

Instrumental Variable Methods for Conditional Effects and Causal Interaction in Voter Mobilization Experiments^{*}

Matthew Blackwell[†]

June 16, 2016

Abstract

In democratic countries, voting is one of the most important ways for citizens to influence policy and hold their representative accountable. And yet, in the United States and many other countries, rates of voter turnout are alarmingly low. Every election cycle, mobilization efforts encourage citizens to vote and ensure that elections reflect the true will of the people. To establish the most effective way of encouraging voter turnout, this paper seeks to differentiate between (1) the *synergy hypothesis* that multiple instances of voter contact increase the effectiveness of a single form of contact, and (2) the *backlash hypothesis* that multiple instances of contact are less effective or even counterproductive. Remarkably, previous studies have been unable to compare these hypotheses because extant approaches to analyzing experiments with noncompliance cannot speak to questions of causal interaction. I resolve this impasse by extending the traditional instrumental variables framework to accommodate multiple treatment-instrument pairs, which allows for the estimation of conditional and interaction effects to adjudicate between synergy and backlash. The analysis of two voter mobilization field experiments provides the first evidence of backlash to follow-up contact and a cautionary tale about experimental design for these quantities.

^{*}Thanks to Peter Aronow, Fred Boehmke, Jake Bowers, Alex Coppock, Don Green, Erin Hartman, Kosuke Imai, Luke Keele, Gary King, Michael LaCour, Jamie Monogan, James Robins, Arturas Rozenas, Maya Sen, Susumu Shikano, Tyler VanderWeele, and Teppei Yamamoto for helpful suggestions and feedback. Special thanks to David Nickerson for sharing his data. All remaining errors are my own.

[†]Assistant Professor, Department of Government, Harvard University, Institute for Quantitative Social Science, 1737 Cambridge Street, Cambridge, MA 02138 (email: mblackwell@gov.harvard.edu, web: mattblackwell.org)

1 Introduction

In democratic countries, casting a ballot is one of the most important ways for ordinary citizens to influence policy or hold their representative accountable. Low voter turnout can skew election results away from the true “will of the people” and result in policies that disproportionately benefit a minority of citizens (see Fowler, 2013, and references therein). To combat or take advantage of this, every election cycle nonprofit groups, political parties, and candidates themselves dedicate tremendous resources to get-out-the-vote (GOTV) efforts. Over the course of the 2012 presidential election, the Barack Obama campaign alone had 2.2 million volunteers make over 150 million door knocks and phone calls designed to encourage (supportive) citizens to vote (Obama for America, 2013), which is remarkable since there were only 127 million votes actually cast in the election.

To combat the effects of low voter turnout on the electoral process, it is crucial to understand the efficacy of these types of GOTV contact and how they might interact. For example, does door-knocking increase turnout? Should it be paired with phone calls or done in isolation? Two rival hypotheses exist in the GOTV literature to answer this latter question: (1) the *synergy hypothesis* that multiple forms of contact combine to produce larger effects than either separately, and (2) the *backlash hypothesis* that additional contact is less effective or even counterproductive compared to a single intervention (Green, McGrath, and Aronow, 2013). These two hypotheses have strong implications for how to maximize voter turnout with a fixed budget. If backlash exists, then GOTV campaigns should focus their efforts on maximizing the number of citizens contacted, whereas if synergy exists, they should attempt more follow-up contact on a smaller group of respondents.

The synergy and backlash hypotheses directly concern the *causal interaction* between multiple forms of contact, a quantity well-understood in traditional randomized experiments with full compliance (VanderWeele, 2015). Unfortunately, experimental studies of GOTV efforts are subject to noncompliance since citizens often fail to answer their door or pick up the phone, which could lead to unmeasured confounding between actual (as opposed to randomized) GOTV contact and voter turnout. Thus, while a large and influential literature in political science has relied on randomized experiments to assess the impact GOTV appeals (Gerber and Green, 2000; Green, Gerber, and Nickerson, 2003; Nickerson, 2007;

Gerber, Green, and Larimer, 2008; Arceneaux and Nickerson, 2009), there has been no study of the interaction between different forms of GOTV contact that has dealt with issues of noncompliance. Gerber and Green (2000) use an intent-to-treat (ITT) analysis to investigate the interaction between the randomization to in-person canvassing and the randomization to telephone canvassing, but non-compliance makes it difficult to ascertain if these ITT interactions reflect the true interaction between *actual* in-person contact and *actual* telephone contact. Furthermore, there is no statistical literature on how one would use instrumental variables, a common approach to handling noncompliance, to estimate these causal interactions.

To remedy this gap, this paper extends the single-treatment instrumental variables framework of Angrist, Imbens, and Rubin (1996) to allow for the identification and estimation of certain conditional and interaction effects of two forms of GOTV contact. Here, each binary randomized assignment to GOTV contact serves as an instrument for the binary treatment of actual GOTV contact and the goal is to identify the joint effects and interactions of the binary treatments. Previous extensions of IV methods to allow for multiple-level treatments (relevant since two binary treatments can be thought of as a single four-level treatment) did not separately identify conditional effects, but rather a weighted average of them that would be difficult to interpret in this context (Angrist and Imbens, 1995). In contrast, I use a principal stratification approach to generalize the concept of compliance types to the setting of two binary treatments and define various “local” quantities of interest, which are average comparisons between pairs of treatment vectors (that is, actual contact) within levels of these compliance types. These include a local average interaction effect (LAIE) that I use to adjudicate between the synergy and backlash hypotheses. I also show that these estimands are nonparametrically identified from the usual IV assumptions and a novel assumptions called the *treatment exclusion restriction*, which ensures that each treatment is only affected by a single instrument. The paper develops consistent and asymptotically normal plug-in estimators for these quantities along with consistent variance estimators. I also show how an interacted version of two-stage least squares (TSLS) can identify the LAIE but cannot identify conditional effects without further assumptions.

I investigate the question of synergy versus backlash in two large GOTV field

experiments that can speak to these questions. In the first, over 11,000 households were randomly assigned to receive telephone canvassing, in-person canvassing, both, or neither ahead of the 1998 general election in New Haven, CT (Gerber and Green, 2000). In this experiment, there is strong variation in conditional effects between compliance groups, indicating that compliance with phone contact, for instance, may be predictive of response to door-to-door canvassing. But this experiment also suffered from very low rates of compliance and an experimental design that minimized the number of respondents that received both forms of treatment. These factors combined to produce very high uncertainty over the conditional and interaction effects.

To overcome issues of low compliance, I analyze a second experiment, conducted by the Youth Voting Coalition (YVC) during the 2004 presidential election, wherein over 26,000 young registered voters across 10 sites were randomly assigned to phone calls from either a volunteer phone bank, a professional phone bank, neither, or both (Nickerson, 2007). The compliance rate in this experiment was far higher than in the New Haven experiment. Here, I find support for the backlash hypothesis among college-aged students—a second round of phone contact produces far weaker mobilizing effects than a single phone call. This represents the first piece of evidence for causal interaction in GOTV efforts. Finally, I discuss how the estimands developed in this paper can be useful for maximizing the cost effectiveness of future GOTV efforts.

2 A Joint Instrumental Variable Framework for GOTV Studies

Using the New Haven experiment as a guide, let $D_{i1} = 1$ indicate that a respondent received a phone canvassing message and $D_{i2} = 1$ indicate that a respondent received in-person canvassing. Furthermore, let $Z_{i1} \in \{0, 1\}$ be random assignment to phone canvassing and Z_{i2} be random assignment to in-person canvassing. Let Y_i be the outcome of interest, whether or not respondent i voted. Define $Y_i(d_1, d_2)$ as the potential outcomes, representing whether or not respondent i would have voted if we set D_{i1} to d_1 and D_{i2} to d_2 (Rubin, 1974). Note that under the usual exclusion restriction (stated formally below), these potential outcomes will not depend on the value of the instruments.

In these experiments, canvassers were unable to contact all respondents assigned to receive contact. For phone contact, this could be caused by incorrect

phone numbers in voter registration records or by respondents refusing to answer the phone. For in-person canvassing, the travel time between physical households could be quite long and could prevent the canvassers from reaching all households intended for contact (Gerber and Green, 2000, p. 655). Furthermore, canvassing in the New Haven experiment ceased at sunset for safety reasons, preventing the canvassers from contacting residents not home during the day. So in these studies D_{ij} may not equal Z_{ij} , but treatments can respond to their instruments, so that $D_{i1}(z_1, z_2)$ is whether or not respondent i would be contacted by phone when assigned to have $Z_{i1} = z_1$ and $Z_{i2} = z_2$ and $D_{i2}(z_1, z_2)$ is defined similarly. Here, I implicitly assume that neither form of actual contact affects the other once the assignments are set so that $D_{i1}(z_1, z_2, d_2) = D_{i1}(z_1, z_2, d'_2) = D_{i1}(z_1, z_2)$. Let $D_{i1}(z_1) = D_{i1}(z_1, Z_{i2})$ and $D_{i2}(z_2) = D_{i2}(Z_{i1}, z_2)$ be the potential outcomes for these variables when only setting one of the two instruments and leaving the other instrument at its observed value. Finally, let $D_i = (D_{i1}, D_{i2})$ and $Z_i = (Z_{i1}, Z_{i2})$. Implicit in the definition of all these potential outcomes is the stable unit treatment value assumption (SUTVA) of no interference between units since the potential outcomes and treatments for unit i only depend on the treatment and instrument values of unit i .

This paper generalizes the single-treatment IV framework of Angrist, Imbens, and Rubin (1996), which makes the following assumptions modified for the presence of two binary treatments:

Assumption 1 (Single-Treatment IV). *For all values $d_1, d_2, z_1, z'_1, z_2, z'_2$ and $j \in \{1, 2\}$, the following statements hold:*

- (i) *Consistency: $Y_i = Y_i(D_{i1}, D_{i2})$, $D_{i1} = D_{i1}(Z_{i1}, Z_{i2})$, and $D_{i2} = D_{i2}(Z_{i1}, Z_{i2})$;*
- (ii) *Randomization of the instruments: (Z_{i1}, Z_{i2}) independent of $\{Y_i(d_1, d_2), D_{i1}(z_1, z_2), D_{i2}(z'_1, z'_2)\}$.*
- (iii) *Outcome exclusion restriction: $Y_i(d_1, d_2, z_1, z_2) = Y_i(d_1, d_2, z'_1, z'_2) = Y_i(d_1, d_2)$.*
- (iv) *Effect of instrument on treatment: $\Pr[D_{ij}(z_j = 1)] \neq \Pr[D_{ij}(z_j = 0)]$ with probability 1.*
- (v) *Monotonicity: $\Pr[D_{ij}(z_j = 1) \geq D_{ij}(z_j = 0)] = 1$.*

These are the standard assumptions of Angrist, Imbens, and Rubin (1996) applied separately to each instrument-treatment pair and allowing for the joint randomization of Z_i . The exclusion restriction is uncontroversial in these GOTV studies since respondents only learn of their random assignment through actual contact. Monotonicity is often, though not always, satisfied by restricting contact to those who are randomly assigned to receive it, which ensures one-sided noncompliance. For the sake of generality, I develop the framework using the more general monotonicity assumption. Note that this framework depends heavily on binary instruments and treatment.

Under Assumption 1, there may be unmeasured confounding between actual (as opposed to assigned) GOTV contact and voter turnout, which renders the average treatment effects of actual contact (D_{ij}) unidentified. Angrist, Imbens, and Rubin (1996) instead identified average causal effects within levels of a respondent’s *compliance type* or *principal strata*, which are defined by how a unit responds to the instrument (Frangakis and Rubin, 2002). In the single-treatment IV case with a binary treatment and instrument, these strata are defined by the vector $(D_i(0), D_i(1))$, which can take on four different values: compliers with $(0, 1)$, always-takers with $(1, 1)$, never-takers with $(0, 0)$, and defiers with $(1, 0)$. Let $C_i \in \{c, a, n, d\}$ denote the compliance type for unit i . Monotonicity implies that there are no defiers in the population and allows for the identification of the local average treatment effect (LATE), which is the average effect of actual contact among compliers: $E[Y_i(1) - Y_i(0) | C_i = c]$. Past analyses of GOTV experiments have leveraged this fact to estimate the LATEs of each form of actual GOTV contact separately. In the GOTV context, compliers are those respondents that are *reachable*—they pick up the phone when called or answer the door when canvassers knock.

While this basic framework has been previously extended beyond single treatments, these endeavors have not addressed the identification or estimation of causal interactions or conditional effects in 2×2 factorial experiments such as the ones considered here. Angrist and Imbens (1995) extended the single-treatment setting to accommodate multiple-level instruments and treatments, focusing on ordinal treatments, where higher values indicate “more” treatment. While it is possible to redefine two binary treatments as a single four-level treatment, their results cannot identify the separate effects of between-level contrasts ($Y_i(1, 0)$ vs.

$Y_i(0, 0)$, for example), but rather a complicated weighted average of these comparisons. Cheng and Small (2006) study noncompliance in randomized experiments to compare the effectiveness of two treatments, but they focus on three-arm trials where both treatments are never assigned together not the full factorial designs of these GOTV studies. Finally, an alternative approach to IV uses a class of models called structural nested mean models (SNMMs), avoiding monotonicity and instead relying homogeneity assumptions on the causal effects to point identify population-average effects (for example, Robins, 1994; Tan, 2010). In the present context, monotonicity holds by design and so I focus on an approach that leverages this fact.

2.1 Joint compliance types, treatment exclusion, and quantities of interest

To generalize single-treatment IV framework, I will define, identify, and estimate average causal effects of actual contact within levels of the two-treatment principal strata. The definition of these principal strata with two forms of contact, which I call the *joint compliance types*, follows a similar logic to the single-treatment case and can be written as a vector of 8 values:

$$(D_{i1}(0, 0), D_{i1}(1, 0), D_{i1}(0, 1), D_{i1}(1, 1), D_{i2}(0, 0), D_{i2}(1, 0), D_{i2}(0, 1), D_{i2}(1, 1)) .$$

This vector completely characterizes how both treatments react to both instruments. The joint compliance types can be reorganized in terms of the single-treatment compliance types by noting that, for example, $\{D_{i1}(0, 0), D_{i1}(1, 0)\}$ defines the compliance type for phone contact when i is assigned to no in-person contact. This vector can take on the four above types—for example, if this vector is $(0, 1)$ then i is a complier for D_{i1} when $z_2 = 0$. Let $C_{i1|z} \in \{c, a, n, d\}$ be the compliance status for D_{i1} when $z_2 = z$ and $C_{i2|z}$ be the compliance status for D_{i2} when $z_1 = z$ for $z \in \{0, 1\}$. We can then characterize the joint compliance types for unit i as $C_i = (C_{i1|0}, C_{i1|1}, C_{i2|0}, C_{i2|1})$, letting $C_i = kmst$ be a shorthand for $C_i = (k, m, s, t)$. With four elements taking on four different values, there are $4^4 = 256$ of these principal strata. One example of these compliance types is the “joint complier” group, $C_i = cccc$, who are reachable by either phone or door-to-door canvassing under any randomization. Another group are those with $C_i = ncnc$ that we might call the “interacted compliers” and who can only be contacted when they are assigned to receive both forms of contact: $D_i(z_1, z_2) = (1, 1)$ if

$(z_1, z_2) = (1, 1)$ and $(0, 0)$ otherwise. For the joint compliers, each treatment only depends on its own instrument, whereas for the interacted compliers, D_{i1} depends on z_2 and D_{i2} depends on z_1 .

To identify and estimate the causal effect of multiple types of actual GOTV contact within joint compliance types, two sets of causal quantities need to be identified: the average potential outcomes within principal strata, $E[Y_i(d_1, d_2)|C_i = kmst]$, and the probability of each principal strata, $\Pr[C_i = kmst]$. To identify these, we will use the observed-data parameters, which consist of the 12 treatment-instrument strata probabilities, $\Pr[D_i = (d_1, d_2)|Z_i = (z_1, z_2)]$, and the 16 average observed outcomes, $E[Y_i|D_i = (d_1, d_2), Z_i = (z_1, z_2)]$. Clearly these 28 observed-data parameters will not be sufficient to identify the 255 compliance probabilities, let alone the full set of causal parameters. Monotonicity in the present setting reduces the number of compliance types to $3^4 = 81$ since there are no defiers, but even this is not enough to achieve identification.

To reduce the number of compliance types even further and identify the causal parameters, I rely on a novel assumption that is justified from the experimental design of the GOTV studies. Namely, if the two forms of canvassing is done by different organizations (as is true both experiments) and if the forms of contact are substantially different (as they are in the New Haven experiment), then it is reasonable to assume that randomization to one form of contact has no impact on the actual receipt of the other form of contact. I call this assumption the *treatment exclusion restriction*:

Assumption 2 (Treatment Exclusion Restriction). *For all values z_1, z'_1, z_2, z'_2 ,*

$$D_{i1}(z_1, z_2) = D_{i1}(z_1, z'_2) = D_{i1}(z_1),$$

$$D_{i2}(z_1, z_2) = D_{i2}(z'_1, z_2) = D_{i2}(z_2).$$

Treatment exclusion says that each instrument only affects its “own” treatment and so the compliance type for one form of contact is unaffected by the randomization of the other contact. This implies that $C_{ij|0} = C_{ij|1} = C_{ij}$ for $j \in \{0, 1\}$ and allows us to simplify the principal strata as $C_i = km$ meaning $C_{i1} = k$ and $C_{i2} = m$, leaving $3^2 = 9$ compliance types and making identification possible. Table 1 shows the mapping between values C_i and the potential contacts,

C_i	$D_{i1}(0)$	$D_{i1}(1)$	$D_{i2}(0)$	$D_{i2}(1)$	LACE(s) for this strata
cc	0	1	0	1	$\tau_{cc,1}(0) = E[Y_i(1,0) - Y_i(0,0) C_i = cc]$ $\tau_{cc,1}(1) = E[Y_i(1,1) - Y_i(0,1) C_i = cc]$ $\tau_{cc,2}(0) = E[Y_i(0,1) - Y_i(0,0) C_i = cc]$ $\tau_{cc,2}(1) = E[Y_i(1,1) - Y_i(1,0) C_i = cc]$
cn	0	1	0	0	$\tau_{cn,1}(0) = E[Y_i(1,0) - Y_i(0,0) C_i = cn]$
nc	0	0	0	1	$\tau_{nc,2}(0) = E[Y_i(0,1) - Y_i(0,0) C_i = nc]$
ca	0	1	1	1	$\tau_{ca,1}(1) = E[Y_i(1,1) - Y_i(0,1) C_i = ca]$
ac	1	1	0	1	$\tau_{ac,2}(1) = E[Y_i(1,1) - Y_i(1,0) C_i = ac]$
aa	1	1	1	1	None
an	1	1	0	0	None
na	0	0	1	1	None
nn	0	0	0	0	None

Table 1: Joint compliance types (principal strata) under treatment exclusion and monotonicity and the local average conditional effects (LACEs) associated with each.

$D_{ij}(z_j)$. In the GOTV experiments, this could be violated if, for instance, actually receiving a phone call changed a person’s propensity to respond to attempts at door-to-door contact. This could occur if a respondent finds a single conversation with an in-person canvasser uncomfortable, discouraging them from answering the phone in future contact attempts so that $C_{i1|0} = c$ and $C_{i1|1} = n$. While this is a strong assumption and must be evaluated in each context, one silver lining is that, under Assumptions 1(i) and (ii), it has testable implications. Conditional on its own instrument, a treatment should be unrelated to the other instrument so that $E[D_{i1}|Z_{i1} = z_1, Z_{i2} = 1] = E[D_{i1}|Z_{i1} = z_1, Z_{i2} = 0]$ and $E[D_{i2}|Z_{i1} = 1, Z_{i2} = z_2] = E[D_{i2}|Z_{i1} = 0, Z_{i2} = z_2]$, which can be tested with, say, a regression of each treatment on both instruments to detect violation of the treatment exclusion restriction.

We can now define the quantities of interest in these GOTV experiments. Their factorial designs allow the investigation of the conditional effect of one form of contact conditional on the other, but noncompliance only allows for identification of these quantities within the above principal strata. I call these the *local average conditional effects*, or LACEs. Each LACE is associated with a particular treatment,

a particular level of the other treatment, and a joint compliance type:

$$\begin{aligned}\tau_{km,1}(d_2) &= E[Y_i(1, d_2) - Y_i(0, d_2)|C_i = km], \\ \tau_{km,2}(d_1) &= E[Y_i(d_1, 1) - Y_i(d_1, 0)|C_i = km].\end{aligned}$$

There are four basic conditional effects with two binary treatments, but not all of them are well-defined for all compliance types. For instance, the contrast $Y_i(1, 0) - Y_i(0, 0)$ is meaningless for always-takers for the second treatment, because this group would logically never observe $Y_i(1, 0)$ or $Y_i(0, 0)$ since they will never have $D_{i2} = 0$. Table 1 shows the basic LACEs in this setup—four for the joint compliers, and one for each of the strata that comply with at least one of the treatments.

The LACEs can provide useful information both to political scientists about how voters make decisions and to candidates and campaigns who can use them to better allocate their resources. For instance, if a campaign has made a failed attempt to call a citizen, then they know that this person is a never-taker for phone calls and can base their decision to canvass their physical address on the LACE that conditions on this information, $\tau_{nc,2}(0)$. The campaign obviously will not know if this respondent is a complier for in-person canvassing, but can use the compliance type probabilities from past studies to derive a prediction for the cost-effectiveness of a second GOTV effort.

The LACEs can be combined to create other interesting estimands. The synergy and backlash hypotheses, for instance, speak to the causal interaction between the two treatments. Thus, I define the *local average interaction effect*, or LAIE, which is the difference in the conditional effects for the joint compliers:

$$\tau_{LAIE} = \tau_{cc,1}(1) - \tau_{cc,1}(0) = E[(Y_i(1, 1) - Y_i(0, 1)) - (Y_i(1, 0) - Y_i(0, 0))|C_i = cc]. \quad (1)$$

This represents the difference in the effect of phone canvassing when it is and is not paired with in-person canvassing, which is the exact quantity needed to adjudicate these two hypotheses. Synergy would imply that $\tau_{LAIE} > 0$, while backlash would imply that $\tau_{LAIE} < 0$. In addition to the LAIE, I also define the *local average joint effect*, or LAJE, as the effect of receiving both treatments for the joint compliers: $\tau_{LAJE} = \tau_{cc,1}(1) + \tau_{cc,2}(0) = E[Y_i(1, 1) - Y_i(0, 0)|C_i = cc]$. Both of these quantities are only defined among the joint compliers since they are functions of

$f_{11 11} = \rho_{cc} + \rho_{aa} + \rho_{ac} + \rho_{ca}$	$f_{11 10} = \rho_{aa} + \rho_{ca}$
$f_{10 11} = \rho_{cn} + \rho_{an}$	$f_{10 10} = \rho_{cc} + \rho_{an} + \rho_{ac} + \rho_{cn}$
$f_{01 11} = \rho_{nc} + \rho_{na}$	$f_{01 10} = \rho_{na}$
$f_{00 11} = \rho_{nn}$	$f_{00 10} = \rho_{nc} + \rho_{nn}$
$f_{11 01} = \rho_{aa} + \rho_{ac}$	$f_{11 00} = \rho_{aa}$
$f_{10 01} = \rho_{an}$	$f_{10 00} = \rho_{an} + \rho_{ac}$
$f_{01 01} = \rho_{cc} + \rho_{nc} + \rho_{na} + \rho_{ca}$	$f_{01 00} = \rho_{na} + \rho_{ca}$
$f_{00 01} = \rho_{cn} + \rho_{nn}$	$f_{00 00} = \rho_{cc} + \rho_{cn} + \rho_{nc} + \rho_{nn}$

Table 2: Relationship between observed strata probabilities, $f_{d_1 d_2 | z_1 z_2} = \Pr[D_i = (d_1, d_2) | Z_i = (z_1, z_2)]$, and compliance types probabilities, $\rho_{km} = \Pr[C_i = km]$, under Assumptions 1 and 2.

multiple LACEs. If a respondent is not reachable by both forms of contact, it is hardly sensible to discuss the interaction or joint effects for them.

It is also possible to aggregate LACEs across compliance types, which can be useful to increase power and to maximize the target population of interest. To do this, we simply take the weighted average of the basic LACEs across the compliance types for which they are defined:

$$\tau_{c-,1}(0) = E[Y_i(1, 0) - Y_i(0, 0) | C_i \in \{cc, cn\}] = (\rho_{cc} + \rho_{cn})^{-1} (\rho_{cc} \tau_{cc,1}(0) + \rho_{cn} \tau_{cn,1}(0)). \quad (2)$$

As we will see below, these weighted averages are the target of TSLS estimation in designs with one-sided noncompliance.

3 Identification and Estimation

To identify these estimands, we must connect the causal parameters with the parameters of the observed data. Table 2 provides a sense of how this works by showing how the joint compliance type probabilities, $\rho_{km} = \Pr[C_i = km]$, relate to the probability of strata defined by combinations of instruments and treatments, $f_{d_1 d_2 | z_1 z_2} = \Pr[D_i = (d_1, d_2) | Z_i = (z_1, z_2)]$, under Assumptions 1 and 2. For instance, one strata of subjects were never contacted and so have $D_i = (0, 0)$, but were randomized to receive both forms of contact and so have $Z_i = (1, 1)$. These

respondents are uniquely identified as never-takers on both treatments:

$$\begin{aligned} f_{00|11} &= \Pr[D_i = (0, 0) | Z_i = (1, 1)] = \Pr[D_{i1}(1) = 0, D_{i2}(1) = 0 | Z_i = (1, 1)] \quad (3) \\ &= \Pr[D_{i1}(1) = 0, D_{i2}(1) = 0] = \rho_{nn} \quad (4) \end{aligned}$$

The first equality comes from consistency, the second from randomization, and the last can be established by noting that nn is the only compliance type in Table 1 that has $D_{ij}(1) = 0$ for all j . Other observed strata are a mixture over several compliance types. We can use similar ideas to connect observed-outcome means to the means of the potential outcomes within principal strata, $E[Y_i(d_1, d_2) | C_i = km]$. Theorem 1 combines these ideas to identify the basic LACEs.

Theorem 1. *Under Assumptions 1 and 2, all of the basic LACEs in Table 1 are non-parametrically identified as:*

$$\begin{aligned} \tau_{cc,1}(0) &= \frac{S_{0110}^1 - S_{0111}^1 - S_{0100}^1 + S_{0101}^1}{f_{11|11} - f_{11|01} - f_{11|10} + f_{11|00}} & \tau_{cc,1}(1) &= \frac{S_{1111}^1 - S_{1110}^1 - S_{1101}^1 + S_{1100}^1}{f_{11|11} - f_{11|01} - f_{11|10} + f_{11|00}} \\ \tau_{cn,1}(0) &= \frac{S_{0111}^1 - S_{0101}^1}{f_{10|11} - f_{10|01}} & \tau_{ca,1}(1) &= \frac{S_{1110}^1 - S_{1100}^1}{f_{11|10} - f_{11|00}} \\ \tau_{cc,2}(0) &= \frac{S_{0101}^2 - S_{0111}^2 - S_{0100}^2 + S_{0110}^2}{f_{11|11} - f_{11|01} - f_{11|10} + f_{11|00}} & \tau_{cc,2}(1) &= \frac{S_{1111}^2 - S_{1110}^2 - S_{1101}^2 + S_{1100}^2}{f_{11|11} - f_{11|01} - f_{11|10} + f_{11|00}} \\ \tau_{nc,2}(0) &= \frac{S_{0111}^2 - S_{0110}^2}{f_{01|11} - f_{01|10}} & \tau_{ac,2}(1) &= \frac{S_{1101}^2 - S_{1100}^2}{f_{11|01} - f_{11|00}} \end{aligned}$$

where $S_{d|z_1 z_2}^j = E[Y_i \mathbb{I}(D_{i,-j} = d) | Z_i = (z_1, z_2)]$, and $\mathbb{I}(\cdot)$ is an indicator function.

The proof for Theorem 1 is given in Appendix C, which follows a similar strategy to Abadie (2003). The quantities in the numerator for the LACEs of treatment j , $S_{d|z_1 z_2}^j$, are expectations of the outcome multiplied by an indicator for the other instrument, which helps select out the appropriate compliance types. Thus, the numerators in these results represent ITT effects of the phone-call and door-to-door randomizations on transformations of the outcome. The denominators represent the estimated probability of being in a particular joint compliance stratum. Theorem 1 implies that the LAIE, LAJE and aggregated LACEs are also identified, since they are simple combinations of the LACEs. The LAIE in particular has an interpretable form. Letting $S_{z_1 z_2} = E[Y | Z_i = (z_1, z_2)]$, we have:

$$\tau_{LAIE} = \frac{(S_{11} - S_{01}) - (S_{01} - S_{00})}{(f_{11|11} - f_{11|01}) - (f_{11|10} - f_{11|00})} \quad (5)$$

Thus, the LAIE is the ratio of the ITT interaction for Z_i and Y_i and an ITT interaction between Z_i and D_i .

All of these identification results have a similar structure: a function of reduced-form relationships divided by a probability of various compliance strata. One potential concern for all of these results is that of weak instruments, where the relationship between the instrument and the treatment is weak (Bound, Jaeger, and Baker, 1995). In the case of a binary instrument and a binary treatment, this is equivalent to low probability of compliance. Thus, estimates of the basic LACEs might have higher variance than the LATE for either of the two treatments separately. This is one reason to focus on the aggregated LACEs compared to ones that focus only a single compliance class. This problem is especially salient in the New Haven experiment below.

Theorem 1 not only identifies the LACEs, the LAJE, and the LAIE, but it also suggests an estimation strategy. With a random sample of N units, $(Y_i, D_{i1}, D_{i2}, Z_{i1}, Z_{i2})$, from an infinite superpopulation, it is possible to estimate each of the expectations above. For example, we replace the quantity $S_{d|z_1 z_2}^1$, with its sample analogue,

$$\hat{S}_{d|z_1 z_2}^1 = \frac{\sum_{i=1}^N Y_i \mathbb{I}(D_{i2} = d) \mathbb{I}(Z_{i1} = z_1) \mathbb{I}(Z_{i2} = z_2)}{\sum_{i=1}^N \mathbb{I}(Z_{i1} = z_1) \mathbb{I}(Z_{i2} = z_2)}.$$

If we replace all expectations in Theorem 1 with their sample quantities, we obtain plug-in estimators for all of the quantities of interest. In Appendix A, I show that, under mild regularity conditions, these estimators are consistent and asymptotically normal and derive a closed-form expression for their asymptotic variance. If the randomization assumption only holds conditional on covariates, it is necessary to modify these estimators to include the covariates. For a set of discrete covariates, X_i , it is possible to calculate each of the above expectations conditional on levels of X_i and then average these stratum-specific effects over the distribution of X_i . I take this approach in the second experiment below, where randomization was stratified by site. With continuous covariates, one could replace the expectations with parametric models and average the estimates over the empirical distribution of X_i .

3.1 Two-stage least squares

In the original analyses of these GOTV experiments, the authors relied on two-stage least squares (TSLS) to estimate the separate effects of both forms of GOTV contact. This approach is justified by Imbens and Angrist (1994), Angrist and Imbens (1995), and Abadie (2003) who document the relationship between TSLS and the LATE. These studies have shown that the TSLS estimand *is* the LATE when the TSLS model includes no covariates and the instrument and treatment are binary. Unfortunately, these results are not immediately applicable to questions of causal interaction between phone and in-person contact. Here, I extend TSLS to allow for interactions and investigate what local conditional or interactive effects it can estimate in the two-treatment setting.

Let \mathbf{D} and \mathbf{Z} be $N \times 4$ design matrices that contain a constant, each instrument or treatment, and their interaction. That is, let $\mathbf{D}_i = (1 \ D_{i1} \ D_{i2} \ D_{i1}D_{i2})$ be a generic row of \mathbf{D} , and $\mathbf{Z}_i = (1 \ Z_{i1} \ Z_{i2} \ Z_{i1}Z_{i2})$ be a generic row of \mathbf{Z} . Let $\mathbf{Y}^\top = (Y_1 \ Y_2 \ \dots \ Y_N)$ be the vector of outcomes. With these definitions in hand, it is possible to define the interacted TSLS estimator (iTSLS), which is simply the application of TSLS to the above design matrices: $\widehat{\boldsymbol{\delta}} = (\mathbf{D}^\top \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{Y}$. This estimator allows for an interaction between the two treatments and is fully saturated in both stages. Under standard regularity conditions, this estimator will converge to $\boldsymbol{\delta} = (E[\mathbf{D}_i^\top \mathbf{Z}_i])^{-1} E[\mathbf{Z}_i^\top Y_i]$. Finally, let the limiting coefficient vector be labelled as $\boldsymbol{\delta}^\top = (\delta_0 \ \delta_1 \ \delta_2 \ \delta_3)$.

How does $\boldsymbol{\delta}$ map onto the estimands defined above? The following theorem shows this connection in the case where the instruments are independent.

Theorem 2. *Suppose that Assumptions 1 and 2 hold and that $Z_{i1} \perp\!\!\!\perp Z_{i2}$ and let $\omega_{ck,1} = \rho_{ck}/(\rho_{cc} + \rho_{cn} + \rho_{ca})$ and $\omega_{kc,2} = \rho_{kc}/(\rho_{cc} + \rho_{nc} + \rho_{ac})$ for $k \in \{c, a, n\}$. Then,*

$$\delta_1 = \omega_{cc,1} \tau_{cc,1}(0) + \omega_{cn,1} \tau_{cn,1}(0) + \omega_{ca,1} [\tau_{ca,1}(1) - \tau_{LAIE}] \quad (6)$$

$$\delta_2 = \omega_{cc,2} \tau_{cc,2}(0) + \omega_{nc,2} \tau_{nc,2}(0) + \omega_{ac,2} [\tau_{ac,2}(1) - \tau_{LAIE}] \quad (7)$$

$$\delta_3 = \tau_{cc,1}(1) - \tau_{cc,1}(0) = \tau_{LAIE}. \quad (8)$$

A proof is provided in Appendix C and a simulation exercise in Appendix B confirms the divergence between the plug-in estimators and the TSLS approach.

The coefficient on the interaction between the two treatments is indeed the interaction effect for joint compliers, τ_{LAIE} . The lower order terms are more complicated. The coefficient on D_{i1} , for instance, is a mixture of three different effects, weighted by the probability of cc , cn , and ca (respectively) conditional being a complier for D_{i1} . The first and second components of this mixture are the LACEs of D_{i1} when $d_2 = 0$ for the joint compliers (cc) group and the complier-never-taker group (cn) respectively. But there is no such LACE for the ca group, so the third component of the mixture is the LACE of D_{i1} when $d_2 = 1$ for the ca group *minus the interaction effect for the joint compliers*. Essentially, TSLS is imputing $Y_i(1, 0) - Y_i(0, 0)$ for a group that could never see this effect (ca) by drawing information from the joint compliers about the interaction between the two treatments. This has two ramifications. First, it will mean that the lower order terms will have no direct interpretation as a causal effect without further assumptions. Second, by combining three compliance groups, TSLS estimands will be less affected by weak instrument issues than the LACEs or LAIE because TSLS draws on a larger group of units.

In both of the experiments I analyze here, there is one-sided noncompliance because respondents are never contacted when randomized to receive no contact. Under this design there are no always-takers, and so the coefficients on the lower order terms will be equal to the aggregated LACEs, $\delta_1 = \tau_{c-,1}(0)$ and $\delta_2 = \tau_{-c,2}(0)$. With or without one-sided noncompliance, it is not possible to take the sum of coefficients to estimate different causal contrasts as is possible with a typical linear model. For example, the LAJE is not equal to $\delta_1 + \delta_2 + \delta_3$ because of treatment effect heterogeneity between compliance groups. As is clear from Theorem 2, each coefficient mixes over effects from different joint compliance types, so that their sum would not reflect an effect for the joint compliers. Overall, interpretation of TSLS requires care in these experiments and may not accurately reflect the LACEs and the LAJE. It is important to note that Theorem 2 relies on independence of the instruments, which is not always satisfied and is not required for the identification of the LACEs.

4 Data Analysis of Two GOTV Experiments

4.1 New Haven 1998 GOTV Experiment

In this section, I use the above methods to analyze the New Haven experiment, focusing on one arm of the study that was a 2×2 factorial design as described above. Another arm of the experiment sent mailers to households, but I omit this arm since there is no measurement of compliance with the mailers treatment. Because the phone treatment was randomized differently depending on the mailer treatment, I subset to households that were randomized to the control condition for the mailers arm, leaving 11,689 households. Randomization was done at the household level and so the outcome in this case is the number of voters in the household that turned out in the 1998 general election. Since the sampled households had at most two residents, the treatment effects can range from -2 to 2. Imai (2005) and Gerber and Green (2005) provide a robust discussion of various methodological and technical issues with this particular experiment, most of which are irrelevant to the current discussion, while Hansen and Bowers (2009) analyze this experiment with randomization-based inference.

The New Haven experiment is a good example of a setting where the above assumptions are met, but that the design and compliance rates of the study make it difficult to draw conclusions. Random assignment of the two forms of contact were carried out independently, with $\Pr[Z_{i1} = 1] = 0.1$ and $\Pr[Z_{i2} = 1] = 0.2$ so that the only 5% of the households were randomly selected to receive both forms of contact. Furthermore, I estimate the compliance probabilities in this setting are $\hat{\rho}_{cc} = 0.117$, $\hat{\rho}_{cn} = 0.217$, and $\hat{\rho}_{nc} = 0.127$, meaning that there are very few joint compliers in this example. Thus, while Assumption 1(iv) is satisfied, the effect of Z_{ij} on D_{ij} is very weak in this experiment. Monotonicity holds here by design since those in the control groups are unable to receive contact and so there are no defiers or always-takers. Finally, as discussed above, treatment exclusion appears plausible given the design, but could be violated if randomized assignment into in-person (phone) contact affected whether phone (in-person) contact was actually made. Because of the randomization, this can be tested. In this case, the effect of in-person randomization on phone contact is -0.013 (95% confidence interval: $[-0.073, 0.048]$) and the effect of phone randomization on in-person contact is 0.042 (95% CI: $[-0.022, 0.106]$). Thus, in this particular context there appears to

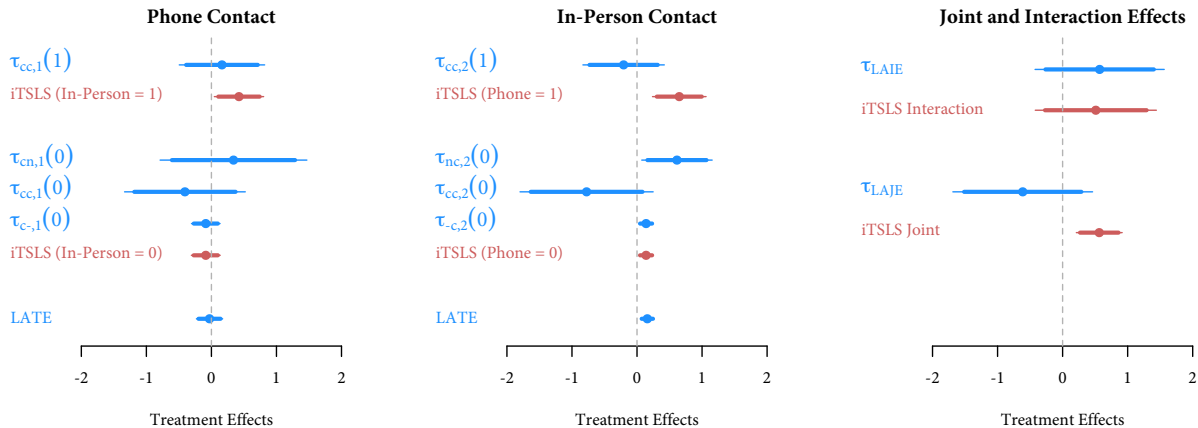


Figure 1: Estimation of LACEs of phone contact (left) and in-person contact (middle), along with the LAJE and LAIE (right). Thin (thick) lines are 95% (90%) confidence intervals, respectively, using the proposed variance estimators proposed. Estimates using an interacted TSLS model and the overall LATE for each treatment are in red. $N = 11,689$ households.

be little evidence of a violation of the treatment exclusion.

Figure 1 shows the estimates of the LACEs, the LAJE, and the LAIE using the plug-in estimators proposed in Section 3 with confidence intervals using the proposed variance estimators from Appendix A. For both treatments, the Figure additionally shows the overall LATE and the effect estimated from the iTSLS approach of Section 3.1. In the original analysis, Gerber and Green (2000) used a standard IV ratio estimator for each instrument-treatment pair separately and I find estimated LATEs that are very similar to theirs—a positive effect of in-person contact and no effect of phone contact. The basic LACE estimates using the above estimators show a great deal of variation across comparisons and across compliance types. For example, the point estimates of the joint-complier effects of the phone and in-person contact alone, $\hat{\tau}_{cc,1}(0)$ and $\hat{\tau}_{cc,2}(0)$, suggest that these forms of contact decrease turnout by a large amount. These point estimates are four to five times larger in magnitude than the estimated effects that aggregate across different compliance groups, $\hat{\tau}_{c-,1}(0)$ and $\hat{\tau}_{c-,2}(0)$. Furthermore, with in-person contact, the effects for phone never-takers, $\hat{\tau}_{nc,2}(0)$ is positive with a 95% confidence interval that excludes zero. The same LACE for the joint compliers is strongly negative, but with a wide confidence interval that includes zero. While it is very difficult to make inferences about the LACEs in this setting, there is suggestive evidence that

there is strong heterogeneity in treatment effects by compliance type.

What about the relationship between the plug-in and iTSLS estimates? Given the one-sided noncompliance and the results of Theorem 2, it is unsurprising that the aggregated LACEs are almost identical to the lower-order term estimates of the iTSLS model. On the other hand, the point estimates $\widehat{\tau}_{cc,1}(1)$ and $\widehat{\tau}_{cc,2}(1)$, which are the effects of each contact when a household has been contacted by the other method, show stark contrasts with the iTSLS approach. Here, the iTSLS estimates are obtained by adding the lower-order and interaction terms. But as discussed above, this combines effect estimates from different subpopulations and therefore will be biased for the $\tau_{cc,j}(1)$. In addition to differences in the point estimates, the iTSLS confidence intervals are also narrower than the plug-in estimators since they draw on more compliance groups for estimation.

Shifting to combinations of the basic LACEs, the plug-in estimator and iTSLS produce almost identical estimates for the LAIE, which is expected given Theorem 2. While the point estimates are in favor of the synergistic hypothesis, the uncertainty is far too large in this experiment to learn much of anything about the interaction. With joint effects, there are more striking differences between the proposed estimators and the iTSLS approach. The plug-in LAJE estimate is strongly negative, -0.613 , though with high uncertainty, while the joint effect from the iTSLS model finds a joint effect of 0.563 and is statistically significant at the typical levels. Thus, the iTSLS and the LAJE have similar magnitudes but in the opposite direction. Once again, though, the uncertainty is far lower using iTSLS. Overall, the results of the New Haven experiment are muddled by massive uncertainty caused by the experimental design and the low rates of compliance.

One concern that we might have in this setting is that the low rates of compliance are driving significant biases or causing our confidence intervals to have poor performance. To assess this possibility, I simulated data from a one-sided noncompliance design with similar compliance rates, sample size, and assignment probabilities to this experiment. In these simulations, presented in Appendix B, I find that the plug-in estimators to have a bias of roughly 12% of the true parameter value but also have confidence intervals with close to nominal coverage or perhaps slight overcoverage. Thus, while the point estimates here may be suspect due to bias from the weak instrument problem, the confidence intervals should

accurately reflect the uncertainty due to low compliance.

4.2 2004 Youth Vote Coalition Experiment

To overcome the issues of low compliance rates in the last experiment, I now turn to an even larger field experiment. Nickerson (2007) reported the results of a multi-site experiment conducted through the Youth Vote Coalition (YVC), a nonpartisan organization dedicated to GOTV efforts for citizens under the age of 30. In each of the 17 counties that served as experimental sites, individual registered voters were randomly assigned in a 2×2 design to receive a call from a volunteer phone bank (Z_{i1}), a professional phone bank (Z_{i2}), both, or neither. Previous research had shown that volunteer phone banks were effective at voter mobilization, but that professional phone banks (like those used in the New Haven experiment) had less efficacy. The outcome of this study was turnout in the 2004 presidential election, as measured by the public voter file. The YVC experiment was designed to compare these two types of calls in the same populations at the same time. For similar reasons to the New Haven experiment, there was noncompliance on calls from both types of phone banks. In the original analysis, Nickerson (2007) used TSLS with no interaction to estimate the LATE for each type of contact separately, but did not consider the potential interaction between the two types of calls. This interaction is important as it could tell us whether multiple forms of GOTV contact are synergistic or whether there is a backlash against additional instances of contact.

To explore these issues, I reanalyze this experiment and focus on college-age respondents (ages 17–22 at the time of the call), who are often the subject of these interventions. As with the New Haven experiment, many of the assumptions here are satisfied by the design of the experiment. In analyzing the data, I found treatment exclusion violations in 3 of the experimental sites and 3 that had effectively assigned only one form of contact. I excluded one other site because it had a joint compliance rate of just 2%, though this has little effect on the estimates below. This left 26,974 respondents spread over 10 experimental sites. In a test of treatment exclusion in this group, the effect of volunteer-bank randomization on professional-bank contact is 0.001 (95% confidence interval: $[-0.015, 0.017]$) and the effect of professional-bank randomization on volunteer-bank contact is -0.006 (95% CI: $[-0.021, 0.009]$). Thus, in the remaining experimental sites, there appears to be




Figure 2: Combined LACE and LAIE estimates among joint compliers across experimental sites from the Nickerson (2007) youth GOTV experiment. Thin (thick) lines are 95% (90%) confidence intervals based on pooled variance across the sites. $N = 26,974$ registered voters aged 17–22.

little evidence of a violation of the treatment exclusion.

Two additional features make the YVC experiment appealing compared to the New Haven experiment. First, each type of phone bank was randomized with probability 0.5 so that each of the four possible assignment vectors had equal probability of occurring. Second, the compliance types are more evenly balanced in this setting, with $\hat{\rho}_{cc} = 0.23$, $\hat{\rho}_{cn} = 0.29$, $\hat{\rho}_{nc} = 0.15$, and $\hat{\rho}_{nn} = 0.33$. Thus, the probability of joint compliance in this experiment is double that of the New Haven experiment. Both of these features should lead to much lower uncertainty compared to the previous experiment. Figure 2 shows the results of this experiment for the joint compliers. Because of the stratified random sampling by site, these results are the averages of the site-specific LACE and LAIE plug-in estimators, weighted by the sample size of site. The confidence intervals are based on the standard combined variance calculation from a stratified randomized experiment (see, for example, Imbens and Rubin, 2015, p. 204), using the variance estimator from Appendix A for the within-site variance estimates.

It is possible to learn much more about the joint effectiveness of GOTV efforts in this experiment. The results show that both the volunteer and professional phone bank calls had a large positive impact on turnout among respondents, at least when they only received one call ($\tau_{cc,1}(0)$ and $\tau_{cc,2}(0)$). The a single professional phone call, for instance, had a 20 percentage-point effect on turnout. The

effects of one phone-bank contact when paired with the other phone bank ($\tau_{cc,1}(1)$ and $\tau_{cc,2}(1)$), on the other hand, were estimated to be negative but not statistically significant. And, in fact, there is a large, negative interaction between these effects as indicated by the LAIE. Thus, there appear to be rather severe diminishing returns to additional instances of GOTV contact in this study, a finding that is in support of the backlash hypothesis and in contrast with much of the previous literature either which has not investigated interactions such as these or have found very small or null ITT interactions. This novel result tells us that in settings like this, it is much more valuable to reach a larger group of individuals rather than concentrate multiple phone calls on a smaller group. It is not clear if these backlash results are due to the two forms of contact being similar (both phone calls) and whether we would see backlash in a properly powered study of other forms of contact.

Finally, these results can be used by campaigns to evaluate the cost-effectiveness of second contact attempts. For instance, the LACE of volunteer contact among noncompliers for professional contact is negative and very close to zero ($\hat{\tau}_{cn,1}(0) = -0.007$, 95% CI: $[-0.053, 0.039]$), whereas the same effect for joint compliers is large and positive as shown in Figure 2. Thus, if a campaign has made a professional phone bank call and failed to reach a respondent, they can predict that a volunteer call will produce very little effect on turnout. Thus, in future GOTV efforts, the results of these analyses could provide valuable information for campaigns attempting to maximize voter outreach.

5 Conclusion

Causal interaction between different forms of GOTV contact has been understudied in part due to a lack of statistical tools to do so. This paper extends the single-treatment IV framework to allow for the estimation of conditional and interaction effects in factorial experiments with noncompliance. With this framework in hand, the paper was able to provide evidence for the backlash hypothesis in GOTV contact, which has large ramifications for how campaigns should attempt to maximize voter turnout. In another experiment, inference is much more difficult due to the weak instrument problem and a suboptimal experimental design.

The statistical framework here is applicable far beyond GOTV studies—it can be applied to any 2×2 factorial design with noncompliance on each factor. The as-

assumptions needed are a combination of the single-treatment IV assumptions and a treatment exclusion restriction, which helps with identification by limiting the number of principal strata. This paper provides plug-in estimators for the LACEs, which are average comparisons between levels of treatment conditional on compliance types. I show that TSLS can be used to estimate the local average interaction effect and, in experiments with one-sided noncompliance, can also be used to estimate aggregated version of the LACEs.

There are several ways in which this statistical framework could be extended. Most obviously, it would be straightforward to allow for $K > 2$ factors in the framework, though the statistical power to estimate LACEs and LAIEs in that setting may be low. It is also important to develop a theoretical understanding of how to bound treatment effects when the treatment exclusion restriction is violated. Another obvious direction for future research is to develop a similar set of results and estimators for situations with non-binary treatments and instruments and with continuous covariates. For instance, one could extend the local average response function approach of Abadie (2003) to a setting with multiple treatments and multiple instruments.

References

- Abadie, Alberto (2003). “Semiparametric instrumental variable estimation of treatment response models.” *Journal of Econometrics* 113.2, pp. 231–263 (cit. on pp. 12, 14, 22).
- Angrist, Joshua D. and Guido W. Imbens (1995). “Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity.” *Journal of the American Statistical Association* 90.430, pp. 431–442 (cit. on pp. 3, 6, 14).
- Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin (1996). “**Identification of Causal Effects Using Instrumental Variables (with discussion)**.” *Journal of the American Statistical Association* 91, pp. 444–455 (cit. on pp. 3, 5, 6).
- Arceneaux, Kevin and David W. Nickerson (2009). “**Who Is Mobilized to Vote? A Re-Analysis of 11 Field Experiments**.” *American Journal of Political Science* 53.1, pp. 1–16 (cit. on p. 3).
- Bound, John, David A Jaeger, and Regina M Baker (1995). “Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogeneous Explanatory Variable is Weak.” *Journal of the American Statistical Association* 90.430, p. 443 (cit. on pp. 13, 30).
- Cheng, Jing and Dylan S Small (2006). “Bounds on causal effects in three-arm trials with non-compliance.” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 68.5, pp. 815–836 (cit. on p. 7).
- Fowler, Anthony (2013). “Electoral and policy consequences of voter turnout: Evidence from compulsory voting in Australia.” *Quarterly Journal of Political Science* 8.2, pp. 159–182 (cit. on p. 2).
- Frangakis, Constantine E. and Donald B. Rubin (2002). “**Principal stratification in causal inference**.” *Biometrics* 58, pp. 21–29 (cit. on p. 6).

- Freedman, David A (2008). “On regression adjustments in experiments with several treatments.” *The Annals of Applied Statistics* 2.1, pp. 176–196 (cit. on p. 34).
- Gerber, Alan S. and Donald P. Green (2000). “The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment.” *American Political Science Review* 94.03, pp. 653–663 (cit. on pp. 2–5, 17).
- (2005). “Correction to Gerber and Green (2000), replication of disputed findings, and reply to Imai (2005).” *American Political Science Review* 99.02, pp. 301–313 (cit. on p. 16).
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer (2008). “Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment.” *The American Political Science Review* 102.1, pp. 33–48 (cit. on p. 3).
- Green, Donald P., Alan S. Gerber, and David W. Nickerson (2003). “**Getting Out the Vote in Local Elections: Results from Six Door-to-Door Canvassing Experiments.**” *Journal of Politics* 65.4, pp. 1083–1096 (cit. on p. 2).
- Green, Donald P., Mary C McGrath, and Peter M Aronow (2013). “Field Experiments and the Study of Voter Turnout.” *Journal of Elections, Public Opinion & Parties* 23.1, pp. 27–48 (cit. on p. 2).
- Hansen, Ben B and Jake Bowers (2009). “Attributing effects to a cluster-randomized get-out-the-vote campaign.” *Journal of the American Statistical Association* 104.487, pp. 873–885 (cit. on p. 16).
- Hoeffding, Wassily (1948). “A Class of Statistics with Asymptotically Normal Distribution.” *Annals of Mathematical Statistics* 19.3, pp. 293–325 (cit. on p. 35).
- Imai, Kosuke (2005). “Do get-out-the-vote calls reduce turnout? The importance of statistical methods for field experiments.” *American Political Science Review* 99.02, pp. 283–300 (cit. on p. 16).
- Imbens, Guido W. and Joshua D. Angrist (1994). “Identification and Estimation of Local Average Treatment Effects.” *Econometrica* 62.2, p. 467 (cit. on p. 14).
- Imbens, Guido W. and Donald B. Rubin (2015). *Causal Inference for Statistics, Social, and Behavioral Sciences*. Cambridge University Press (cit. on p. 20).
- Lehmann, E L (1999). *Elements of Large-Sample Theory*. Springer (cit. on pp. 26, 34, 35).

- Nickerson, David W. (2007). “Quality Is Job One: Professional and Volunteer Voter Mobilization Calls.” *American Journal of Political Science* 51.2, pp. 269–282 (cit. on pp. 2, 4, 19, 20).
- Obama for America (2013). *2012 Obama Campaign Legacy Report*. Tech. rep. (cit. on p. 2).
- Robins, James M. (1994). “Correcting for non-compliance in randomized trials using structural nested mean models.” *Communications in Statistics* 23.8, 2379–2412 (cit. on p. 7).
- Rubin, Donald B. (1974). “Estimating causal effects of treatments in randomized and nonrandomized studies.” *Journal of educational Psychology* 66.5, p. 688 (cit. on p. 4).
- Tan, Zhiqiang (2010). “Marginal and Nested Structural Models Using Instrumental Variables.” *Journal of the American Statistical Association* 105.489, pp. 157–169 (cit. on p. 7).
- VanderWeele, Tyler (2015). *Explanation in causal inference: methods for mediation and interaction*. Oxford University Press (cit. on p. 2).

A Asymptotic properties of the plug-in estimators

To develop the asymptotic properties of the the plug-in estimators, I focus on the plug-in estimator for $\widehat{\tau}_{cc,1}(0)$, but the development of other quantities follows similarly. The estimator is:

$$\widehat{\tau}_{cc,1}(0) = \frac{\widehat{S}_{0|10}^1 - \widehat{S}_{0|11}^1 - \widehat{S}_{0|00}^1 - \widehat{S}_{0|01}^1}{\widehat{f}_{1|11} - \widehat{f}_{1|01} - \widehat{f}_{1|10} + \widehat{f}_{1|00}}, \quad (9)$$

where

$$\widehat{f}_{1|z_1 z_2} = \frac{\sum_{i=1}^N D_{i1} D_{i2} \mathbb{I}(Z_{i1} = z_1) \mathbb{I}(Z_{i2} = z_2)}{\sum_{i=1}^N \mathbb{I}(Z_{i1} = z_1) \mathbb{I}(Z_{i2} = z_2)}.$$

To derive the asymptotic properties of this estimator, I focus on the case of random sampling from an infinite superpopulation so that the samples defined by each combination of Z_{i1} and Z_{i2} are themselves simple random samples from a population distribution. The instrument group sizes, $N_{z_1 z_2} = \sum_{i=1}^N \mathbb{I}(Z_{i1} = z_1) \mathbb{I}(Z_{i2} = z_2)$, grow with N and their fraction of the sample converges to $N_{z_1 z_2}/N \rightarrow \theta_{z_1 z_2}$. Let $\sigma_{z_1 z_2}^2 = \text{var}[Y_i(1 - D_{i2})|Z_i = (z_1, z_2)]$ and $\nu_{z_1 z_2}^2 = \text{var}[D_{i1} D_{i2}|Z_i = (z_1, z_2)]$. Under standard results on sampling (see Theorem 6.1.4 of Lehmann, 1999), it is straightforward to derive the asymptotic variances of the numerator and denominator of (9) as

$$V^S = \frac{\sigma_{10}^2}{\theta_{10}} + \frac{\sigma_{11}^2}{\theta_{11}} + \frac{\sigma_{00}^2}{\theta_{00}} + \frac{\sigma_{01}^2}{\theta_{01}}, \quad (10)$$

and

$$V^f = \frac{\nu_{11}^2}{\theta_{11}} + \frac{\nu_{01}^2}{\theta_{01}} + \frac{\nu_{10}^2}{\theta_{10}} + \frac{\nu_{00}^2}{\theta_{00}}, \quad (11)$$

respectively. Theorem 3 gives the regularity conditions under which (9) will be asymptotically normal and gives an expression for the asymptotic variance of the estimator. Here, I focus on the LAJE, but results for the LACEs follow in a similar fashion.

Theorem 3. *Suppose that the conditions of Theorem 1 hold and that, in addition, we have (i) for all $(z_1, z_2) \in \{0, 1\} \times \{0, 1\}$, $N_{z_1 z_2} \rightarrow \infty$ as $N \rightarrow \infty$ with $N_{z_1 z_2}/N \rightarrow \theta_{z_1 z_2}$ for some $\theta_{z_1 z_2} \in (0, 1)$ and $N_{z_1 z_2} > 0$ for all N ; (ii) conditional on $(Z_{i1}, Z_{i2}) = (z_1, z_2)$, the sample consists of independent draws from $(Y_i, D_{i1}, D_{i2})|(Z_{i1}, Z_{i2}) =$*

(z_1, z_2) of size $N_{z_1 z_2}$; (iii) $\rho_{cc} > \eta$ for some $\eta > 0$; and (iv) $E[|Y_i(d_1, d_2)|^2] < L < \infty$, for all d_1, d_2 . Then,

$$\sqrt{N}(\widehat{\tau}_{cc,1}(0) - \tau_{cc,1}(0)) \xrightarrow{d} \mathcal{N}(0, V), \quad (12)$$

where,

$$V = \frac{1}{\rho_{cc}^2} V^S + \frac{S^2}{\rho_{cc}^4} V^f - 2 \frac{S}{\rho_{cc}^3} \text{cov}(\widehat{S}, \widehat{\rho}_{cc}), \quad (13)$$

and

$$\begin{aligned} S &= S_{0|10}^1 - S_{0|11}^1 - S_{0|00}^1 + S_{0|01}^1 \\ \widehat{S} &= \widehat{S}_{0|10}^1 - \widehat{S}_{0|11}^1 - \widehat{S}_{0|00}^1 + \widehat{S}_{0|01}^1 \\ \widehat{\rho}_{cc} &= \widehat{f}_{1|11} - \widehat{f}_{1|01} - \widehat{f}_{1|10} + \widehat{f}_{1|00}. \end{aligned}$$

While this theorem focuses on one of these estimators estimator, all of the plug-in estimators for the LACEs will follow a similar form. The covariance in (16) is the covariance between the numerator and the denominator of the ratio estimator. Note that with an infinite superpopulation, the covariance between $\widehat{S}_{0|z_1 z_2}^1$ and $\widehat{f}_{1|z'_1 z'_2}$ will be zero unless $z_1 = z'_1$ and $z_2 = z'_2$. This is because otherwise the estimators are calculated on different independent random samples. Thus, it follows that:

$$\begin{aligned} \text{cov}(\widehat{S}, \widehat{\rho}_{cc}) &= - \left(\frac{\text{cov}(Y_i(1 - D_{i2}), D_{i1} D_{i2} | Z_i = (1, 0))}{\theta_{10}} + \frac{\text{cov}(Y_i(1 - D_{i2}), D_{i1} D_{i2} | Z_i = (1, 1))}{\theta_{11}} \right. \\ &\quad \left. + \frac{\text{cov}(Y_i(1 - D_{i2}), D_{i1} D_{i2} | Z_i = (0, 0))}{\theta_{00}} + \frac{\text{cov}(Y_i(1 - D_{i2}), D_{i1} D_{i2} | Z_i = (0, 1))}{\theta_{01}} \right). \end{aligned} \quad (14)$$

Note the minus sign that occurs here because the signs of the quantities related to a specific instrument value are flipped between the numerator and the denominator.

To consistently estimate the variance in Theorem 3, one can replace the population parameters, variances, and covariances with their sample counterparts. Let $\widehat{\sigma}_{z_1 z_2}^2$ be the sample variance of $Y_i(1 - D_{i2})$ among those units with $Z_{i1} = z_1$ and $Z_{i2} = z_2$, then a consistent estimator for V^S is

$$\widehat{V}^S = \frac{N \widehat{\sigma}_{11}^2}{N_{11}} + \frac{N \widehat{\sigma}_{10}^2}{N_{10}} + \frac{N \widehat{\sigma}_{01}^2}{N_{01}} + \frac{N \widehat{\sigma}_{00}^2}{N_{00}}, \quad (15)$$

where we replace the sample proportion $N_{z_1 z_2} / N$ for the limit, $\theta_{z_1 z_2}$. Following the same line of reasoning with $\widehat{v}_{z_2 z_2}^2$ being the sample variance of $D_{i1} D_{i2}$ among

those with $Z_{i1} = z_1$ and $Z_{i2} = z_2$, gives a consistent estimator, \widehat{V}^f , for V^f . Finally, it is possible to replace the covariances in (14) with the sample covariances of $Y_i(1 - D_{i2})$ and $D_{i1}D_{i2}$ in each combination of treatments to derive a consistent estimator for (14). Combining these consistent estimators together yields a consistent estimator of the asymptotic variance due to Slutsky's theorem:

$$\widehat{V} = \frac{1}{\widehat{\rho}_{cc}^2} \widehat{V}^S + \frac{\widehat{S}^2}{\widehat{\rho}_{cc}^4} \widehat{V}^f - 2 \frac{\widehat{S}}{\widehat{\rho}_{cc}^3} \widehat{\text{cov}}(\widehat{S}, \widehat{\rho}_{cc}). \quad (16)$$

Of course, it is often straightforward to use the nonparametric bootstrap to estimate confidence intervals as well. In Appendix B, I simulate the performance of this variance estimator, which appears to do well even at small sample sizes, and I further explore the use of a GMM estimator that attempts to leverage the additional information inherent in Table 2 to generate a more efficient estimator for these quantities. These simulations show that the GMM estimator produces very similar results to the above plug-in estimator.

B Simulation study

This appendix investigates the finite-sample performance of the plug-in estimator with a simulation experiment. I generate the potential outcomes with heterogeneous effects that depend on the compliance type across two treatment-instrument pairs.

$$Y_i(d_1, d_2) = \beta_{i0} + \beta_{i1}d_1 + \beta_{i2}d_2 + \beta_{i3}d_1d_2 + \varepsilon_i \quad (17)$$

Here, the exact parameters for each unit depends on the compliance type for each unit, as given in Table 3 and $\varepsilon_i \sim N(0, 1)$. I generate each instrument separately from a completely randomized design, so that exactly half of the units have $Z_{ij} = 1$ and treatments are generated according to their compliance types. This design easily satisfies the randomization and exclusion restriction assumptions. To generate the compliance types, I set ρ_{cc} and then all other compliance probabilities have $(1 - \rho_{cc})/8$, with no probability for defiers in either treatment.

I apply two estimators to these simulated data sets: the LAJE estimator from Section 3 and the iTSLS estimator from Section 3.1. From this second estimator, I calculate the implied joint effect of the two treatments ($\delta_1 + \delta_2 + \delta_3$) and use standard variance calculations of joint effects to calculate confidence intervals.

Compliance Type	Simulation Parameters
$C_{i1} = c, C_{i2} = c$	$\beta_{ik} \sim \mathcal{N}(1, 1)$
$C_{i1} = c, C_{i2} = a$	$\beta_{ik} \sim \mathcal{N}(0.2, 1)$
$C_{i1} = c, C_{i2} = n$	$\beta_{ik} \sim \mathcal{N}(0.1, 1)$
$C_{i1} = a, C_{i2} = c$	$\beta_{ik} \sim \mathcal{N}(-0.1, 1)$
$C_{i1} = a, C_{i2} = a$	$\beta_{ik} \sim \mathcal{N}(0.2, 1)$
$C_{i1} = a, C_{i2} = n$	$\beta_{ik} \sim \mathcal{N}(0.2, 1)$
$C_{i1} = n, C_{i2} = c$	$\beta_{ik} \sim \mathcal{N}(-0.4, 1)$
$C_{i1} = n, C_{i2} = a$	$\beta_{ik} \sim \mathcal{N}(0.2, 1)$
$C_{i1} = n, C_{i2} = n$	$\beta_{ik} \sim \mathcal{N}(0.7, 1)$

Table 3: Simulation parameters that depend on the compliance status

	N	80% Joint Compliance			40% Joint Compliance		
		Bias	RMSE	95% CI Coverage	Bias	RMSE	95% CI Coverage
LAJE	200	-0.017	0.563	0.939	0.122	1.456	0.955
	1000	-0.007	0.231	0.961	0.007	0.559	0.946
	2500	-0.004	0.153	0.950	0.012	0.346	0.953
iTSLs	200	-0.176	0.539	0.930	-0.651	1.055	0.863
	1000	-0.153	0.264	0.894	-0.679	0.769	0.510
	2500	-0.148	0.206	0.782	-0.666	0.708	0.173

Table 4: Simulation results for the performance of the LAJE estimator and iTSLs for various samples sizes and probabilities of joint compliance. Based on 1,000 simulations.

The results of this simulation study, based on 1,000 simulations, are presented in Table 4 with the bias, root mean square error (RMSE), and the coverage of the 95% confidence intervals. A few notable patterns emerge from this table. First, the performance of all estimators, in terms of bias or RMSE are worse as the probability of joint compliance falls. With half the probability of joint compliance, the RMSE of the estimators is more than double for some sample sizes. Second, for the LAJE, the performance of the estimator generally improves with higher sample sizes, with RMSE and bias decreasing with fairly stable coverage for the confidence intervals. The iTSLs estimator, on the other hand, degrades as the sample size in-

creases, with bias remaining fixed and the coverage probability of the confidence intervals falling well below their nominal rate. This is to be expected given the bias derived in Section 3.1. In simulations not reported here, the bias in the TSLS estimation of the LACEs is very similar to what is predicted from the identification results in Section 3.1. With only 40% compliance and a sample size of 2,500, for instance, the 95% confidence interval for the iTSLS joint effect only covers the true value in 17.3% of the simulations. Finally, the combination of small samples and high noncompliance can lead to poor performance of the LAJE estimator, which is a common feature of instrumental variable methods (Bound, Jaeger, and Baker, 1995).

Finally, in a separate simulation, I check the performance of these estimators in a situation very similar to the New Haven experiment. I use $\rho_{cc} = 0.1$, $\rho_{cn} = \rho_{nc} = 0.2$, and $\rho_{nn} = 0.5$ with the other compliance types set to probability 0, as in the one-sided design. I use $N = 11,500$ which is slightly lower than the sample I analyzed in the main text and use the low treatment probabilities of $\Pr[Z_{i1} = 1] = 0.1$ and $\Pr[Z_{i2} = 1] = 0.2$. The simulations are otherwise the same as above. I find that the plug-in estimators have good performance in this setting, with the $\hat{\tau}_{LAIE}$ having a bias of -0.144 (relative to a true value of 1), a RMSE of 1.56, and a 95% confidence interval coverage of 0.953. The joint effect estimator, $\hat{\tau}_{LAJE}$ has a similar performance, with a bias of -0.126 , a RMSE of 1.073, and coverage of 0.966. This is roughly the same as performance of the LAJE with 40% joint compliance and an N of 200. The estimators for other quantities give similar results. Thus, it appears that even in the face of very low joint compliance rates, these estimator give sensible confidence intervals with roughly correct coverage and relatively little bias with samples this large.

I investigate the performance of a GMM estimator with another simulation study. Using the same data generating process as above, with a joint compliance of 0.4 and an $N = 1000$, apply a GMM estimator using moment conditions of the following type:

$$E[Z_{i1}Z_{i2}(Y_iD_{i1}D_{i1} - \psi_{11|cc}\rho_{cc} - \psi_{11|aa}\rho_{aa} - \psi_{11|ac}\rho_{ac} - \psi_{11|ca}\rho_{ca})] = 0, \quad (18)$$

$$E[Z_{i1}Z_{i2}(D_{i1}D_{i2} - \rho_{cc} - \rho_{aa} - \rho_{ac} - \rho_{ca})] = 0. \quad (19)$$

Here, $\psi_{d_1d_2|jk} = E[Y_i(d_1, d_2)|C_i = (j, k)]$. These conditions follow from the proof of Theorem 1 in Appendix C. There are 28 such moment restrictions—12 for the

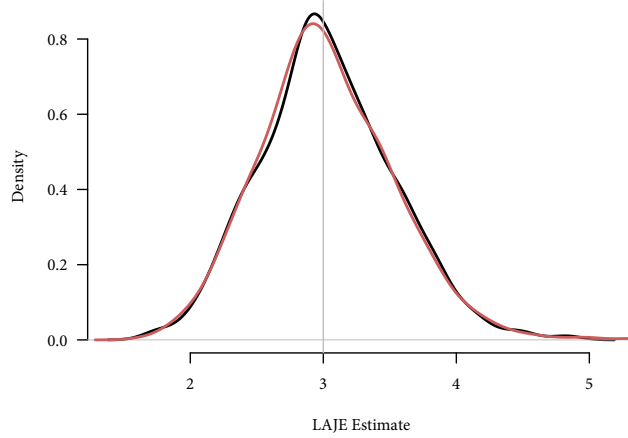


Figure 3: Monte Carlo sampling distributions of the GMM estimator (black) and the plug-in estimator (red) of the LAJE when $N = 1000$ and the probability of joint compliance is 0.4. The true LAJE of 3 is marked by a grey vertical line.

treatments and 16 for the outcomes—which implies that a GMM estimator is overidentified since there are only 24 parameters—8 for the treatment and 16 for the potential outcomes. I use the optimal two-step GMM estimator for these restrictions. I compare this estimator to the plug-in LAJE estimator in the main text. The results of this simulation can be seen in Figure 3, where the black line is the density of the simulated sampling distribution of the GMM estimator and the red line is the simulated sampling distribution of the plug-in estimator. What is clear is that the two sampling distributions are almost identical. This indicates that the efficiency gains from GMM may be limited, though exploring this avenue more fully is an interesting direction for future research.

C Proofs

Proof of Theorem 1. Let $\psi_{d_1 d_2 | km} = E[Y_i(d_1, d_2) | C_i = km]$ be the compliance type specific potential outcome means. Also let $S_{d_1 d_2 | z_1 z_2} = E[Y_i \mathbb{I}(D_{i1} = d_1) \mathbb{I}(D_{i2} = d_2) | Z_i = (z_1, z_2)]$, noting that because the D_{ij} are binary we have $S_{d | z_1 z_2}^1 = S_{1d | z_1 z_2} + S_{0d | z_1 z_2}$.

First, note that given the usual consistency and exclusion restriction assump-

tion, we can write the observed outcome as a function of the potential outcomes and the treatment conditions:

$$Y_i = Y_i(0, 0) + D_{i1}[Y_i(1, 0) - Y_i(0, 0)] + D_{i2}[Y_i(0, 1) - Y_i(0, 0)] + D_{i1}D_{i2}([Y_i(1, 1) - Y_i(0, 1)] - [Y_i(1, 0) - Y_i(0, 0)]). \quad (20)$$

Given this, we have the following:

$$S_{d_1 d_2 | z_1 z_2} = E[Y_i(d_1, d_2) \mathbb{I}(D_{i1} = d_1) \mathbb{I}(D_{i2} = d_2) | Z_i = (z_1, z_2)]. \quad (21)$$

From these statements, it is possible to use the law of iterated expectations over compliance classes, noting that expectations will only be non-zero for a subset of the classes, as shown in Table 2. For instance, $S_{11|11}$ will only take positive values in the compliance classes that have $D_i = (1, 1)$ when $Z_i = (1, 1)$, which are $C_i \in \{cc, aa, ac, ca\}$. Expanding this logic to all of these regression functions:

$$S_{11|11} = \psi_{11|cc} \rho_{cc} + \psi_{11|aa} \rho_{aa} + \psi_{11|ac} \rho_{ac} + \psi_{11|ca} \rho_{ca} \quad (22)$$

$$S_{11|10} = \psi_{11|aa} \rho_{aa} + \psi_{11|ca} \rho_{ca} \quad (23)$$

$$S_{11|01} = \psi_{11|aa} \rho_{aa} + \psi_{11|ac} \rho_{ac} \quad (24)$$

$$S_{11|00} = \psi_{11|aa} \rho_{aa} \quad (25)$$

$$S_{10|11} = \psi_{10|cn}\rho_{cn} + \psi_{10|an}\rho_{an} \quad (26)$$

$$S_{10|10} = \psi_{10|cc}\rho_{cc} + \psi_{10|an}\rho_{an} + \psi_{10|ac}\rho_{ac} + \psi_{10|cn}\rho_{cn} \quad (27)$$

$$S_{10|01} = \psi_{10|an}\rho_{an} \quad (28)$$

$$S_{10|00} = \psi_{10|an}\rho_{an} + \psi_{10|ac}\rho_{ac} \quad (29)$$

$$S_{01|11} = \psi_{01|nc}\rho_{nc} + \psi_{01|na}\rho_{na} \quad (30)$$

$$S_{01|10} = \psi_{01|na}\rho_{na} \quad (31)$$

$$S_{01|01} = \psi_{01|cc}\rho_{cc} + \psi_{01|nc}\rho_{nc} + \psi_{01|na}\rho_{na} + \psi_{01|ca}\rho_{ca} \quad (32)$$

$$S_{01|00} = \psi_{01|na}\rho_{na} + \psi_{01|ca}\rho_{ca} \quad (33)$$

$$S_{00|11} = \psi_{00|nn}\rho_{nn} \quad (34)$$

$$S_{00|10} = \psi_{00|nc}\rho_{nc} + \psi_{00|nn}\rho_{nn} \quad (35)$$

$$S_{00|01} = \psi_{00|cn}\rho_{cn} + \psi_{00|nn}\rho_{nn} \quad (36)$$

$$S_{00|00} = \psi_{00|cc}\rho_{cc} + \psi_{00|cn}\rho_{cn} + \psi_{00|nc}\rho_{nc} + \psi_{00|nn}\rho_{nn} \quad (37)$$

Similar logic gives us the identification results presented in Table 2. With these in hand, it is easy to show that, for instance,

$$\begin{aligned} & S_{0|10}^1 - S_{0|11}^1 - S_{0|00}^1 + S_{0|01}^1 \\ &= (S_{10|10} + S_{00|10}) - (S_{10|11} + S_{00|11}) - (S_{10|00} + S_{00|00}) + (S_{10|01} + S_{00|01}) \\ &= (\psi_{10|cc} - \psi_{00|cc})\rho_{cc} \end{aligned}$$

Using Table 2, it is possible to identify the joint complier probability using the various observed treatment-instrument strata:

$$\rho_{cc} = f_{11|11} - f_{11|10} - f_{11|01} + f_{11|00}$$

Combining these two steps, dividing both sides by ρ_{cc} , and noting that $\tau_{cc,1}(0) = \psi_{10|cc} - \psi_{00|cc}$, gives us the first result: $\tau_{cc,1}(0) = (S_{0|10}^1 - S_{0|11}^1 - S_{0|00}^1 + S_{0|01}^1) / (f_{11|11} - f_{11|01} - f_{11|10} + f_{11|00})$. The other results follow from the same logic. \square

Proof of Theorem 3. First note that condition (iv) implies that both Y_i and $Y_i(1 - D_{i2})$ have finite second moments. Let $Y_i^+ = Y_i(1 - D_{i2})$ and $Y_i^+(z_1, z_2) = Y_i(D_{i1}(z_1), D_{i2}(z_2))(1 -$

$D_{i2}(z_2)$). Note that by consistency we have

$$\hat{S}_{0|z_1 z_2}^1 = \frac{1}{N_{z_1 z_2}} \sum_{Z_i=(z_1, z_2)} Y_i^+ = \frac{1}{N_{z_1 z_2}} \sum_{Z_i=(z_1, z_2)} Y_i^+(z_1, z_2).$$

Note that the numerator can be written as:

$$\hat{S}_N = \hat{S}_{0|10}^1 - \hat{S}_{0|11}^1 - \hat{S}_{0|00}^1 - \hat{S}_{0|01}^1.$$

I rely on Proposition 1 of Freedman (2008) and the infinite superpopulation to note the following:

$$E[\hat{S}_{0|z_1 z_2}^1] = E[Y_i^+(z_1, z_2)] = S_{0|z_1 z_2}^1 \quad (38)$$

$$\text{var}[\hat{S}_{0|z_1 z_2}^1] = \frac{1}{N_{z_1 z_2}} \text{var}[Y_i^+(z_1, z_2)] = \frac{1}{N_{z_1 z_2}} \sigma_{z_1 z_2}^2 \quad (39)$$

$$\text{cov}(\hat{S}_{0|z_1 z_2}^1, \hat{S}_{0|z'_1 z'_2}^1) = 0 \text{ unless } z_1 = z'_1, z_2 = z'_2. \quad (40)$$

Since each $\hat{S}_{0|z_1 z_2}^1$ is unbiased for $S_{0|z_1 z_2}^1$, we have $E[\hat{S}_N] = S_{0|10}^1 - S_{0|11}^1 - S_{0|00}^1 + S_{0|01}^1$ and the variance will be:

$$\text{var}[\hat{S}_N] = \frac{\sigma_{10}^2}{N_{10}} + \frac{\sigma_{11}^2}{N_{11}} + \frac{\sigma_{00}^2}{N_{00}} + \frac{\sigma_{01}^2}{N_{01}}.$$

To derive the asymptotic distribution of this quantity, note that \hat{S}_N can be seen as a multiple-sample U -statistic with kernel

$$\phi(Y_i^+(1, 0), Y_j^+(1, 1), Y_s^+(0, 0), Y_t^+(0, 1)) = Y_i^+(1, 0) - Y_j^+(1, 1) - Y_s^+(0, 0) + Y_t^+(0, 1),$$

where the different subscripts indicate each value is from a different (z_1, z_2) population. Relying on conditions (i), (ii), and (iv), we can apply Theorem 6.1.4 of Lehmann (1999) to find that

$$\sqrt{N}(\hat{S}_N - S) \xrightarrow{d} \mathcal{N}(0, V^S),$$

where,

$$V^S = \frac{\sigma_{11}^2}{\theta_{11}} + \frac{\sigma_{10}^2}{\theta_{10}} + \frac{\sigma_{01}^2}{\theta_{01}} + \frac{\sigma_{00}^2}{\theta_{00}}.$$

By a similar logic, we have

$$\sqrt{N}(\hat{\rho}_{cc} - \rho_{cc}) \xrightarrow{d} \mathcal{N}(0, V^f),$$

where

$$V^f = \frac{v_{10}^2}{\theta_{10}} + \frac{v_{11}^2}{\theta_{11}} + \frac{v_{00}^2}{\theta_{00}} + \frac{v_{01}^2}{\theta_{01}}.$$

Furthermore, since both \hat{S}_N and $\hat{\rho}_{cc}$ are U -statistics, then as shown by Hoeffding (1948), the two U -statistics converge in distribution jointly (see also, Theorem 6.1.6 of Lehmann, 1999). Specifically, we have

$$\sqrt{N} \begin{pmatrix} \hat{S}_N - S \\ \hat{\rho}_{cc} - \rho_{cc} \end{pmatrix} \xrightarrow{d} \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} V^S & \text{cov}[\hat{S}_N, \hat{\rho}_{cc}] \\ \text{cov}[\hat{S}_N, \hat{\rho}_{cc}] & V^f \end{bmatrix} \right). \quad (41)$$

The proof is completed by applying the multivariate delta method to the transformation $\hat{\tau}_{cc,1}(0) = \frac{\hat{S}_N}{\hat{\rho}_{cc}}$. □

To prove the next theorem, it will be useful to have a slight generalization of the Frisch-Waugh-Lovell theorem as applied to TSLS models with multiple endogenous variables. This will allow characterizing coefficients of a larger TSLS model in a simple manner. Suppose that \mathbf{Z} is an $N \times k$ matrix of instruments and \mathbf{X} is an $N \times k$ matrix of treatments, so that an IV model is just identified. Each matrix is equivalently partitioned by columns so that \mathbf{Z}_1 and \mathbf{X}_1 are $N \times k_1$ submatrices of \mathbf{Z} and \mathbf{X} , respectively, and \mathbf{Z}_2 and \mathbf{X}_2 are $N \times k_2$, where $k = k_1 + k_2$. Let \mathbf{b}_1 and \mathbf{b}_2 be the coefficients associated with each set of treatments in a TSLS regression. That is, these quantities solve the following set of normal equations:

$$\begin{bmatrix} \mathbf{Z}_1^\top \mathbf{X}_1 & \mathbf{Z}_1^\top \mathbf{X}_2 \\ \mathbf{Z}_2^\top \mathbf{X}_1 & \mathbf{Z}_2^\top \mathbf{X}_2 \end{bmatrix} \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{Z}_1^\top \mathbf{Y} \\ \mathbf{Z}_2^\top \mathbf{Y} \end{bmatrix} \quad (42)$$

Finally, let $\mathbf{M}_1 = (\mathbf{I} - \mathbf{X}_1(\mathbf{Z}_1^\top \mathbf{X}_1)^{-1} \mathbf{Z}_1^\top)$ be a residual-maker; that is, \mathbf{M}_1 creates residuals from a TSLS regression of \mathbf{Y} on \mathbf{X}_1 using \mathbf{Z}_1 as instruments. Note that these residuals are formed by multiplying the TSLS coefficients by the treatments here. Let $\mathbf{Y}^* = \mathbf{M}_1 \mathbf{Y}$ and $\mathbf{X}_2^* = \mathbf{M}_1 \mathbf{X}_2$.

Lemma 1. *Under suitable regularity conditions for TSLS,*

$$\mathbf{b}_2 = (\mathbf{Z}_2^\top \mathbf{X}_2^*)^{-1} \mathbf{Z}_2^\top \mathbf{Y}^*. \quad (43)$$

Proof. From the first set of normal equations from (42), we can solve for \mathbf{b}_1 :

$$\mathbf{b}_1 = (\mathbf{Z}_1^\top \mathbf{X}_1)^{-1} \mathbf{Z}_1^\top (\mathbf{Y} - \mathbf{X}_2 \mathbf{b}_2)$$

Stating the second set of normal equations, plugging in the value of \mathbf{b}_1 , and rearranging gives the desired result:

$$\begin{aligned} \mathbf{Z}_2^\top \mathbf{X}_1 \mathbf{b}_1 + \mathbf{Z}_2^\top \mathbf{X}_2 \mathbf{b}_2 &= \mathbf{Z}_2^\top \mathbf{Y} \\ \mathbf{Z}_2^\top (\mathbf{I} - \mathbf{X}_1 (\mathbf{Z}_1^\top \mathbf{X}_1)^{-1} \mathbf{Z}_1^\top) \mathbf{X}_2 \mathbf{b}_2 &= \mathbf{Z}_2^\top (\mathbf{I} - \mathbf{X}_1 (\mathbf{Z}_1^\top \mathbf{X}_1)^{-1} \mathbf{Z}_1^\top) \mathbf{Y} \\ \mathbf{b}_2 &= (\mathbf{Z}_2^\top \mathbf{M}_1 \mathbf{X}_2)^{-1} \mathbf{Z}_2^\top \mathbf{M}_1 \mathbf{Y} \end{aligned}$$

□

Proof of Theorem 2. Let $\pi_{z_1 z_2} = \Pr[Z_1 = z_1, Z_2 = z_2]$. To simplify notation let $D_{i3} = D_{i1} D_{i2}$ and $Z_{i3} = Z_{i1} Z_{i2}$. Furthermore, let $\Delta(A_i | B_i) = E[A_i | B_i = 1] - E[A_i | B_i = 0]$ be average (associational) effect of B_i on A_i . Note that, like expectation, $\Delta(\cdot | B_i)$ is linear in the first argument. By Lemma 1 and the fact that both the instruments are binary, it is possible to express this coefficient from a partitioned regression approach:

$$\delta_1 = \frac{\Delta(Y_i^* | Z_{i1})}{\Delta(D_{i1}^* | Z_{i1})} = \frac{E[Y_i^* | Z_{i1} = 1] - E[Y_i^* | Z_{i1} = 0]}{E[D_{i1}^* | Z_{i1} = 1] - E[D_{i1}^* | Z_{i1} = 0]}, \quad (44)$$

where,

$$Y_i^* = Y_i - \beta_2 D_{i2} - \beta_3 D_{i3}, \quad (45)$$

$$D_{i1}^* = D_{i1} - \alpha_2 D_{i2} - \alpha_3 D_{i3}. \quad (46)$$

Here, the coefficients β_2 and β_3 are the coefficients on D_{i2} and D_{i3} from a TSLS regression of Y_i on these two variables, instrumented by Z_{i2} and Z_{i3} . These quantities also need to be derived and can be with an additional application of Lemma 1, below. The quantities β_2 and α_2 will not contribute to δ_1 because, due to independence of the treatments and treatment exclusion, $\Delta(D_{i2} | Z_{i1}) = E[D_{i2} | Z_{i1} = 1] - E[D_{i2} | Z_{i1} = 0] = 0$. It is helpful to state a handful of intermediate identification results. I omit the proofs of these statements, but they follow easily from the

assumptions of the theorem and sometimes the law of iterated expectations:

$$\Delta(D_{i1}|Z_{i1}) = \rho_{cc} + \rho_{ca} + \rho_{cn} \quad (47)$$

$$\Delta(D_{i1}|Z_{i2}) = 0 \quad (48)$$

$$\Delta(D_{i1}|Z_{i3}) = \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) (\rho_{cc} + \rho_{ca} + \rho_{cn}) \quad (49)$$

$$\Delta(D_{i2}|Z_{i1}) = 0 \quad (50)$$

$$\Delta(D_{i2}|Z_{i2}) = \rho_{cc} + \rho_{ac} + \rho_{nc} \quad (51)$$

$$\Delta(D_{i2}|Z_{i3}) = \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) (\rho_{cc} + \rho_{ac} + \rho_{nc}) \quad (52)$$

$$\Delta(D_{i3}|Z_{i1}) = (\pi_{01} + \pi_{11}) \rho_{cc} + \rho_{ca} \quad (53)$$

$$\Delta(D_{i3}|Z_{i2}) = (\pi_{10} + \pi_{11}) \rho_{cc} + \rho_{ac} \quad (54)$$

$$\Delta(D_{i3}|Z_{i3}) = \rho_{cc} + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) \rho_{ca} + \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) \rho_{ac} \quad (55)$$

$$\begin{aligned} \Delta(Y_i|Z_{i1}) &= (\pi_{01} + \pi_{11}) \tau_{cc,1|1} \rho_{cc} + (\pi_{00} + \pi_{10}) \tau_{cc,1|0} \rho_{cc} \\ &\quad + \tau_{ca} \rho_{ca} + \tau_{cn} \rho_{cn} \end{aligned} \quad (56)$$

$$\begin{aligned} \Delta(Y_i|Z_{i2}) &= (\pi_{10} + \pi_{11}) \tau_{cc,2|0} \rho_{cc} + (\pi_{00} + \pi_{01}) \tau_{cc,2|0} \rho_{cc} \\ &\quad + \tau_{ac} \rho_{ac} + \tau_{nc} \rho_{nc} \end{aligned} \quad (57)$$

$$\begin{aligned} \Delta(Y_i|Z_{i3}) &= \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) [\tau_{ac} \rho_{ac} + \tau_{nc} \rho_{nc}] \\ &\quad + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) [\tau_{ca} \rho_{ca} + \tau_{cn} \rho_{cn}] \\ &\quad + \frac{1}{1 - \pi_{11}} [\pi_{10} \tau_{cc,2|1} \rho_{cc} + \pi_{01} \tau_{cc,1|1} \rho_{cc} + \pi_{00} \tau_{LAJE} \rho_{cc}] \end{aligned} \quad (58)$$

Here, I use τ_{jk} for those compliance classes where there is only one effect defined, such as $\tau_{ca} = \tau_{ca,1}(1)$. To further save space, for the joint complier effects, I write $\tau_{cc,1|d_2} = \tau_{cc,1}(d_2)$.

To derive δ_1 , I begin with the denominator, which first requires the derivation of α_3 . By Lemma 1, this coefficient can be written as:

$$\alpha_3 = \frac{\Delta(D_{i1} - \alpha_{12} D_{i2} | Z_{i3})}{\Delta(D_{i3} - \alpha_{32} D_{i2} | Z_{i3})} \quad (59)$$

Here, the α_{ab} is the coefficient from the TSLS regression of D_{ia} on D_{ib} , using Z_{ib} as a instrument. This quantity relies on two others, which can be derived from the

above intermediate results as

$$\alpha_{1|2} = \frac{\Delta(D_{i1}|Z_{i2})}{\Delta(D_{i2}|Z_{i2})} = 0, \quad (60)$$

and

$$\alpha_{3|2} = \frac{\Delta(D_{i3}|Z_{i2})}{\Delta(D_{i2}|Z_{i2})} = \frac{(\pi_{10} + \pi_{11})\rho_{cc} + \rho_{ac}}{\rho_{cc} + \rho_{ac} + \rho_{nc}}. \quad (61)$$

Given these results, the numerator of α_3 can be written as:

$$\Delta(D_{i1} - \alpha_{1|2}D_{i2}|Z_{i3}) = \Delta(D_{i1}|Z_{i3}) \quad (62)$$

$$= \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) (\rho_{cc} + \rho_{ca} + \rho_{cn}) \quad (63)$$

The denominator of α_3 can also be derived from the above results:

$$\begin{aligned} \Delta(D_{i3} - \alpha_{3|2}D_{i2}|Z_{i3}) &= \Delta(D_{i3}|Z_{i3}) - \alpha_{3|2}\Delta(D_{i2}|Z_{i3}) \\ &= \rho_{cc} + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) \rho_{ca} + \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) \rho_{ac} \\ &\quad - \frac{(\pi_{10} + \pi_{11})\rho_{cc} + \rho_{ac}}{(\rho_{cc} + \rho_{ac} + \rho_{nc})} \left[\left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) (\rho_{cc} + \rho_{ac} + \rho_{nc}) \right] \\ &= \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) [\rho_{cc} + \rho_{ca}] \end{aligned}$$

Here I used independence of the instruments which implies $(\pi_{10} + \pi_{11})(\pi_{10} + \pi_{00}) = \Pr[Z_{i1} = 1] \Pr[Z_{i2} = 0] = \pi_{10}$

This gives us $\alpha_3 = (\rho_{cc} + \rho_{ca} + \rho_{cn})/(\rho_{cc} + \rho_{ca})$. With this, it is possible to derive the denominator of δ_1 :

$$\Delta(D_{i1}^*|Z_{i1}) = \Delta(D_{i1}|Z_{i1}) - \alpha_2\Delta(D_{i2}|Z_{i1}) - \alpha_3\Delta(D_{i3}|Z_{i1}) \quad (64)$$

$$= (\rho_{cc} + \rho_{ca} + \rho_{cn}) - \frac{(\rho_{cc} + \rho_{ca} + \rho_{cn})}{(\rho_{cc} + \rho_{ca})} [(\pi_{01} + \pi_{11})\rho_{cc} + \rho_{ca}] \quad (65)$$

$$= (\rho_{cc} + \rho_{ca} + \rho_{cn}) \left(1 - \frac{(\pi_{01} + \pi_{11})\rho_{cc} + \rho_{ca}}{(\rho_{cc} + \rho_{ca})} \right) \quad (66)$$

$$= (\rho_{cc} + \rho_{ca} + \rho_{cn})(\pi_{10} + \pi_{00}) \left(\frac{\rho_{cc}}{\rho_{cc} + \rho_{ca}} \right) \quad (67)$$

Moving on to the numerator of δ_1 , we now need to derive β_3 , which by Lemma 1 can be written as:

$$\beta_3 = \frac{\Delta(Y_i - \beta_{y|2}D_{i2}|Z_{i3})}{\Delta(D_{i3} - \alpha_{3|2}D_{i2}|Z_{i3})} \quad (68)$$

Here, $\alpha_{3|2}$ was already derived above so the only quantity needed is $\beta_{y|2}$, which is the coefficient on D_{i2} from a TSLS regression of Y_i on D_{i2} using Z_{i2} as instrument and by Lemma 1 is

$$\begin{aligned}\beta_{y|2} &= \frac{\Delta(Y_i|Z_{i2})}{\Delta(D_{i2}|Z_{i2})} \\ &= \frac{(\pi_{10} + \pi_{11})\tau_{cc,2}(1)\rho_{cc} + (\pi_{00} + \pi_{01})\tau_{cc,2}(0)\rho_{cc} + \tau_{ac}\rho_{ac} + \tau_{nc}\rho_{nc}}{\rho_{cc} + \rho_{ac} + \rho_{nc}}.\end{aligned}$$

With this, the numerator of β_3 can be written as:

$$\begin{aligned}\Delta(Y_{i1} - \beta_{y|2}D_{i2}|Z_{i3}) &= \Delta(Y_i|Z_{i3}) - \beta_{y|2}\Delta(D_{i2}|Z_{i3}) \\ &= \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}}\right) [\tau_{ac,2}(1)\rho_{ac} + \tau_{nc,2}\rho_{nc}] + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}}\right) [\tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}] \\ &\quad + \frac{1}{1 - \pi_{11}} [\pi_{10}\tau_{cc,2|1}\rho_{cc} + \pi_{01}\tau_{cc,1|1}\rho_{cc} + \pi_{00}\tau_{LAJE}\rho_{cc}] \\ &\quad - \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}}\right) [(\pi_{10} + \pi_{11})\tau_{cc,2|1}\rho_{cc} + (\pi_{00} + \pi_{01})\tau_{cc,2|0}\rho_{cc} + \tau_{ac}\rho_{ac} + \tau_{nc}\rho_{nc}] \\ &= \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}}\right) [\tau_{cc,1|1}\rho_{cc} + \tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}]\end{aligned}$$

The third equality uses the fact that $\tau_{LAJE} - \tau_{cc,2|0} = \tau_{cc,1|1}$. Combining this with the results on the denominator (which is the same as for α_3), we find that:

$$\beta_3 = \frac{\tau_{cc,1|1}\rho_{cc} + \tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}}{\rho_{cc} + \rho_{ca}} \quad (69)$$

Finally, we can combine these results to derive the numerator of δ_1 :

$$\begin{aligned}\Delta(Y_i^*|Z_{i1}) &= \Delta(Y_i|Z_{i1}) - \beta_2\Delta(D_{i2}|Z_{i1}) - \beta_3\Delta(D_{i3}|Z_{i1}) \\ &= (\pi_{01} + \pi_{11})\tau_{cc,1|1}\rho_{cc} + (\pi_{00} + \pi_{10})\tau_{cc,1|0}\rho_{cc} + \tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn} \\ &\quad - \frac{\tau_{cc,1|1}\rho_{cc} + \tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}}{\rho_{cc} + \rho_{ca}} [(\pi_{01} + \pi_{11})\rho_{cc} + \rho_{ca}] \\ &= \left(\frac{\rho_{cc}}{\rho_{cc} + \rho_{ca}}\right) (\pi_{10} + \pi_{00}) \left[\frac{\rho_{ca}}{\rho_{cc}}\tau_{cc,1|0}\rho_{cc} + \tau_{cc,1|0}\rho_{cc} + \tau_{cn}\rho_{cn} + \tau_{ca}\rho_{ca} - \frac{\rho_{ca}}{\rho_{cc}}\tau_{cc,1|1}\rho_{cc} \right] \\ &= \left(\frac{\rho_{cc}}{\rho_{cc} + \rho_{ca}}\right) (\pi_{10} + \pi_{00}) [(\rho_{cc}\tau_{cc,1|0} + \rho_{cn}\tau_{cn} + \rho_{ca}(\tau_{ca} - (\tau_{cc,1|1} - \tau_{cc,1|0})))]\end{aligned}$$

Combining the numerator and denominator of δ_1 gives the desired result:

$$\begin{aligned}\delta_1 &= \left(\frac{\rho_{cc}}{\rho_{cc} + \rho_{ca} + \rho_{cn}}\right) \tau_{cc,1|0} + \left(\frac{\rho_{cn}}{\rho_{cc} + \rho_{ca} + \rho_{cn}}\right) \tau_{cn} \\ &\quad + \left(\frac{\rho_{ca}}{\rho_{cc} + \rho_{ca} + \rho_{cn}}\right) [\tau_{ca} - (\tau_{cc,1|1} - \tau_{cc,1|0})]\end{aligned} \quad (70)$$

By symmetry, the coefficient on D_{i2} in this TSLS setup is:

$$\begin{aligned} \delta_2 = & \left(\frac{\rho_{cc}}{\rho_{cc} + \rho_{ac} + \rho_{nc}} \right) \tau_{cc,2|0} + \left(\frac{\rho_{nc}}{\rho_{cc} + \rho_{ac} + \rho_{nc}} \right) \tau_{nc} \\ & + \left(\frac{\rho_{ac}}{\rho_{cc} + \rho_{ac} + \rho_{nc}} \right) [\tau_{ac} - (\tau_{cc,2|1} - \tau_{cc,2|0})] \end{aligned} \quad (71)$$

To derive the coefficient on the interaction term, I use a similar approach. By Lemma 1, we have

$$\delta_3 = \frac{\Delta(\tilde{Y}_i^*|Z_{i3})}{\Delta(\tilde{D}_{i3}^*|Z_{i3})}, \quad (72)$$

where,

$$\tilde{Y}_i^* = Y_i - \tilde{\beta}_1 D_{i2} - \tilde{\beta}_2 D_{i2}, \quad (73)$$

$$\tilde{D}_{i3}^* = D_{i3} - \tilde{\alpha}_1 D_{i2} - \tilde{\alpha}_2 D_{i2}. \quad (74)$$

Here, $\tilde{\beta}_1$ and $\tilde{\beta}_2$ are the coefficients from a TSLS regression of Y_i on both treatments and both instruments without any interactions. Given treatment exclusion and independence of the treatments, here $\tilde{\beta}_1 = \tau_{LATE,1}$ and $\tilde{\beta}_2 = \tau_{LATE,2}$. Using Lemma 1 again, we can derive the coefficients in the numerator as:

$$\begin{aligned} \tilde{\alpha}_1 &= \frac{\Delta(D_{i3} - \tilde{\alpha}_{3|2} D_{i2} | Z_{i1})}{\Delta(D_{i1} - \tilde{\alpha}_{1|2} D_{i2} | Z_{i1})} = \frac{\Delta(D_{i3} | Z_{i1})}{\Delta(D_{i1} | Z_{i1})} \\ &= \frac{(\pi_{01} + \pi_{11})\rho_{cc} + \rho_{ca}}{\rho_{cc} + \rho_{ca} + \rho_{cn}} \end{aligned} \quad (75)$$

$$\tilde{\alpha}_2 = \frac{(\pi_{10} + \pi_{11})\rho_{cc} + \rho_{ac}}{\rho_{cc} + \rho_{ac} + \rho_{nc}} \quad (76)$$

To derive the denominator of δ_3 , I expand it and

$$\Delta(\tilde{D}_{i3}^*|Z_{i3}) = \Delta(D_{i3}|Z_{i3}) - \tilde{\alpha}_1 \Delta(D_{i1}|Z_{i3}) - \tilde{\alpha}_2 \Delta(D_{i2}|Z_{i3}) \quad (77)$$

$$\begin{aligned} &= \rho_{cc} + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) \rho_{ca} + \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) \rho_{ac} \\ &\quad - \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) ((\pi_{01} + \pi_{11})\rho_{cc} + \rho_{ca}) \end{aligned} \quad (78)$$

$$\begin{aligned} &\quad - \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) ((\pi_{10} + \pi_{11})\rho_{cc} + \rho_{ac}) \\ &= \frac{\pi_{00}}{1 - \pi_{11}} \rho_{cc} \end{aligned} \quad (79)$$

Thus, the denominator of δ_3 is the probability of joint compliance scaled.

For the numerator of δ_3 , I take a similar strategy:

$$\begin{aligned}
\Delta(\tilde{Y}_i^*|Z_{i3}) &= \Delta(Y_i|Z_{i3}) - \tilde{\beta}_1 \Delta(D_{i1}|Z_{i3}) - \tilde{\beta}_2 \Delta(D_{i2}|Z_{i3}) & (80) \\
&= \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) [\tau_{ac}\rho_{ac} + \tau_{nc}\rho_{nc}] + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) [\tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}] \\
&\quad + \frac{1}{1 - \pi_{11}} [\pi_{10}\tau_{cc,2|1}\rho_{cc} + \pi_{01}\tau_{cc,1|1}\rho_{cc} + \pi_{00}\tau_{LAJE}\rho_{cc}] \\
&\quad - \frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} ((\pi_{01} + \pi_{11})\tau_{cc,1|1}\rho_{cc} + (\pi_{00} + \pi_{10})\tau_{cc,1|0}\rho_{cc} + \tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}) \\
&\quad - \frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} ((\pi_{10} + \pi_{11})\tau_{cc,2|1}\rho_{cc} + (\pi_{00} + \pi_{01})\tau_{cc,2|0}\rho_{cc} + \tau_{ac}\rho_{ac} + \tau_{nc}\rho_{nc}) & (81)
\end{aligned}$$

$$\begin{aligned}
&= \left(\frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} \right) [\tau_{ac}\rho_{ac} + \tau_{nc}\rho_{nc}] + \left(\frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} \right) [\tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}] \\
&\quad + \frac{1}{1 - \pi_{11}} [\pi_{10}\tau_{cc,2|1}\rho_{cc} + \pi_{01}\tau_{cc,1|1}\rho_{cc} + \pi_{00}\tau_{LAJE}\rho_{cc}] & (82) \\
&\quad - \frac{\pi_{01}}{1 - \pi_{11}} \tau_{cc,1|1}\rho_{cc} - \frac{\pi_{00}}{1 - \pi_{11}} \tau_{cc,1|0}\rho_{cc} - \frac{\pi_{01} + \pi_{00}}{1 - \pi_{11}} [\tau_{ca}\rho_{ca} + \tau_{cn}\rho_{cn}] \\
&\quad - \frac{\pi_{10}}{1 - \pi_{11}} \tau_{cc,2|1}\rho_{cc} - \frac{\pi_{00}}{1 - \pi_{11}} \tau_{cc,2|0}\rho_{cc} - \frac{\pi_{10} + \pi_{00}}{1 - \pi_{11}} [\tau_{ac}\rho_{ac} + \tau_{nc}\rho_{nc}] \\
&= \frac{\pi_{00}}{1 - \pi_{11}} [\tau_{LAJE}\rho_{cc} - \tau_{cc,1|0}\rho_{cc} - \tau_{cc,2|0}\rho_{cc}] & (83)
\end{aligned}$$

$$= \left(\frac{\pi_{00}}{1 - \pi_{11}} \right) \rho_{cc} [\tau_{cc,1|1} - \tau_{cc,1|0}] & (84)$$

The third equality here used independence of the instruments.

Finally, we have,

$$\delta_3 = \tau_{cc,1|1} - \tau_{cc,1|0} = \tau_{LAIE}, & (85)$$

which finishes the proof. \square