2019. "The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence." *Political Behavior* 41: 135–163.

Cite this article: Markovich, Z. (2026). Estimating the Local Average Treatment Effect Without the Exclusion Restriction. *Political Analysis*, 38–52. https://doi.org/10.1017/pan.2025.10005

https://doi.org/10.1017/pan.2025.10005 Published online by Cambridge University Press