NEJM Evidence

Published March 25, 2025
NEJM Evid 2025;4(4)
DOI: 10.1056/EVIDctw2400204

CLINICAL TRIALS WORKSHOP | INTERFACE SERIES

# Instrumental Variables in Randomized Trials

Joshua D. Angrist, Ph.D.,[1] Carol Gao, B.A.,[2] Peter Hull, Ph.D.,[3] and Robert W. Yeh, M.D.[4]

## Abstract

Many randomized clinical trials fail to play out as intended: some participants assigned to the treatment group remain untreated, while others assigned to the control group cross over and receive treatment. In such settings, intention-to-treat analyses that compare participants by treatment assignment are diluted by noncompliance, while per-protocol analyses that compare participants by treatment received are contaminated by selection bias. Instrumental variables methods can address both problems. We explain the rationale for instrumental variables estimation in clinical trials and illustrate instrumental variables methods through an analysis of the effect of revascularization on quality of life. We argue that instrumental variables analysis should be central to pragmatic trials of all kinds, strategy trials in particular, and emerging "nudge trials" that encourage specific health-related behaviors in large populations.

## Introduction

Randomized clinical trials (RCTs) seek to provide unbiased estimates of causal effects for a wide range of medical interventions. By randomly selecting a group of participants to receive treatment, an RCT is designed to yield balanced comparisons. In many trials, however, treatment is not delivered as intended: some participants randomly assigned to treatment opt out (a problem known as *nonadherence*), while other participants randomly assigned to control cross over to receive treatment. In such trials, the intended treatment is randomly assigned but treatment received is not. Trialists ignore nonadherence and crossovers at their peril: when trial participants deviate from random assignment, comparisons on the basis of treatment received are likely to generate biased estimates of treatment effects.

The trialist's first line of defense against such selection bias is to compare participants by the treatment they were assigned instead of the treatment they received. The resulting estimates capture intention-to-treat (ITT) effects, also called intention-to-screen effects in the case of screening trials. We refer to both as *intentions analysis*. Intentions analysis eliminates selection bias, but at a cost: ITT and intention-to-screen effects are diluted by noncompliance with treatment assignment, whether through nonadherence or crossovers. Consequently, intentions analysis provides a clouded picture of per-protocol effects. Moreover, even when therapeutic or screening effects are roughly constant across trial participants, intentions analysis generates estimates that will differ across trials when compliance rates differ.

The dilution and increased variability inherent in intentions analysis diminish the impact of — and benefits garnered from — costly RCTs. For example, intentions analyses have generated contested findings regarding the effectiveness of revascularization in the recent International Study of Comparative Health Effectiveness with Medical and Invasive

*The author affiliations are listed at the end of the article.*

*Dr. Yeh can be contacted at ryeh@bidmc.harvard.edu or at Division of Cardiovascular Medicine, Department of Medicine, Smith Center for Outcomes Research, Beth Israel Deaconess Medical Center, 375 Longwood Avenue, 4th Floor Boston, MA 02215.*

*This article was updated on April 3, 2025 at evidence.nejm.org.*

Approaches (ISCHEMIA) trial comparing conservative and invasive strategies for management of coronary artery disease.[1,2] Many ISCHEMIA participants assigned to conservative treatment crossed over to revascularization, while many assigned to invasive treatment were not revascularized as planned. Nonadherence in colorectal cancer screening trials has likewise led some observers to question the clinical relevance of screening trial findings.[3]

Instrumental variables methods attempt to rescue trials from the pitfalls of unfulfilled intentions. In the face of nonadherence and crossovers, instrumental variables methods yield easily computed and readily interpreted per-protocol effects. In particular, these methods estimate a local average treatment effect (LATE): the average causal effect for participants induced into treatment or screening by random assignment. This is a per-protocol effect for a subset of treated individuals, often the majority. Instrumental variables methods are especially well suited to pragmatic trials in everyday clinical settings, "strategy trials" like ISCHEMIA that assign flexible clinical pathways, and emerging "nudge trials."[4]

In a trial with no stratification or other deviations from simple random assignment, the simplest instrumental variables estimator divides an intentions effect (either ITT or intention-to-screen) by the treatment-group-versus-control-group difference in the share treated. In an instrumental variables analysis of an RCT with deviations from treatment assigned, a dummy variable indicating treatment assigned is said to be an instrument for treatment received. Instrumental variables methods presume that random assignment to the treatment group causes some participants to get treated, does not inhibit treatment for anyone, and affects outcomes solely by changing the likelihood of treatment receipt. These conditions — called *first-stage, monotonicity,* and *exclusion* — are plausible in most RCTs.

While instrumental variables analyses crop up occasionally in medical research, instrumental variables methods remain obscure to many medical professionals who might fruitfully apply them. This article aims to remedy this. We explain key instrumental variables assumptions, interpret instrumental variables estimates in clinical settings, and demonstrate the use of diagnostic tools. The next section on Understanding Clinical Instrumental Variables describes the rationale for instrumental variables analysis of data from clinical trials, while the section on Instrumental Variables in Action — Revascularization Effects illustrates instrumental variables methods by estimating the effect of revascularization on cardiac patients' health status. We conclude with some comments on expanding the use of instrumental variables in the clinical literature.

# Understanding Clinical Instrumental Variables

## INSTRUMENTAL VARIABLES THEORY

Consider a trial similar to the ISCHEMIA trial, which randomly assigned patients with coronary artery disease to an invasive revascularization strategy. Trial participants not assigned to revascularization are meant to be treated conservatively, with medical therapy alone. Let dummy variable $Z_i \in \{0, 1\}$ indicate treatment assignment to trial participants indexed by $i$, and let dummy variable $T_i \in \{0, 1\}$ indicate treatment received. For clinical or other reasons, some participants assigned to be revascularized remain unrevascularized, in which case $Z_i$=1 but $T_i$=0. At the same time, some assigned controls opt for revascularization, in which case $Z_i$=0 but $T_i$=1. In the ISCHEMIA trial, 12% of trial participants who were assigned to the conservative group were nevertheless revascularized the following year, while 20% of those assigned to the invasive group were not revascularized. Finally, let $Y_i$ be an outcome measured 1 year after assignment; in our ISCHEMIA analysis, this is the Seattle Angina Questionnaire (SAQ) health index and its components, measured for all participants, with scores ranging from 0 to 100. Higher scores indicate better health. The first column of Table 1 reports control means and standard deviations for the two SAQ outcomes — a quality of life score and an angina frequency score — analyzed here.

Treatment effects in this setting can be described with a simple causal model. The model posits the existence of potential outcomes $Y_i(1)$ and $Y_i(0)$, representing person $i$'s SAQ health score with and without the treatment, respectively. Both potential outcomes are defined for all participants. For instance, someone who would be in poor health when treated conservatively and better health when revascularized has $Y_i(1) > Y_i(0)$. The causal effect of treatment on this person equals $Y_i(1) - Y_i(0)$, a positive number. Someone else might have the same quality of life regardless of treatment, in which case their $Y_i(1) - Y_i(0)$ equals zero. Individual treatment effects necessarily remain hidden, since they contrast observed and counterfactual potential outcomes for the same person. Our goal, therefore, is to estimate average treatment effects for groups of trial participants.[5]

Random assignment seeks to ensure that participants assigned to invasive and conservative groups have the same average characteristics. Specifically, because random assignment makes treatment offers independent of

**Table 1. Estimated Effects of Revascularization on 1-Year SAQ Scores.***

| | Control Mean (1) | Reduced Form (ITT) (2) | First Stage (Compliance) (3) | Instrumental Variables Estimate (4) | Per-Protocol (As-Treated) Estimate (5) |
|---|---|---|---|---|---|
| SAQ quality of life score | 76.47 [23.20] | 3.98 (0.674) | 0.683 (0.011) | 5.83 (0.989) | 3.70 (0.680) |
| SAQ angina frequency score | 90.36 [15.94] | 3.69 (0.421) | — | 5.39 (0.616) | 3.95 (0.424) |

*This table reports estimated effects of revascularization on self-assessed Seattle Angina Questionnaire scores 1 year after random assignment. Treatment is defined as receiving revascularization. Column 1 reports mean scores and standard deviations for the group assigned to conservative care. Column 2 reports reduced-form (intention-to-treat) effects; column 3 reports first-stage effects of assignment on revascularization (compliance); column 4 reports instrumental variables estimates of local average treatment effects (LATE) of revascularization on compliers; computed using two-stage least squares; and column 5 reports corresponding as-treated estimates based on treatment received. All estimates are from models that control for baseline angina frequency and enrollment regions. Standard deviations in column 1 appear in brackets. Robust standard errors appear in parentheses. ITT denotes intention to treat; SAQ, Seattle Angina Questionnaire; SD, standard deviation; and SE, standard error.

potential outcomes, we have, for either treated ($t=1$) or untreated ($t=0$):

$$\begin{aligned}&\text{Avg. } Y_i(t) \text{ if assigned treatment}\\ &\quad = \text{Avg. } Y_i(t) \text{ if assigned control}\\ &\quad = \text{Avg. } Y_i(t) \text{ for everyone,}\end{aligned}$$

where by "Avg." we mean the average that we would see for groups of participants in a trial with a large number of participants (formally, this is a conditional expectation). As shorthand for this key fact, we say that random assignment *balances* potential outcomes.

Intentions analysis compares average outcomes conditional on treatment assignment, $Z_i$. When everyone is treated as intended, $T_i$=$Z_i$ for all $i$. In this case, the difference in average outcomes between participants assigned treatment and participants assigned control equals the average causal effect of treatment. To show this, we interpret assignment-based comparisons as follows:

$$\begin{aligned}&\textit{Intention to treat effect}\\ &= \text{Avg. } Y_i \text{ if assigned treatment}\\ &\qquad - \text{Avg. } Y_i \text{ if assigned control}\\ &= \text{Avg. } Y_i(1) \text{ if assigned treatment}\\ &\qquad - \text{Avg. } Y_i(0) \text{ if assigned control}\\ &= \text{Avg. } Y_i(1) \text{ for everyone}\\ &\qquad - \text{Avg. } Y_i(0) \text{ for everyone}\\ &= \text{Avg. } Y_i(1) - Y_i(0) \text{ for everyone.}\end{aligned} \tag{1}$$

The second equality here uses the fact that when $T_i$=$Z_i$, we see $Y_i(1)$ for everyone assigned treatment and we see $Y_i(0)$ for everyone assigned control. The third equality uses the fact that random assignment balances potential outcomes. The upshot is that in a trial where everyone is treated as intended, the difference in average outcomes conditional on treatment assigned equals the average causal effect of treatment.

In pragmatic and strategy trials, many who win the treatment-assignment lottery fail to collect their prize (especially when the prize is a time-consuming, risky, or unpleasant medical procedure). When $T_i \neq Z_i$ for some trial participants, intentions analysis no longer reveals the average effect of treatment. Rather, under typical trial conditions, ITT effects are diluted by including outcomes of some patients who did not receive treatment in the treatment group.

This dilution is easily seen when causal effects are constant. Suppose, for now, that the difference in outcome with treatment and without treatment is the same for everyone; call this difference $\beta$. Then:

$$Y_i = Y_i(0) + \beta \times T_i,$$

since $Y_i(1)$=$Y_i(0)$+$\beta$ for all participants. Substituting this quantity for $Y_i$ in the first line of expression (1), we have:

$$\begin{aligned}&\textit{ITT effect}\\ &= \text{Avg. } (Y_i(0) + \beta \times T_i \text{ if assigned treatment})\\ &\qquad - \text{Avg. } (Y_i(0) + \beta \times T_i \text{ if assigned control}).\end{aligned}$$

Because random assignment balances potential outcomes, average $Y_i(0)$ for participants assigned treatment and participants assigned control are the same. This implies:

$$\begin{aligned}&\textit{ITT effect}\\ &= \beta \times (\text{Avg. } T_i \text{ if assigned treatment})\\ &\qquad - \beta \times (\text{Avg. } T_i \text{ if assigned control})\\ &= \beta \times (\text{Avg. } T_i \text{ if assigned treatment}\\ &\qquad - \text{Avg. } T_i \text{ if assigned control})\\ &= \beta \times \textit{Compliance}.\end{aligned} \tag{2}$$

In this expression,

$$\begin{aligned} &\textit{Compliance} \\ &= \text{Avg. } T_i \text{ if assigned treatment} \\ &\qquad - \text{Avg. } T_i \text{ if assigned control} \\ &= \text{Share treated if assigned treatment} \\ &\qquad - \text{Share treated if assigned control.} \end{aligned}$$

In general, treatment rates among those assigned treatment exceed treatment rates among those assigned control. In a trial with nonadherence and crossovers, the difference between these rates — compliance — is a number between zero and one.

In an instrumental variables analysis, compliance is known as the *first stage*. When this is less than one, ITT effects are attenuated or diluted relative to the causal effect, $\beta$. Suppose, for instance, that half of participants assigned treatment remain untreated while 10% assigned control get treated anyway. In this case, the first stage is 0.5−0.1=0.4. Differences in mean outcomes between those assigned treatment and those assigned control must therefore be driven by the 40% of patients induced to treatment by random assignment. For everyone else, treatment is unchanged by random assignment, so the causal effect of assignment on these participants is zero.

The fact that ITT effects equal $\beta$ times compliance allows us to solve the dilution problem. Specifically, we obtain $\beta$ by dividing:

$$\begin{aligned} &\frac{\textit{ITT effect}}{\textit{Compliance}} \\ &= \frac{\text{Avg. } Y_i \text{ if assigned treatment} - \text{Avg. } Y_i \text{ if assigned control}}{\text{Share treated if assigned treatment} - \text{Share treated if assigned control}} \\ &= \beta. \end{aligned} \tag{3}$$

This ratio of differences in averages is the heart of instrumental variables analysis.

Econometricians call the difference in average outcomes by treatment assigned the *reduced form*. As we have seen, in an RCT, the reduced form is an ITT effect capturing the impact of being randomly assigned to treatment on outcomes. In practice, of course, clinical impact need not be constant. The LATE theorem[5,6] interprets instrumental variables estimates in a setting where $Y_i(1) - Y_i(0)$ differs from one person to another (in other words, when there are heterogeneous treatment effects). This theorem says that, under conditions likely to hold in an RCT, the ratio of the reduced form to the first stage reveals an average causal effect for the population of trial participants who comply with treatment assignment. That is:

$$\frac{\textit{ITT effect}}{\textit{Compliance}} = \text{Avg. } Y_i(1) - Y_i(0) \text{ for compliers (LATE).} \tag{4}$$

Trial compliers are participants who get treated when assigned treatment but do not get treated when assigned control. The LATE theorem says that in a world where treatment effects differ from one person to another, instrumental variables estimates can be interpreted as the average treatment effect for this group. The first stage is the share of the trial population who are compliers.

In addition to presuming random assignment to treatment groups, equation (4) hinges on three assumptions. The first-stage assumption requires treatment assignment to increase the share actually treated; compliance cannot be zero. Second, a monotonicity assumption precludes scenarios in which assignment to the treatment group perversely inhibits treatment receipt. The third assumption, called exclusion, holds when treatment assignment changes outcomes solely by changing treatment status.

The first two assumptions are typically uncontroversial in RCTs. In particular, in the ISCHEMIA trial, participants assigned to the treatment group were much more likely to be revascularized than participants assigned to the control group, generating positive compliance and a nonzero first stage. Moreover, assignment to the ISCHEMIA treatment group surely facilitated (rather than inhibited) revascularization, ensuring monotonicity. The question of exclusion is more subtle. In unblinded trials like ISCHEMIA, exclusion rules out scenarios in which treatment assignment is a revivifying morale-booster even for the untreated. It is hard to see such a scenario as relevant for quality of life outcomes related to angina and shortness of breath. Exclusion fails in some cancer screening trials if the invitation to screen for one type of cancer leads to unrelated treatments such as vaccination. This possibility highlights the importance of post-assignment data collection on a wide range of health behaviors, including those not targeted in the trial.

As suggested by expression (3), the simplest instrumental variables estimator is a ratio of differences in means. In an RCT with partial compliance, instrumental variables estimates can be computed by dividing estimated ITT effects (the instrumental variables reduced form) by compliance (the instrumental variables first stage). This calculation is illustrated in Figure 1 for the ISCHEMIA trial (detailed in Instrumental Variables in Action — Revascularization

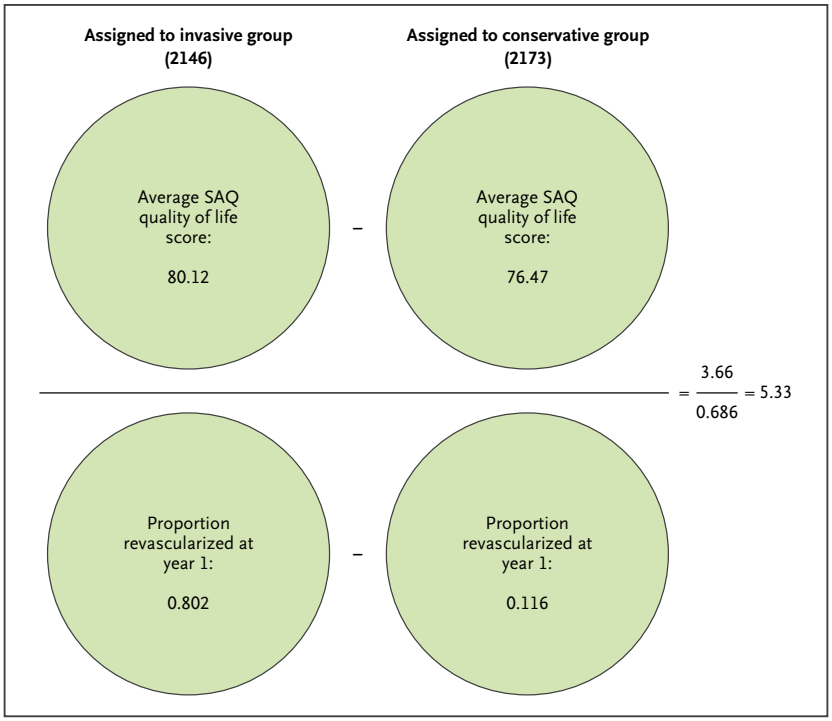


Figure 1. Simple Instrumental Variables Estimates of the Effect of Revascularization on Year 1 Seattle Angina Questionnaire Quality of Life Score.

This figure illustrates instrumental variables estimation of the effect of revascularization on quality of life. The numerator compares outcome means for trial participants assigned treatment (invasive group) and assigned control (conservative group), an intention-to-treat effect. This is divided by the difference in revascularization rates between the two groups. SAQ denotes Seattle Angina Questionnaire.

Effects, below). One year after random assignment, the group assigned treatment had an SAQ quality of life score about 3.7 points higher than those assigned control. Compliance equals 0.69, an estimate obtained by subtracting the share revascularized when assigned control (about 0.12) from the share revascularized when assigned treatment (about 0.8). The resulting estimate of the effects of revascularization on compliers equals about 5.3, almost 50% higher than the corresponding ITT effect.

While exhibits like Figure 1 are useful explanatory tools, in practice, instrumental variables estimates and the associated statistical tests are most easily computed using two-stage least squares (2SLS). 2SLS is a powerful and flexible instrumental variables estimator that allows for covariates such as dummies for randomization strata and measures of baseline health, and for which statistical inference is straightforward. 2SLS can also be used to estimate subgroup effects via interaction terms, and to efficiently combine and reconcile data from different trial sites or centers, allowing different compliance rates in each.[7,8] The instrumental variables estimates discussed in the following subsection were computed using 2SLS with controls for baseline angina frequency and enrollment region.

## INSTRUMENTAL VARIABLES FOR CLINICAL PRACTICE

In some trials, there are few crossovers from control to treatment, so noncompliance with treatment assigned reflects nonadherence in the treatment group only. In European colorectal cancer screening trials, for instance, adherence rates in the assigned-treatment group are as low as 50%, but few trial participants assigned control are routinely sent for colonoscopy or sigmoidoscopy. This contrasts with U.S. screening norms that prescribe colonoscopies for most people over 50 years of age. When noncompliance arises solely from nonadherence to treatment, LATE captures the average causal effect of screening or treatment on all who are screened or treated. For binary outcomes, the reciprocal of LATE can then be interpreted

as the number needed to screen to prevent one case for all screened patients.[8]

The clinical relevance of instrumental variables estimates is illuminated by a description of the relevant complier population. At first blush, this seems challenging. With crossovers as well as nonadherence, (as in ISCHEMIA) it is impossible to identify individual compliers, since we cannot know whether people assigned $Z_i$=1 would have been treated had they been assigned $Z_i$=0 (and vice versa).[9] How, then, can an analyst be confident that instrumental variables estimates are clinically relevant for a broad population?

Fortunately, complier characteristics are easily computed and compared with the characteristics of other populations of interest. To see how this works, let $X_i$ denote a characteristic of individual $i$, such as age or baseline health. In trials without randomization strata, complier mean $X_i$ can be obtained by replacing the ITT effect on outcome $Y_i$ in expression (4) with an ITT effect on $T_i \times X_i$:

$$\frac{(\text{Avg. } T_i \times X_i \text{ if assigned treatment} - \text{Avg. } T_i \times X_i \text{ if assigned control})}{\text{Compliance}} = \text{Avg. } X_i \text{ for compliers.} \quad (5)$$

This formula (derived in Angrist and Hull[8]) can be used to compute and contrast complier means with mean $X_i$ for all treated participants or some other population. These comparisons reveal, for instance, which demographic groups are well represented among compliers and the extent to which compliers' baseline health is similar to that of other populations of interest.

It is also instructive to contrast the instrumental variables approach with widely reported as-treated and per-protocol estimates that compare participants conditional on screening or treatment received (discussed, for instance, in the Catheter Ablation versus Antiarrhythmic Drug Therapy for Atrial Fibrillation [CABANA] trial[10]). An as-treated analysis compares participants grouped by treatment received ($T_i$), effectively discarding information on random assignment ($Z_i$); traditional per-protocol analysis retains only observations for which treatment received matches treatment assigned. Since $T_i$ is not randomly assigned, both comparisons are subject to the sort of selection bias that confounds nonrandomized cohort studies. By comparing individuals on the basis of randomly assigned $Z_i$, instrumental variables methods avoid selection bias while still estimating a per-protocol effect.

Figure 2 contrasts alternative analysis strategies for a trial in which treatment received deviates from treatment assigned, highlighting the shared apples-to-apples nature of the comparisons that underpin instrumental variables and ITT analyses.

## Instrumental Variables in Action — Revascularization Effects

The ISCHEMIA trial randomly assigned 5179 patients with moderate-to-severe cardiac ischemia to one of two care strategies. Patients assigned to the invasive group (assigned treatment) underwent diagnostic coronary angiography and subsequent revascularization when feasible — through angioplasty, stenting, or coronary artery bypass surgery — as well as medical therapy. Conservative group patients (assigned control) were to receive medical therapy alone with possible invasive treatment when medical therapy was deemed inadequate.[1,2]

In practice, many ISCHEMIA participants were not treated as intended. As noted in the discussion of Figure 1, most patients assigned treatment were revascularized (whether by angioplasty or bypass surgery), but 20% in this group received medical therapy alone. At the same time, 12% of those randomly assigned to control nevertheless underwent revascularization. Differences in clinical outcomes or quality of life between randomization groups can reasonably be attributed to differences in revascularization rates (ISCHEMIA patients were treated in the same hospitals regardless of treatment assigned, so participants in the two assignment groups likely received the same standard of care).

As can be seen in the first column of Table 1, average SAQ scores among those assigned to control range from around 76 for quality of life to 90 for angina frequency. Estimated ITT effects on these outcomes, reported in column 2 of the table, range from 3.7 to 4.0, all different from zero (standard errors for these estimates appear in parentheses; ITT estimates in the first row of the table differ from those in Figure 1 because the former control for covariates). Instrumental variables estimates for each of the two outcomes in the table, computed by 2SLS controlling for baseline angina frequency and enrollment region, exceed ITT estimates by 46%. This is a clinically important finding: the benefits of revascularization for trial participants who were revascularized as a result of the trial, on the order of 5.4 to 5.8, are markedly greater than previously reported and lie in a range considered clinically meaningful (usually >5 points on the SAQ scale). As is typical of instrumental variables analyses of RCTs, the statistical significance of LATE estimates roughly matches that of ITT estimates.

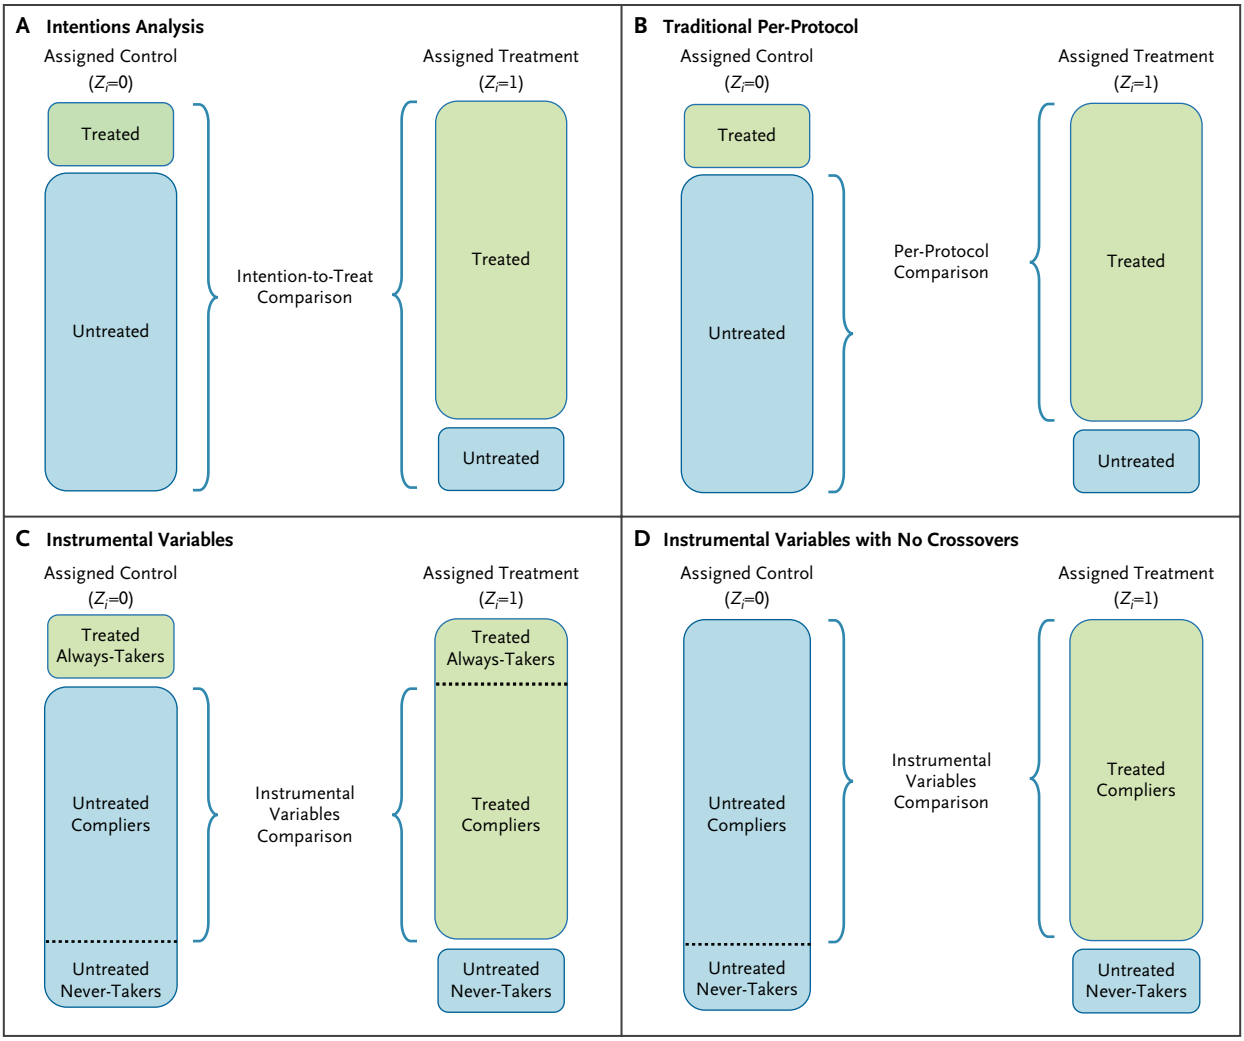


Figure 2. Alternative Within-Trial Comparisons.

This figure indicates the groups compared in alternative analysis strategies for causal effects in a randomized trial. Panels A, B, and C are drawn for a scenario in which the study population consists of two thirds compliers and one sixth each of always-takers and never-takers. In Panel D, there are no always-takers; all treated participants are compliers. The distribution of complier, always-taker, and never-taker status is independent of assigned treatment (indicated by $Z$), so blocks on each side of each panel are the same size. Intentions analysis, pictured in Panel A in the figure, compares all assigned treatment ($Z_i$=1) with all assigned control ($Z_i$=0). This comparison mixes treated and untreated on both sides, diluting the measure of the effects of treatment or screening. Per-protocol comparisons, described in Panel B, condition both on random assignment and the intervention chosen. Because treatment choices are nonrandom, groups on each side of the per-protocol comparison likely differ, leading to selection bias (confounding). Panels C and D picture the comparisons underpinning instrumental variables analysis. A subset of the treated are treated because they are compliers assigned $Z_i$=1. Likewise, a subset of controls remains untreated because they are compliers assigned $Z_i$=0. In general, compliers are trial participants who are treated as randomly assigned. Complier populations on both sides of an instrumental variables comparison are identical, differing only by whether they were randomly assigned to $Z_i$=1 or $Z_i$=0. Although not labeled as such in any data set, the instrumental variables formula finds this group. In Panels A, B, and C, one sixth of participants are *always-takers* who are treated regardless of assignment, while one sixth of participants are *never-takers* who remain untreated regardless of assignment. Compliance in this scenario therefore equals 5/6–1/6=2/3. In Panel D, which depicts a scenario with no crossovers, compliance equals 5/6–0=5/6. With no always-takers, everyone who is treated is a complier.

It is noteworthy that complier baseline health measures are virtually indistinguishable from those of the full sample of trial participants. This can be seen in the first two columns of Table 2 (the complier means in column 2 are computed as described in the section on Instrumental Variables for Clinical Practice.) The fact that complier baseline health

**Table 2. Baseline Sample Characteristics at 1-Year Follow-Up.***

| Characteristic | Sample | Compliers | Always-Takers | Never-Takers |
|---|---|---|---|---|
| SAQ quality of life score | 61.9 | 61.8 (0.79) | 49.4 (1.71) | 64.7 (1.33) |
| SAQ angina frequency score | 81.5 | 81.2 (0.62) | 71.8 (1.45) | 84.6 (0.87) |
| Female | 0.23 | 0.22 (0.012) | 0.21 (0.026) | 0.30 (0.022) |
| Black† | 0.04 | 0.03 (0.006) | 0.05 (0.013) | 0.05 (0.011) |
| White | 0.73 | 0.72 (0.012) | 0.74 (0.028) | 0.75 (0.021) |
| Age at random assignment | 64.3 | 63.8 (0.27) | 63.9 (0.57) | 66.7 (0.45) |
| Population share | | 0.68 | 0.12 | 0.20 |
| Share among treated | | 0.74 | 0.26 | |
| Share among nontreated | | 0.63 | | 0.37 |

* This table compares baseline health for compliers, always-takers, and never-takers, the share of always-takers among those treated, and the share of never-takers among those untreated. Complier means are computed using two-stage least squares, instrumenting $T_i \times X_i$ with treatment assigned, as described by expression (5) in the text. Estimates control for enrollment regions. Overall, 50% of patients were offered treatment, and 46% of patients received treatment. Robust standard errors appear in parentheses. SAQ denotes Seattle Angina Questionnaire.

† Race is recorded as Black, White or Other in the trial data. The race characteristics reported in Table 2 are binary.

matches that of the full trial sample bolsters the case for seeing instrumental variables estimates as clinically relevant for the population at large.

Table 2 also compares the baseline health of *always-takers* (those assigned to control who were nevertheless revascularized) and *never-takers* (those assigned to treatment who were not revascularized). Evidently, always-takers have baseline health well below that of compliers, while never-takers are a little healthier. This helps explain why as-treated estimates reported in the last column of Table 1 are below the corresponding set of instrumental variables estimates. As-treated estimates compare all treated with all untreated, and 26% of the treated are always-takers while 37% of the untreated are never-takers (statistics reported in the bottom rows of Table 2). As baseline characteristics summarized in the table show, these facts confound as-treated comparisons since always-takers (necessarily treated) enter the trial sicker than never-takers (necessarily untreated).

## Conclusion

A 2021 report[11] on the influential CABANA trial commented:

> The CABANA clinical trial...was difficult to interpret because of nonadherence with the treatment protocol that resulted from substantial crossover between groups.

Instrumental variables methods solve this problem, offering a clear path from randomized intentions to the impact of treatment or screening itself. In an RCT with nonadherence but no control group crossovers, instrumental variables estimation captures an average causal effect for the entire population that is treated or screened. Otherwise, instrumental variables identify effects on the population of protocol compliers — that is, the group induced into treatment or screening by the trial. This effect, along with the associated reduced-form and first-stage estimates, should be central to any report of trial findings.

The ISCHEMIA trial analyzed here is a leading example of a strategy trial. These and other similarly structured trials (such as pragmatic cancer screening trials) fit the instrumental variables template well. Nudge trials (sometimes called "encouragement trials") may also prove a fruitful domain for instrumental variables applications. A nudge trial aims to increase the likelihood of treatment by randomly assigning emails, texts, or letters to patients in a large population of clinical interest, without actively recruiting trial participants. Examples include trials encouraging older adults to get influenza vaccines[12,13] and a recent trial delivering nudges intended to increase cardiovascular medication adherence.[14] The instrumental variables framework converts small nudge effects into possibly much larger vaccination or medication effects, revealing otherwise hidden health benefits for nudge compliers.

Instrumental variables methods come to clinical research by way of econometrics. Clinicians and medical researchers may therefore see these methods as complicated or mysterious. Yet, the instrumental variables framework is rooted in the concern with selection bias that epidemiology

and medicine share with econometrics. We hope this article demystifies instrumental variables and promotes wider clinical application of this powerful tool.

## Disclosures

Author disclosures are available at evidence.nejm.org.

This article was prepared using International Study of Comparative Health Effectiveness with Medical and Invasive Approaches (ISCHEMIA) Research Materials obtained from the National Heart, Lung, and Blood Institute (NHLBI) Biologic Specimen and Data Repository Information Coordinating Center and does not necessarily reflect the opinions or views of ISCHEMIA or the NHLBI.


## Acknowledgments

We thank the National Cancer Institute for access to data collected by the Prostate, Lung, Colorectal, and Ovarian Cancer Screening Trial.



## Author Affiliations

[1] Department of Economics and NBER, MIT, Cambridge, MA

[2] Operations Research Center, MIT, Cambridge, MA

[3] Department of Economics and NBER, Brown University, Providence, RI

[4] Division of Cardiovascular Medicine, Department of Medicine, Smith Center for Outcomes Research, Beth Israel Deaconess Medical Center, Boston


## References


1. Spertus JA, Jones PG, Maron DJ, et al. Health-status outcomes with invasive or conservative care in coronary disease. N Engl J Med 2020;382:1408-1419. DOI: 10.1056/NEJMoa1916370.
2. Maron DJ, Hochman JS, Reynolds HR, et al. Initial invasive or conservative strategy for stable coronary disease. N Engl J Med 2020;382:1395-1407. DOI: 10.1056/NEJMoa1915922.
3. Winawer SJ. Colonoscopy screening and colorectal cancer incidence and mortality. N Engl J Med 2023;388:376-376. DOI: 10.1056/NEJMc2215192.
4. Johansen ND, Modin D, Nealon J, et al. A pragmatic randomized feasibility trial of influenza vaccines. NEJM Evid 2023;2(2). DOI: 10.1056/EVIDoa2200206.
5. Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. J Am Stat Assoc 1996;91:444-455. DOI: 10.1080/01621459.1996.10476902.
6. Imbens GW, Angrist JD. Identification and estimation of local average treatment effects. Econometrica 1994;62:467-475. DOI: 10.2307/2951620.
7. Angrist JD, Pischke JS. Mostly harmless econometrics: an empiricist's companion. Princeton, NJ: Princeton University Press, 2009.
8. Angrist JD, Hull P. Instrumental variables methods reconcile intention-to-screen effects across pragmatic cancer screening trials. Proc Natl Acad Sci U S A 2023;120:e2311556120. DOI: 10.1073/pnas.2311556120.
9. Hernán MA, Robins JM. Per-protocol analyses of pragmatic trials. N Engl J Med 2017;377:1391-1398. DOI: 10.1056/NEJMsm1605385.
10. Packer DL, Mark DB, Robb RA, et al. Effect of catheter ablation vs antiarrhythmic drug therapy on mortality, stroke, bleeding, and cardiac arrest among patients with atrial fibrillation: the CABANA randomized clinical trial. JAMA 2019;321:1261-1274. DOI: 10.1001/jama.2019.0693.
11. Smith VA, Coffman CJ, Hudgens MG. Interpreting the results of intention-to-treat, per-protocol, and as-treated analyses of clinical trials. JAMA 2021;326:433-434. DOI: 10.1001/jama.2021.2825.
12. Hirano K, Imbens GW, Rubin DB, Zhou XH. Assessing the effect of an influenza vaccine in an encouragement design. Biostatistics 2000;1:69-88. DOI: 10.1093/biostatistics/1.1.69.
13. Johansen ND, Vaduganathan M, Bhatt A, et al. Electronic nudges to increase influenza vaccination uptake in Denmark: a nationwide, pragmatic, registry-based, randomised implementation trial. Lancet 2023;401:1103-1114. DOI: 10.1016/S0140-6736(23)00349-5.
14. Glasgow RE, Knoepke CE, Magid D, et al. The NUDGE trial pragmatic trial to enhance cardiovascular medication adherence: study protocol for a randomized controlled trial. Trials 2021;22:528. DOI: 10.1186/s13063-021-05453-9.