

NBER WORKING PAPER SERIES

INSTRUMENTAL VARIABLES WITH TIME-VARYING EXPOSURE:
NEW ESTIMATES OF REVASCULARIZATION EFFECTS ON QUALITY OF LIFE

Joshua Angrist
Bruno Ferman
Carol Gao
Peter Hull
Otavio L. Tecchio
Robert W. Yeh

Working Paper 33296
<http://www.nber.org/papers/w33296>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2024, Revised July 2026

We thank Michal Kolesar and Amanda Kowalski for helpful comments; Fatima Djalalova provided exceptional research assistance. This manuscript was prepared using ISCHEMIA Research Materials obtained from the NHLBI Biologic Specimen and Data Repository Information Coordinating Center and does not necessarily reflect the opinions or views of ISCHEMIA, the NHLBI, or the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w33296>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Joshua Angrist, Bruno Ferman, Carol Gao, Peter Hull, Otavio L. Tecchio, and Robert W. Yeh. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Instrumental Variables with Time-Varying Exposure: New Estimates of Revascularization Effects on Quality of Life

Joshua Angrist, Bruno Ferman, Carol Gao, Peter Hull, Otavio L. Tecchio, and Robert W. Yeh

NBER Working Paper No. 33296

December 2024, Revised July 2026

JEL No. C26, C32, I10

ABSTRACT

This paper develops instrumental variables (IV) estimators for dynamic causal effects in randomized trials with imperfect compliance. These methods are applied to a randomized trial that assigned patients with ischemic heart disease to either an invasive treatment arm centered on revascularization or a control group meant to receive non-invasive medical therapy. As is common in such “strategy trials,” many participants assigned to treatment remained untreated while many assigned to control crossed over into treatment. Protocol non-compliance causes ITT estimates to diverge from the effect of treatment received, while conventional per-protocol analyses that condition on treatment received are compromised by selection bias. Extending the static potential-outcomes IV framework, the methods here identify average causal effects of treatment for dynamic compliers, the set of trial participants who comply with trial protocol at different follow-up horizons. IV estimates of revascularization effects on compliers’ quality of life are markedly larger and more sustained than previously reported ITT and per-protocol estimates. We also show how to estimate average characteristics and marginal potential outcome means for dynamic compliers. These results are used to explain confounding in as-treated per-protocol estimates.

Joshua Angrist
Massachusetts Institute of Technology
Department of Economics
and NBER
angrist@mit.edu

Peter Hull
Brown University
Department of Economics
and NBER
peter_hull@brown.edu

Bruno Ferman
Sao Paulo School of Economics, FGV
bruno.ferman@fgv.br

Otavio L. Tecchio
Massachusetts Institute of Technology
Department of Economics
otecchio@mit.edu

Carol Gao
Massachusetts Institute of Technology
carolgao@mit.edu

Robert W. Yeh
Beth Israel Lahey Health
ryeh@bidmc.harvard.edu

1 Introduction

Many patients with ischemic heart disease undergo revascularization, an invasive strategy involving percutaneous coronary intervention (PCI) and/or coronary artery bypass grafting (CABG). These are invasive, resource-intensive, and potentially risky procedures. Alternatively, ischemia patients may be treated conservatively with a combination of lifestyle changes and drugs. Whether the benefits of revascularization outweigh the associated risks and costs remains controversial.

The ISCHEMIA trial aimed to assess the benefits of revascularization in a randomized clinical trial (RCT) contrasting an invasive strategy centered on revascularization with conservative treatment involving medical therapy alone. The intention-to-treat (ITT) estimates that are the focus of ISCHEMIA analyses to date compare trial participants by assigned treatment arm. These show no statistically significant difference in mortality, while ITT estimates of assignment effects on angina frequency and other Seattle Angina Questionnaire (SAQ) domain scores indicate modest and fading effects on quality of life and angina frequency [Maron et al., 2020; Spertus et al., 2020a,b].

These findings notwithstanding, there’s more to learn from the ISCHEMIA trial than has been previously reported. ITT estimates are affected by substantial treatment group non-adherence (i.e., some assigned to invasive treatment were not revascularized) along with control group crossovers (i.e., some assigned to conservative treatment were revascularized). In other words, assignment to the invasive treatment group boosted revascularization rates substantially, but not deterministically. Comparisons by treatment assigned reflect a mix of causal effects of revascularization for trial participants who comply with the trial protocol and null effects for participants whose revascularization status is unchanged by random assignment [Angrist and Hull, 2023]. Conventional per protocol estimators attempt to undo the resulting dilution by conditioning on treatment received. But treatment received is typically far from randomly assigned: patients and their doctors choose whether to revascularize, a choice determined in part by patient frailty and prognosis.

The ISCHEMIA trial ran for five years, a period during which trial participants assigned conservative medical therapy were free to cross over to revascularization. The timing of noncompliance varied across participants in both arms of the ISCHEMIA trial. Most assigned to revascularize did so shortly after random assignment. Some control-group crossovers were also revascularized shortly after random assignment, but many waited years before invasive treatment. Consequently, in each follow-up wave, control-group crossovers include a mix of recently-revascularized patients and patients who were revascularized several years ago. This time-varying exposure has implications for estimating dynamic treatment effects, which might emerge if revascularization benefits fade, and is common in strategy trials.¹

These problems are addressed here with the method of instrumental variables (IV), using random assignment as an instrument for time-varying exposure to generate new estimates of revascularization effects on quality-of-life. Our methodological discussion starts with a theoretical result motivated by compliance behavior seen in the ISCHEMIA data. This approach exploits the fact

¹For example, the recent British Heart Foundation SENIOR-RITA trial is affected by non-adherence and time-varying crossover rates of magnitudes similar to those in ISCHEMIA [Kunadian et al., 2024].

that random assignment induces immediate revascularization for participants assigned to invasive treatment, but does not otherwise change the distribution of revascularization exposure by much. Following Angrist, Santos and Tecchio [2025], we show this behavior is sufficient (when added to standard IV assumptions) to identify effects of immediate exposure averaged over all possible levels of counterfactual exposure.

We then turn to identification of incremental and cumulative average causal effects of time-varying exposure, for each possible exposure value, under a set of identifying assumptions well-suited to strategy trials like ISCHEMIA. This identification strategy exploits novel *wave ignorability* assumptions that leverage the repeated measurement of outcomes across multiple follow-up waves. Both strategies identify average causal effects of revascularization exposure for patients revascularized or revascularized sooner as a result of invasive assignment. Finally, we show how to estimate mean characteristics of variables like baseline SAQ scores for trial participants—labeled *dynamic compliers*—whose exposure is changed by random assignment. These theoretical results should be applicable in a wide range of settings. An appendix extends identification results for dynamic compliers to marginal distributions of potential outcomes.

Application of these new IV tools to the ISCHEMIA trial reveals quality of life gains from revascularization that are substantially greater and more sustained than previously reported. IV estimates suggest revascularization yields sustained gains in SAQ summary scores of around four points over five years, while revascularization boosts angina frequency scores by a sustained three points after an initial-year effect spike of 5-6 points. Conventional ITT estimates of revascularization effects on SAQ and angina scores, by contrast, fade to around two points over a five-year horizon. As-treated per-protocol estimates similarly fade. We show that the latter result is explained by the growing share of relatively sick control-group crossovers and the fact that treatment-group nonadherents are an increasingly healthy group.

Our work builds on a large literature in statistics, biostatistics, and epidemiology that develops estimators for causal effects of time-varying treatment exposure (including, for instance, Hernán and Robins [2020]; Robins [1994]; Robins and Hernán [2008]). Much of this literature relies on sequential conditional independence assumptions, such as those used in foundational g-methods and structural nested models [Robins, 1997, 1998, 2000, 2004] and related approaches [e.g., Luo, Ruberg and Qu, 2022; Olarte Parra, Daniel and Bartlett, 2023; Qu et al., 2020; Rytgaard, Gerds and van der Laan, 2022]. IV methods, by contrast, rely on the initial randomization of treatment offers along with exclusion and monotonicity assumptions likely to hold in RCTs. To these we add either a (testable) restriction on compliance patterns or an event-study-type restriction on average causal effects. Sequential conditional independence requires that, conditional on treatment and covariate history, subsequent treatment decisions are as good as randomized with respect to future potential outcomes. This seems unattractive in our setting, since decisions to discontinue or modify treatment over time are likely to depend on unmeasured and evolving patient characteristics (a point noted in a related context by Michael et al. [2024]).

The identification strategies developed here are also related to work using instrumental variables

to identify causal effects in clinical trials. In trials with a static treatment and imperfect compliance, initial treatment assignment is a natural instrument for treatment received [Angrist and Hull, 2023].² Identification of dynamic causal effects, however, typically requires stronger assumptions than those needed to identify a single static causal effect [Hernán and Robins, 2006; Hernán and Robins, 2017]. Our analysis differs in two ways from previous applications of IV to models with dynamic non-compliance. First, we require only a single static instrument (derived from initial random assignment) while earlier work typically relies on a time-varying instrument [see, e.g., Michael et al., 2024; Robins and Hernán, 2008]. Second, we consider point-in-time treatment effects for outcomes like quality-of-life scores while earlier work focuses on parametric survival models for absorbing states like death [e.g., Martinussen et al., 2017; Michiels, Vandebosch and Vansteelandt, 2024; Robins and Tsiatis, 1991; Shi et al., 2021; Yende-Zuma, Mwambi and Vansteelandt, 2019; Ying and Tchetgen Tchetgen, 2023].³

In a conceptual framework closely related to ours, Bowden et al. [2025] identifies dynamic treatment effects in the context of a parametric model with a single static instrument and time-varying nonadherence. In contrast with Bowden et al. [2025], we identify dynamic causal effects without parametric restrictions. Our analysis also differs from Bowden et al. [2025] in allowing for control-group crossovers, a signal feature of ISCHEMIA and many other RCTs with partial compliance.⁴ Also closely related is Rose and Shem-Tov [2021], which explores the identifying power of a single instrument when dynamic compliers are moved on the extensive margin only (i.e., from untreated to positive treatment). Our theoretical analysis begins with an analogous assumption in which dynamic compliers are moved from positive to maximal exposure. To our knowledge, this restriction is new to appear in IV analyses of clinical trials.⁵

Beyond ISCHEMIA, the identification strategies developed here are well-suited to applications in which a time-invariant instrument affects time-varying treatment exposure. Examples from economics include studies of dynamic causal effects of fertility on labor market outcomes (e.g., Bronars and Grogger [1994]; Lundborg, Plug and Rasmussen [2017, 2024]) and time-varying effects of job training programs, as in Ba et al. [2017]. Recent work using in-vitro fertilization as an instrument for childbirth obtains theoretical results related to our Theorem 2 [Bensnes, Huitfeldt and Leuven, 2023; Gallen et al., 2024] while Vohra and Goldin [2024] give partial identification results for the effect of going from the smallest (shortest) to the largest (or longest) treatment exposure. Wave ignorability, the link between cumulative and incremental average causal effects in

²IV tools have long been applied to randomized trials in social science and medicine, though sometimes without reference to IV. For example, Bloom [1984] and Newcombe [1988] derive adjustments for randomized trials with control-group crossovers while Baker and Lindeman [1994] use maximum likelihood to derive IV-type adjustments for nonadherence for Bernoulli outcomes. Baker and Lindeman [2024] characterize this shared intellectual history.

³Simulation studies comparing causal effect estimators applied to survival data from randomized trials include Morden et al. [2011] and Latimer et al. [2018]. Other frameworks for time-varying effects in randomized trials include the Bayesian principal stratification models developed in Lin, Ten Have and Elliott [2008] and Mattei et al. [2024].

⁴Appendix B.1 extends the Bowden et al. [2025] parametric model to accommodate control-group crossovers. This appendix details limitations of the Bowden et al. [2025] parametric setup.

⁵Ferman and Tecchio [2023] explores the identifying power of assumptions like wave ignorability for a variety of dynamic treatment effect models. Assumptions limiting the dependence of causal effects on calendar time also appear in the biostatistics literature (e.g., Jiang, Chen and Ding [2023]).

a potential outcomes model identified by IV, and tools for analyzing marginal potential outcomes seem useful in these and other settings.

The rest of this paper is organized as follows. The next section introduces the ISCHEMIA trial and describes the evolution of compliance behavior over the life of the trial. Section 3 develops new IV identification strategies, following a brief review of IV methods applied to clinical trials with a static Bernoulli treatment. This section also shows how to identify the average characteristics of dynamic compliers. Section 4 applies this theoretical framework to ISCHEMIA. Section 5 summarizes and discusses broader implications.

2 The ISCHEMIA Trial

The ISCHEMIA trial randomized 5,179 patients with moderate-to-severe cardiac ischemia to one of two care strategies. Patients assigned to the invasive treatment arm were meant to undergo diagnostic coronary angiography and subsequent revascularization when feasible (through PCI or CABG) while also receiving medical therapy. Conservative-arm patients were meant to receive medical therapy alone, with possible invasive treatment when medical therapy was deemed inadequate [Maron et al., 2020; Spertus et al., 2020b]. The trial followed patients for up to five years after random assignment.

Our analysis of ISCHEMIA is organized around four scientific goals or questions. The first is to estimate effects of revascularization exposure on SAQ summary scores and angina frequency one to five years after randomization. The second is to determine whether these effects are short-lived or persistent. The third is to compare dynamic IV estimates of cumulative exposure effects with ITT and conventional as-treated/per-protocol analyses. Fourth, we aim to estimate the baseline characteristics of dynamic compliers—the patient group whose revascularization exposure is shifted by random assignment.

These matters are analyzed using data collected over the five-year life of the trial. Sample sizes decrease over time since participants who enrolled in the trial later contribute fewer observations ahead of the last follow-up date in December 2018. Roughly 4,300 participants contribute to the analysis of quality of life in wave 1, a number that falls to 670 by wave 5. The proportion with follow-up data is similar in the two arms of the trial. Falling sample sizes over time are due primarily to censoring by trial end date, as can be seen in Supplement Table 2 in Spertus et al. [2020b].⁶

Because SAQ outcomes are defined only for living trial participants observed at follow-up, the IV estimands detailed in the next section are subject to the usual concerns arising from truncation by patient death. These concerns are mitigated, however, by earlier analyses of ISCHEMIA showing that mortality is unchanged by treatment assignment (e.g., Bangalore et al. [2020, 2022]; Maron et al. [2020]). Balanced survival does not guarantee the absence of survival-related selection bias, especially for IV estimands defined by latent compliance groups. Given balanced and high follow-up rates, however, adjustments based on randomized-arm mortality differences would be expected to

⁶Estimates that control for expected follow-up time are virtually indistinguishable.

have little effect on ITT (and hence IV) estimates. Our estimates can therefore be interpreted as causal effects among participants alive and observed at follow-up.

ISCHEMIA participants assigned to the invasive arm enjoyed a moderately higher quality of life as a result. Estimated ITT effects on SAQ summary scores and angina frequency scores, plotted in Figure 1 and reported in Table 1, show modest treatment-induced gains in each wave. These effects are estimated by differences in outcome means, with covariate adjustment.⁷ Estimated ITT effects on the summary score range from just above 3 in the first wave to 1.7 in wave 4; angina gains initially peak a little higher in the first wave and fall a little lower in wave 3. These estimates can be compared to outcome means of 86-94, with standard deviations of 13-15 (see Table 1).

The ISCHEMIA trial is characterized by high and variable rates of dynamic noncompliance with experimental treatment assignment, reported in Table 1. Revascularization rates in the control group increase from 12% in wave 1 to 29% in wave 5. At the same time, only 80% of those assigned invasive were initially revascularized. Revascularization rates in the group assigned invasive increased from 80% in wave 1 to 83% in wave 5. Importantly, the difference in the likelihood of revascularization by assignment status, reported in the third column of the table, is well below one. Moreover, this gap falls from a high of 68% in the first wave to 54% by wave 5. We show below that these high rates of nonadherence and crossover likely dilute the ITT estimates in Figure 1.

Changes in treatment take-up over time induce differences in revascularization exposure, defined as the number of years since revascularization. Figure 2 charts the distribution of exposure conditional on assignment for each wave. These plots indicate that most invasive-arm patients who were revascularized were revascularized shortly after random assignment: in each wave w , most of those assigned to treatment have 0 or w years of exposure. The next section shows how such immediate compliance helps identify average dynamic effects.

3 IV for Clinical Exposure

3.1 Conceptual Framework

Consider a setting like ISCHEMIA in which trial participants are randomized to either an invasive treatment like revascularization or a conservative alternative like medical therapy. Let $Z \in \{0, 1\}$ denote a dummy variable indicating invasive assignment. Quality of life outcomes are measured in each of $\bar{w}$ annual follow-up waves; Y_w denotes quality of life measured in wave $w \in \{1, 2, 3, \dots, \bar{w}\}$. Treatment is assigned some time ahead of wave 1.⁸

As we’ve seen, post-assignment revascularization exposure is time-varying: many participants assigned conservative were revascularized later, while some assigned invasive were never revascularized or revascularized years after random assignment. Let $T_w \in \{0, 1, 2, \dots, w\}$ denote *revascu-*

⁷Specifically, ITTs come from ordinary least squares regressions of outcomes on assigned treatment, controlling for baseline outcomes and enrollment regions to boost precision.

⁸ISCHEMIA collected data every few months in the first year and then semi-annually [Spertus et al., 2020b]. Our analysis is limited to relatively extensive annual follow-ups. Spertus et al. [2020b] uses a Bayesian framework to estimate high-frequency time-varying exposure effects on an ITT basis.

larization exposure, defined as the years a patient has lived since revascularization, as seen in wave w . Participants in any wave are either never revascularized ($T_w = 0$), revascularized shortly after assignment ($T_w = w$), revascularized in the year of observation ($T_w = 1$), or revascularized at some other time in between ($1 < T_w < w$).

Potential outcomes determine heterogeneous causal effects of revascularization in each wave. Specifically, let $Y_w(t)$ denote a participant's outcome in wave w when they've lived $T_w = t$ years since revascularization. Once a participant is revascularized, they're revascularized forever. Formally, this means that for all $w \in \{1, \dots, \bar{w} - 1\}$, $T_w = t$ implies $T_{w+1} = t + 1$. This treatment pattern mirrors a staggered adoption setup in event-study models.⁹ Measured in wave w , the incremental individual causal effect of revascularization exposure is $Y_w(t) - Y_w(t-1)$. This notation encompasses the difference in quality of life a patient experiences when revascularized this year versus last year. The cumulative effect of revascularization for $t \leq w$ years of exposure observed in wave w is $Y_w(t) - Y_w(0)$. The cumulative effect of revascularization exposure for a participant revascularized shortly after random assignment is $Y_w(w) - Y_w(0)$.

Just as potential outcomes are indexed against treatment received, it's useful to index potential treatment against treatment assigned. To that end, let $T_w(z)$ denote revascularization exposure in wave w when $Z = z$. This is defined for all trial participants, regardless of assignment. The effect of invasive assignment on revascularization exposure in wave w is therefore $T_w(1) - T_w(0)$.

Given an exclusion restriction, random assignment makes Z independent of potential outcomes and potential treatments. This *independence assumption* is formalized as:

Assumption 1 (Independence). *For each wave $w \in \{1, \dots, \bar{w}\}$, the random variables $Y_w(0), \dots, Y_w(w), T_w(1), T_w(0)$ are jointly independent of Z .*

The exclusion restriction in this case asserts that assignment to the invasive arm affects outcomes solely by increasing the likelihood of revascularization.¹⁰ This assumption is plausible in the IS-CHEMIA trial, since randomization to the invasive treatment likely had no direct effects on outcomes. Importantly, the exclusion restriction allows assignment to the invasive treatment to affect outcomes via the *timing* of revascularization.

As in Imbens and Angrist [1994] and Angrist and Imbens [1995], we assume that invasive assignment either induces revascularization, advances the date of revascularization, or leaves revascularization exposure unchanged. This is formalized as:

Assumption 2 (Monotonicity). *For each $w \in \{1, \dots, \bar{w}\}$, $T_w(1) \geq T_w(0)$ almost surely.*

Given monotonicity, trial participants can be classified in each wave as *compliers* (with $T_w(1) > T_w(0)$), *always-takers* (with $T_w(1) = T_w(0) > 0$), or *never-takers* (with $T_w(1) = T_w(0) = 0$). For

⁹See, e.g., Athey and Imbens [2022]; Borusyak, Jaravel and Spiess [2024]; Callaway and Sant'Anna [2021].

¹⁰Exclusion is formalized by double indexing potential outcomes as in Angrist, Imbens and Rubin [1996]. Let $Y_w(t, z)$ denote a participant's wave- w potential outcome given t years of exposure and assignment z . The exclusion restriction says that $Y_w(t, z) = Y_w(t, z')$ for each $t \leq w$, z , and $z' \neq z$.

always- and never-takers, invasive assignment leaves revascularization exposure unchanged. For compliers, invasive assignment increases exposure.¹¹

Invasive assignment is also assumed to generate a *first stage* in wave 1, making the instrument relevant for treatment received shortly after assignment. This restriction is:

Assumption 3 (Relevance). $0 < P[Z = 1] < 1$ and $E[T_1 | Z = 1] > E[T_1 | Z = 0]$.

Under Assumption 1, instrument relevance implies a positive wave-1 first-stage effect on treatment:

$$E[T_1 | Z = 1] - E[T_1 | Z = 0] = E[T_1(1) - T_1(0)] > 0.$$

Because wave-1 compliers have $T_1(1) > T_1(0)$, the wave-1 first-stage is the probability of wave-1 compliance under Assumption 2. Moreover, because treatment is irreversible, Assumption 3 implies the existence of compliers such that $T_w(1) = w > T_w(0)$ in each wave. In other words, a subset of compliers in each wave, w , has w years of revascularization exposure when assigned invasive.

Under Assumptions 1-3, a simple Wald-type IV estimand using wave-1 data identifies the average causal effect of one year of revascularization exposure for wave-1 compliers. Specifically, in wave 1, revascularization exposure, T_1 , is a Bernoulli treatment that indicates participants revascularized shortly after random assignment. The Imbens and Angrist [1994] local average treatment effect (LATE) theorem applied to wave-1 data implies that:

$$\frac{E[Y_1 | Z = 1] - E[Y_1 | Z = 0]}{E[T_1 | Z = 1] - E[T_1 | Z = 0]} = E[Y_1(1) - Y_1(0) | T_1(1) > T_1(0)]. \quad (1)$$

The causal parameter on the right-hand side of this expression is the average effect of revascularization exposure, $Y_1(1) - Y_1(0)$, on wave-1 compliers. This is obtained by dividing the wave-1 *reduced form* by the wave-1 first stage. The former is defined as the difference in average quality of life between participants assigned invasive and participants assigned conservative, measured in the first wave. The first stage in the Wald denominator is given by the corresponding difference in wave-1 revascularization rates.

The reduced form is an ITT effect that compares outcomes by treatment assigned, that is, by the instrument, Z . The first stage is the difference in treatment rates by assignment status; this is also the share of the trial population for whom treatment status is changed by random assignment. Since this share is between zero and one, ITT effects in a static or short-run analysis of trial data are diluted relative to the effect of revascularization itself. Intuitively, because revascularization assignment is assumed to affect outcomes solely by inducing revascularization among compliers, ITT is diluted by averaging in causal effects of zero for always- and never-takers.

Conventional as-treated analyses of clinical trials ignore random assignment. Rather, as-treated analyses compare outcomes conditional on treatment received, typically with covariate controls (see, for instance, the Smith, Coffman and Hudgens [2021] analysis of the CABANA trial comparing treatments for atrial fibrillation). Our as-treated estimates come from ordinary least squares (OLS)

¹¹Appendix B.2 shows that Assumption 2 is equivalent to assuming revascularization decisions arise from a non-parametric latent-index model.

regressions of Y_w on T_w , with a few baseline controls. Non-random treatment take-up likely biases such estimates; in the LATE framework, this bias can be understood as stemming from differences in health between compliers, always-takers, and never-takers [Angrist, 2004; Aronow and Green, 2013]. We substantiate this view with a comparison of group characteristics in Section 4.3 below.

3.2 IV for Immediate Exposure Effects

In randomized with repeated follow-ups, random assignment changes the distribution of time-varying exposure. Monotonicity requires only that exposure be weakly increasing with invasive assignment. As Figure 2 suggests, however, most ISCHEMIA participants assigned invasive were revascularized in the first year following assignment, with few induced by random assignment to revascularize in later years. Participants assigned conservative, by contrast, flow into revascularization over the life of the trial. This empirical observation motivates an identifying assumption requiring compliers who are assigned invasive to be revascularized shortly after random assignment. This *immediate compliers only* (IMCO) assumption is formalized as:

Assumption 4 (IMCO). *Assignment either shifts participants into treatment immediately or has no effect: in each wave $w \in \{1, \dots, \bar{w}\}$, $P[T_w(1) = w \mid T_w(1) > T_w(0)] = 1$.*

Assumption 4 says that everyone for whom revascularization exposure is changed by random assignment is revascularized in the first wave. The following result shows that this condition is sufficient to identify LATE averaged over immediate complier types. In particular, instrumenting T_1 with Z in wave w identifies a causal effect averaged over all immediate compliers seen in wave w .

Theorem 1 (Immediate-Exposure IV Under IMCO). *Suppose Assumptions 1-3 and IMCO hold. Consider the IV estimator that, for a given w , uses Z to instrument T_1 in*

$$Y_w = \gamma_w + \tau_w T_1 + u_w, \quad (2)$$

where γ_w and τ_w are parameters and the IV estimand is defined by $E[u_w|Z] = 0$. Then, coefficient τ_w in this IV estimand equals

$$\tau_w = \sum_{t=0}^{w-1} \omega_{wt} E[Y_w(w) - Y_w(t) \mid T_w(1) = w, T_w(0) = t], \quad (3)$$

where $\omega_{wt} \equiv P[T_w(1) = w, T_w(0) = t \mid T_1(1) > T_1(0)]$ and $\sum_{t=0}^{w-1} \omega_{wt} = 1$. Moreover, for each $t < w$, ω_{wt} is identified by:

$$\omega_{wt} = -\frac{E[\mathbf{1}[T_w = t] \mid Z = 1] - E[\mathbf{1}[T_w = t] \mid Z = 0]}{E[T_1 \mid Z = 1] - E[T_1 \mid Z = 0]}. \quad (4)$$

This is an IV estimand using Z to instrument T_1 with dependent variable $-\mathbf{1}[T_w = t]$.

Proof. See Appendix A.1. □

Estimand τ_w in Theorem 1 averages effects of increasing exposure from t to w for each $t < w$. For instance, instrumenting T_1 in wave 3 averages cumulative effects from $0 \rightarrow 3$, $1 \rightarrow 3$, and $2 \rightarrow 3$. We refer to τ_w as an *immediate exposure effect* because $t = w$ in any wave w implies exposure in wave 1.¹² The immediate exposure estimand in (3) provides a useful summary measure of immediate revascularization. This estimand does not, however, reveal how revascularization effects evolve with increasing exposure. We turn next to estimands that identify these dynamic effects.¹³

3.3 Identification and Estimation of Dynamic Causal Effects

Dynamic exposure effects are identified with the aid of an assumption that limits systematic heterogeneity in treatment effects across waves:

Assumption 5 (Wave Ignorability). *Average incremental effects for compliers are common across waves: for each pair of waves w and $v \leq w$, and for each exposure time $t \leq v$,*

$$E[Y_w(t) - Y_w(t-1) \mid T_w(1) \geq t > T_w(0)] = E[Y_v(t) - Y_v(t-1) \mid T_v(1) \geq t > T_v(0)] \equiv \lambda_t. \quad (5)$$

This assumption restricts and defines incremental exposure effects for compliers, denoted λ_t . In particular, wave ignorability allows λ_t to vary freely with exposure time while requiring incremental effects to be independent of the wave in which they're seen. For instance, in wave w , the initial incremental effect of revascularization, $Y_w(1) - Y_w(0)$, may differ from the incremental effect of one year of early exposure, $Y_w(w) - Y_w(w-1)$. Assumption 5 also allows effects to differ for compliers, always-takers, and never-takers in each wave. Wave ignorability mirrors restrictions implicit in conventional event study regression models, which typically index dynamic causal effects by event time rather than calendar time or treatment cohort (see, e.g., Miller [2023] for a survey).

The following theorem shows that, in the context of the LATE framework, wave ignorability is sufficient to identify the set of incremental causal effects, λ_t .

Theorem 2 (Dynamic Incremental Causal Effects). *For each $t \in \{1, \dots, \bar{w}\}$, let $D_{wt} \equiv \mathbf{1}[T_w \geq t]$ indicate exposure of at least t periods as of wave w . Stack the data across waves and consider the two-stage least squares (2SLS) estimator that uses Z and $Z \times \mathbf{1}[j = w]$ with $j \in \{2, \dots, \bar{w}\}$ to instrument the set of D_{wt} in the linear model*

$$Y_w = \mu + \sum_{j=2}^{\bar{w}} \alpha_j \mathbf{1}[j = w] + \sum_{t=1}^{\bar{w}} \lambda_t D_{wt} + \eta_w, \quad (6)$$

where μ and α_j are constants and the corresponding 2SLS estimand is defined by $E[\eta_w \mid Z] = 0$ for all w . Given Assumptions 1-3 and 5, the vector of λ_t defined in expression (5) equals the set of λ_t in equation (6).

¹²Theorem 1 is implied by Angrist, Santos and Tecchio [2025] Theorem 2.3. Appendix A.1 gives a direct proof.

¹³IMCO is similar in spirit to the Rose and Shem-Tov [2021] restriction allowing extensive margin compliers only (EMCO). With an ordered treatment under EMCO, everyone whose exposure is increased by assignment moves from $T_w(0) = 0$ to positive $T_w(1)$.

Proof. See Appendix A.2. □

Theorem 2 is proved using the Angrist and Imbens [1995] average causal response (ACR) theorem. The ACR theorem shows that the wave-specific IV estimand using Z to instrument T_w with outcome Y_w can be written as a weighted average of incremental complier revascularization effects for each period up to w . Under wave ignorability, these incremental effects are common across w and can therefore be obtained by solving a system of linear equations linking wave-specific reduced forms and first stages. The 2SLS estimator described in Theorem 2 implements this solution.¹⁴

Beyond average incremental effects of an additional year of exposure, it's useful to know the average causal effect of each level of exposure relative to a common reference outcome with no exposure—for ISCHEMIA, this means no revascularization. These cumulative exposure effects can be obtained under a somewhat stronger version of Assumption 5:

Assumption 6 (Strong Wave Ignorability). *Average incremental effects for compliers at a given level of exposure are common across waves for compliers at any possible margin: for each t, v, w, v' , and t' such that $t \leq v \leq w$ and $t \leq t' \leq v' \leq w$,*

$$E[Y_w(t) - Y_w(t-1) \mid T_w(1) \geq t' > T_w(0)] = E[Y_v(t) - Y_v(t-1) \mid T_{v'}(1) \geq t' > T_{v'}(0)] \equiv \lambda_t.$$

Assumption 6 implies Assumption 5 (to see this, set $v' = v$ and $t' = t$ in Assumption 6). In general, however, Assumption 6 adds to Assumption 5 the requirement that a given incremental causal effect be the same at each follow-up wave for complier groups defined over every relevant level of exposure (indexed by t'), not just the level at which the incremental effect in question is measured (indexed by t). Like Assumption 5, however, Assumption 6 asserts that the period or wave in which effects are measured is irrelevant for incremental causal effects at a given level of exposure.

Under strong wave ignorability, average cumulative causal effects are identified and equal to a 2SLS estimand described in the following theorem.

Theorem 3 (Dynamic Cumulative Causal Effects). *Under Assumptions 1-3 and 6:*

i. For all w and $t \leq w$, dynamic cumulative causal effects are given by

$$E[Y_w(t) - Y_w(0) \mid T_w(1) \geq t > T_w(0)] = \sum_{i=1}^t \lambda_i \equiv \Lambda_t. \quad (7)$$

ii. For each $t \in \{1, \dots, \bar{w}\}$, let $R_{wt} \equiv \mathbf{1}[T_w = t]$. Stack the data across waves and consider a 2SLS estimator that uses Z and $Z \times \mathbf{1}[j = w]$ with $j \in \{2, \dots, \bar{w}\}$ to instrument the set of R_{wt} in

¹⁴The ACR theorem identifies a weighted average of incremental causal effects for partially overlapping groups of compliers, while Theorem 2 identifies average incremental effects for a fixed subpopulation of compliers immediately induced into treatment. To see this, note that for any $t \leq w$, wave ignorability and irreversible treatment imply $\lambda_t = E[Y_w(t) - Y_w(t-1) \mid T_w(1) \geq t > T_w(0)] = E[Y_t(t) - Y_t(t-1) \mid T_t(1) = t > T_t(0)] = E[Y_t(t) - Y_t(t-1) \mid T_1(1) > T_1(0)]$. The second equality here is a consequence of wave ignorability and the fact that $T_t(1) \geq t$ if and only if $T_t(1) = t$; the third equality comes from the fact that, given irreversibility, $T_t(1) = t$ if and only if $T_1(1) = 1$.

the linear model

$$Y_w = \phi + \sum_{j=2}^{\bar{w}} \delta_j \mathbf{1}[j = w] + \sum_{t=1}^{\bar{w}} \Lambda_t R_{wt} + \varepsilon_w, \quad (8)$$

where ϕ and δ_j are constants and the 2SLS estimand is defined by $E[\varepsilon_w | Z] = 0$ for all w . Then the vector of Λ_t defined in (7) equals the set of Λ_t in equation (8).

Proof. See Appendix A.3. □

The first part of Theorem 3 says that cumulative revascularization effects are identified by the sum of incremental revascularization effects. Strong wave ignorability is key to this result: each λ_t captures incremental causal effects for different complier groups. A scenario where trial participants have 0-2 years of exposure highlights the role played by Assumption 6 in this context. In this case, the relevant incremental effects in wave 2 are $\lambda_1 = E[Y_2(1) - Y_2(0) | T_2(1) \geq 1 > T_2(0)]$ and $\lambda_2 = E[Y_2(2) - Y_2(1) | T_2(1) \geq 2 > T_2(0)]$, which average over different complier groups. The appendix proof shows that under strong wave ignorability, effects for these groups are the same at a given increment: $E[Y_2(1) - Y_2(0) | T_2(1) \geq 1 > T_2(0)] = E[Y_2(1) - Y_2(0) | T_2(1) \geq 2 > T_2(0)]$. Hence, $\lambda_1 + \lambda_2 = \Lambda_2$.¹⁵ The second part of Theorem 3 shows that this summation is implemented by instrumenting indicators for the full set of exposure levels, R_{wt} .

3.4 Characterizing Dynamic Compliers and Always-Takers

Are IV estimates of causal effects clinically relevant? This question is answered in part by describing the baseline health (and other characteristics) of different sorts of compliers. In particular, complier characteristics inform assessments of the external validity of IV estimates. IV estimates are highly relevant when the associated complier characteristics are similar to those of a population of interest, typically the study population in an RCT. Moreover, because the treated population includes always-takers, the contrast between complier and always-taker characteristics illuminates the biases we should expect in as-treated per-protocol estimates.

In a dynamic setting, covariate means for compliers and always-takers are identified by the following theorem, which is implied by Theorem A1 in Appendix A.4. The appendix theorem covers mean potential outcomes as well as means for never-takers.

Theorem 4 (Dynamic Characterizations). *Let X denote a baseline covariate, fixed across waves and unchanged by assignment. Suppose that $(X, T_w(1), T_w(0))$ is independent of Z for all $w \in \{1, \dots, \bar{w}\}$ and that Assumptions 2-3 hold. Then:*

¹⁵Equality of incremental causal effects across complier groups simplifies Λ_t in settings with no control-group crossovers in the first wave. By setting $v = v' = t' = w$ in Assumption 6 we have $\lambda_t = E[Y_w(t) - Y_w(t-1) | T_w(1) \geq w > T_w(0)] = E[Y_w(t) - Y_w(t-1) | T_1(1) > T_1(0)]$ since treatment is irreversible; hence $\Lambda_t = \sum_{i=1}^t \lambda_i = E[Y_w(t) - Y_w(0) | T_1(1) > T_1(0)]$. This shows that, when $P[T_1(0) = 1] = 0$, cumulative effect Λ_t equals the average effect of t periods of exposure on the immediately-treated.

i. Immediate complier means are given by

$$E[X | T_1(1) > T_1(0)] = \frac{E[\mathbf{1}[T_1 = 1] \times X | Z = 1] - E[\mathbf{1}[T_1 = 1] \times X | Z = 0]}{E[\mathbf{1}[T_1 = 1] | Z = 1] - E[\mathbf{1}[T_1 = 1] | Z = 0]}. \quad (9)$$

ii. Immediate always-taker means are given by

$$E[X | T_1(1) = T_1(0) = 1] = E[X | T_1 = 1, Z = 0]. \quad (10)$$

Moreover, if Assumption 4 also holds, we have that:

iii. Complier means are constant across waves and equal to immediate complier means:

$$E[X | T_w(1) > T_w(0)] = E[X | T_1(1) > T_1(0)] \quad (11)$$

for all $w \in \{1, \dots, \bar{w}\}$;

iv. For each $w > 1$ and $t \in \{0, \dots, w - 1\}$ such that

$$E[\mathbf{1}[T_w = t] | Z = 1] - E[\mathbf{1}[T_w = t] | Z = 0] \neq 0,$$

disaggregated complier means are given by

$$E[X | T_w(1) = w, T_w(0) = t] = \frac{E[\mathbf{1}[T_w = t] \times X | Z = 1] - E[\mathbf{1}[T_w = t] \times X | Z = 0]}{E[\mathbf{1}[T_w = t] | Z = 1] - E[\mathbf{1}[T_w = t] | Z = 0]}; \quad (12)$$

v. For each $w > 1$, later always-taker means are given by

$$E[X | w > T_w(1) = T_w(0) \geq 1] = E[X | w > T_w \geq 1, Z = 1]. \quad (13)$$

Finally, marginal always-taker means, which average immediate and later always-takers, can be obtained using

$$\begin{aligned} E[X | T_w(1) = T_w(0) \geq 1] \\ = \pi_w E[X | T_1 = 1, Z = 0] + (1 - \pi_w) E[X | w > T_w \geq 1, Z = 1], \end{aligned} \quad (14)$$

where $\pi_w = E[\mathbf{1}[T_w = w] | Z = 0] / \{E[\mathbf{1}[T_w = w] | Z = 0] + E[\mathbf{1}[1 \leq T_w < w] | Z = 1]\}$ is the share of immediate always-takers among all always-takers.

Proof. See Appendix A.4. □

Since identification of average baseline characteristics for immediate compliers and immediate always-takers is analogous to identification of characteristics in a static IV setup, parts (i) and (ii) of Theorem 4 follow from Imbens and Rubin [1997] and Abadie [2003]. The rest of the theorem uses IMCO (Assumption 4) to identify complier and always-takers means in a dynamic framework. Specifically, equation (12) disaggregates compliers based on the exposure level they attain when assigned conservative. These complier means can be computed using an IV estimand that takes

$\mathbf{1}[T_w = t] \times X$ as the outcome while instrumenting $\mathbf{1}[T_w = t]$. Equation (13) recovers average baseline characteristics of those revascularized after the first wave regardless of treatment assignment. Equation (14) combines (10) and (13) to recover the average characteristics of all always-takers by wave w .

4 Revascularization Effects on Quality of Life

We estimate causal effects of revascularization on quality-of-life outcomes by 2SLS, clustering standard errors on patient ID when waves are stacked.¹⁶ The assumptions underpinning this analysis are that treatment assignment is independent of potential outcomes and potential exposure and that some participants are induced to revascularize immediately by assignment to the invasive arm of the trial. The exclusion restriction implicit in Assumption 1 rules out scenarios in which assignment to the invasive strategy improves quality of life with no change in revascularization. This seems plausible in an RCT focused on revascularization as the prelude to additional cardiac care. Patients assigned to the invasive arm were also much more likely to be revascularized in the first year, ensuring the assignment instrument is relevant.

Monotonicity (Assumption 2) restricts the effect of random assignment on revascularization exposure. Importantly, however, monotonicity allows unobserved determinants of revascularization including symptoms, comorbidities, quality-of-life history, physician idiosyncrasies and so on. Monotonicity precludes only scenarios in which assignment to the invasive strategy delays or prevents revascularization—a scenario that seems unlikely. This restriction is supported by the ISCHEMIA trial protocol and by the exposure distributions in Appendix Figure B.1, which shows that being assigned invasive induces earlier revascularization for all exposure values by wave 5. This first order stochastic dominance of the distribution of exposure for those assigned invasive relative to those assigned control is a testable implication of Assumption 2.

4.1 Immediate Exposure Effects under IMCO

Theorem 1 offers a natural starting point for IV estimation of the causal effects of revascularization exposure. Assuming all compliers are induced to revascularize immediately (Assumption 4), IV using random assignment to instrument T_1 identifies τ_w , an average causal effect for all complier groups. Figure 3 plots IV estimates of τ_w (implemented by 2SLS), along with the corresponding OLS estimates of immediate-exposure effects; these are computed by regressing SAQ outcomes on an immediate-exposure dummy in each wave in regression models that include the covariates used for IV. In this context, OLS is an as-treated per-protocol estimator that discards random assignment and is therefore susceptible to unobserved confounding.

¹⁶Appendix B.3 establishes semiparametric efficiency of the sample analog of the 2SLS (IV) estimand in Theorem 2 (proofs for Theorems 1 and 3 are similar). Our empirical specifications add baseline covariates (enrollment region and baseline angina frequency scores) to boost precision. Appendix B.4 presents a simulation experiment comparing statistical properties of IV, OLS, and a g-method estimator for a data generating process modeled on ISCHEMIA.

IV estimates of immediate-revascularization effects are mostly larger than the ITT estimates in Figure 1 and Table 1. ITT estimates for summary scores range from 1.7 – 3.2 points; for angina frequency, ITT estimates fall from a high of almost 3.7 in wave 1 to 1.8 or less in later waves. 2SLS estimates of SAQ summary score effects, by contrast, start around 4 and return to this level in wave 5. Angina frequency gains decline from around 5.5 in the first wave to a low of 2 in wave 3, increasing to around 3 in waves 4 and 5.

OLS (as-treated) estimates of immediate-revascularization effects fall over time, declining to a level close to and statistically indistinguishable from zero by wave 5. Divergence between OLS and 2SLS estimates signals unobserved confounding in the as-treated analysis. Differences in underlying health between compliers and others are documented in Figure 4, which plots average potential outcomes for *immediate* compliers, always-takers, and never-takers (defined by $T_1(z)$) by wave. These estimates use Appendix Theorem A1.¹⁷

Complier mean SAQ summary and angina frequency scores are precisely estimated at around 90 and 95, respectively. Always- and never-takers are mostly sicker than compliers, with always-takers generally the worst off (estimated non-complier means are less precise than estimated complier means because non-complier groups comprise much smaller shares of the study population). Never-taker average health approaches that of compliers by wave 5. This likely reflects the fact that sicker patients are likely to be revascularized even when assigned conservative, while patients assigned invasive who go without revascularization are an increasingly healthy group.

4.2 Dynamic Effects under Wave Ignorability

2SLS estimates of cumulative exposure effects, plotted in Figure 5, suggest that revascularization generates a persistent improvement in quality of life. These estimates were computed using equation (8) in Theorem 3. Estimated exposure effects on SAQ summary scores are remarkably stable at around four. Revascularization is estimated to improve angina scores by about 5.5 points initially, with long-run gains sustained at around 3 points. Dotted lines in the figure mark average estimated exposure effects on summary scores over waves 1-5 and average estimated exposure effects over waves 2-5 for angina scores. For summary scores, a chi-square test of the hypothesis of equal cumulative effects yields a p-value of 0.14. For angina frequency, a chi-square test of constant effects in waves 2-5 generates a p-value of 0.86.¹⁸

2SLS estimates of exposure effects are markedly larger—and statistically distinguishable from—the corresponding OLS estimates. 2SLS estimates, OLS estimates, and the difference between them along with associated standard errors appear in Table 2 for both SAQ outcomes.¹⁹ As can be seen in columns 3 and 6 of the table, differences by type of estimate are statistically

¹⁷Theorem A1 encompasses Theorem 4, below. The latter is used in Section 4.3 to assess the clinical relevance of IV estimates for dynamic compliers.

¹⁸A test of the null hypothesis of equal angina score effects for all five exposure values generates $\chi^2(4) = 36$, a decisive rejection.

¹⁹Unlike the specification test comparing IV and OLS estimates introduced in Hausman [1978], these standard errors are computed allowing unrestricted correlation between estimators.

significant at conventional levels for most of the summary score estimates and for one of the angina score estimates. For both outcomes, joint tests of equality generate decisive rejections. It’s also noteworthy that all of the 2SLS estimates in the table are significantly different from zero while the wave-5 OLS estimates are not.

IV estimates computed using Theorem 3 are robust to empirically relevant deviations from strong wave ignorability. This is demonstrated in Appendix B.5, which shows that an extended model allowing for differences in average cumulative exposure effects across waves generates estimates similar to those reported in Table 2.

4.3 Contrasting Treated Populations

As a brief illustration of Theorem 4, Table 3 compares mean baseline SAQ scores for the full sample, compliers, and always-takers. Complier means are computed using IMCO; by Theorem 4(iii), these are immediate complier means, with differences over time due to attrition. The immediate and marginal always-taker means reported in columns 4-5 of Table 3 are computed using Theorem 4(ii) and Theorem 4(v). Marginal always-taker means average the baseline health of immediate and later always-takers.

As is apparent from the first two columns in Table 3, mean baseline SAQ scores for compliers are virtually indistinguishable from average baseline scores in the full study sample. This suggests IV estimates of revascularization effects are likely to be relevant for the ISCHEMIA study population. And, consistent with the average potential outcomes plotted in Figure 4, estimated mean baseline SAQ scores for always-takers are substantially lower than complier means.

Estimated immediate always-taker means in column 4 are similar to the marginal always-take means in column 5. Disaggregated complier means are omitted since over 80% of the compliers allowed under IMCO move from no exposure to immediate exposure in all waves; similarly, we do not report means for the relatively small group of later always-takers.²⁰ The gap between complier and always-taker baseline health is around 10 points, comparable to wave-1 standard deviations of baseline scores of around 19 points.

5 Summary and Conclusions

This paper develops instrumental variables methods for settings with dynamic treatment exposure and outcomes measured in repeated follow-up waves. Variation in exposure time and repeated follow-ups feature in countless randomized trials. In strategy trials and other pragmatic trials, exposure variation is outside trialists’ control once the assignment die has been cast. Initial assignment notwithstanding, treatment exposure is not random.

²⁰The share of compliers with $T_w(0) = t$ as a share of all compliers is identified for all $t < w$ by equation (3) in Theorem 1. Similar to equation (13), the share of later always-takers (i.e., those with $1 \leq T_w(1) = T_w(0) < w$) is identified by $E[\mathbf{1}[1 \leq T_w < w] \mid Z = 1]$ under IMCO. To see this, note that $E[\mathbf{1}[1 \leq T_w < w] \mid Z = 1] = E[\mathbf{1}[1 \leq T_w(1) < w]] = P[1 \leq T_w(1) = T_w(0) < w]$, where the first equality follows from independence and the second from $T_w(1) < w \implies T_w(1) = T_w(0)$ under IMCO.

Dynamic treatment effects have long posed a challenge for clinical research. Treatment choices made after random assignment are likely correlated with post-trial potential outcomes. Our instrumental variables framework identifies average causal effects as a function of exposure time without imposing either conditional independence of treatment exposure and potential outcomes or constant treatment effects. The identification results developed and applied here exploit novel identifying assumptions that we’ve called IMCO and wave ignorability. IMCO is motivated by compliance dynamics observed in ISCHEMIA data; wave ignorability mirrors event-study models that presume exposure effects are unrelated to calendar time. Along with standard IV assumptions, IMCO and wave ignorability deliver causal effects of time-varying exposure that are easily estimated by 2SLS.

When applied to ISCHEMIA trial data, these tools reveal quality of life improvements that are substantially greater and more persistent than previously reported ITT and per-protocol estimates would suggest. IV estimates imply revascularization of patients with ischemic heart disease yields sustained gains in SAQ summary scores on the order of four points. Moreover, after an initial bump up to 5.5, IV estimates of effects on angina frequency scores are stable at around three points.

In view of these estimates, it’s noteworthy that conventional as-treated per-protocol (OLS) estimates decline markedly over time. Dynamic latent-group characterizations reveal this to be an artifact of the poor health of always-takers and improving health of never-takers. In particular, control group crossovers who are revascularized regardless of assignment are far sicker than compliers. At the same time, never-takers, who forgo revascularization when assigned invasive, constitute an increasingly healthy group. Consequently, the gap in SAQ scores by treatment received shrinks.

The finding that as-treated per-protocol estimates are misleading is not unique to ISCHEMIA. In an analysis of data from a randomized mammography screening trial, [Kowalski \[2023\]](#) shows mammography always-takers to be healthier than screening compliers, leading as-treated (i.e., as-screened) estimates to exaggerate gains from screening.²¹

Finally, our ISCHEMIA results weigh against the view that IV complier populations are likely to be unrepresentative of any clinical population of interest (see, for instance, [Hernán and Robins \[2017\]](#)). The baseline health of ISCHEMIA compliers is similar to that of the overall study population. Of course, this is an empirical result and not a theorem. But this finding too is not unique to ISCHEMIA. In a static IV analysis of colorectal cancer screening trials, [Angrist and Hull \[2023\]](#) report mean complier characteristics much like those of the population invited for cancer screening. We expect the IV tools developed here and elsewhere to generate new, clinically useful findings in a range of applications to come.

References

- Abadie A.** 2003. “Semiparametric Instrumental Variable Estimation of Treatment Response Models.” *Journal of Econometrics*, 113(2): 231–263. Publisher: Elsevier.
- Angrist J.** 2004. “Treatment Effect Heterogeneity in Theory and Practice.” *Economic Journal*, 114(494): C52–C83.
- Angrist JD, Hull P.** 2023. “Instrumental Variables Methods Reconcile Intention-to-Screen Effects Across Pragmatic Cancer Screening Trials.” *Proceedings of the National Academy of Sciences*, 120(51): e2311556120.


²¹[Einav et al. \[2020\]](#) likewise documents nonrandom selection in mammography.

- Angrist JD, Imbens GW.** 1995. "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity." *Journal of the American Statistical Association*, 90(430): 431–442.
- Angrist JD, Imbens GW, Rubin DB.** 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association*, 91(434): 444–455.
- Angrist J, Santos A, Tecchio O.** 2025. "One Instrument, Many Treatments: Instrumental Variables Identification of Multiple Causal Effects." National Bureau of Economic Research Working Paper 34607.
- Aronow PM, Green DP.** 2013. "Sharp Bounds for Complier Average Potential Outcomes in Experiments with Noncompliance and Incomplete reporting." *Statistics & Probability Letters*, 47(2-3): 677–679.
- Athey S, Imbens GW.** 2022. "Design-based Analysis in Difference-In-Differences Settings with Staggered Adoption." *Journal of Econometrics*, 226(1): 62–79.
- Ba BA, Ham JC, LaLonde RJ, Li X.** 2017. "Estimating (Easily Interpreted) Dynamic Training Effects from Experimental Data." *Journal of Labor Economics*, 35(S1): S149–S200.
- Baker SG, Lindeman KS.** 1994. "The paired availability design: a proposal for evaluating epidural analgesia during labor." *Statistics in Medicine*, 13(21): 2269–2278.
- Baker SG, Lindeman KS.** 2024. "Multiple discoveries in causal inference: LATE for the party." *Chance*, 37(2): 21–25.
- Bangalore S, et al.** 2020. "Management of Coronary Disease in Patients with Advanced Kidney Disease." *New England Journal of Medicine*, 382(17): 1608–1618.
- Bangalore S, et al.** 2022. "Clinical and Quality-of-Life Outcomes Following Invasive vs Conservative Treatment of Patients With Chronic Coronary Disease Across the Spectrum of Kidney Function." *JAMA Cardiology*, 7(8): 825–835.
- Bensnes S, Huitfeldt I, Leuven E.** 2023. "Reconciling Estimates of the Long-Term Earnings Effect of Fertility." Institute of Labor Economics IZA Discussion Papers 16174.
- Bloom HS.** 1984. "Accounting for no-shows in experimental evaluation designs." *Evaluation Review*, 8(2): 225–246.
- Borusyak K, Jaravel X, Spiess J.** 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *The Review of Economic Studies*, 91(6): 3253–3285.
- Bowden J, et al.** 2025. "Instrumental Variable Methods to Target Hypothetical Estimands With Longitudinal Repeated Measures Data: Application to the STEP 1 Trial." *Statistics in Medicine*. Open access.
- Bronars SG, Grogger J.** 1994. "The Economic Consequences of Unwed Motherhood: Using Twin Births as a Natural Experiment." *The American Economic Review*, 84(5): 1141–1156.
- Callaway B, Sant’Anna PH.** 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*, 225(2): 200–230.
- Chen X, Santos A.** 2018. "Overidentification in Regular Models." *Econometrica*, 86(5): 1771–1817.
- Cinelli C, Hazlett C.** 2025. "An omitted variable bias framework for sensitivity analysis of instrumental variables." *Biometrika*, 112(2): asaf004.
- Conley TG, Hansen CB, Rossi PE.** 2012. "Plausibly Exogenous." *The Review of Economics and Statistics*, 94(1): 260–272.
- Einav L, Finkelstein A, Oostrom T, Ostriker A, Williams H.** 2020. "Screening and Selection: The Case of Mammograms." *American Economic Review*, 110(12): 3836–70.
- Ferman B, Tecchio O.** 2023. "Identifying Dynamic LATEs with a Static Instrument." *arXiv 2305.18114*.
- Gallen Y, Joensen JS, Johansen ER, Veramendi GF.** 2024. "The Labor Market Returns to Delaying Pregnancy." Available at SSRN 4554407.
- Hansen LP.** 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, 50(4): 1029–1054.
- Hausman JA.** 1978. "Specification Tests in Econometrics." *Econometrica*, 46(6): 1251–1271.
- Hernán MA, Robins JM.** 2006. "Instruments for Causal Inference: An Epidemiologist’s Dream?" *Epidemiology*, 17(4): 360–372.
- Hernán MA, Robins JM.** 2020. *Causal Inference: What If*. Chapman & Hall/CRC.
- Hernán MA, Robins JM.** 2017. "Per-protocol Analyses of Pragmatic Trials." *New England Journal of Medicine*,

- 377(14): 1391–1398.
- Imbens GW, Angrist JD.** 1994. “Identification and Estimation of Local Average Treatment Effects.” *Econometrica*, 62(2): 467–475.
- Imbens GW, Rubin DB.** 1997. “Estimating Outcome Distributions for Compliers in Instrumental Variables Models.” *The Review of Economic Studies*, 64(4): 555–574.
- Jiang Z, Chen S, Ding P.** 2023. “An Instrumental Variable Method for Point Processes: Generalized Wald Estimation Based on Deconvolution.” *Biometrika*, 110(4): 989–1008.
- Kowalski AE.** 2023. “Behaviour within a Clinical Trial and Implications for Mammography Guidelines.” *The Review of Economic Studies*, 90(1): 432–462.
- Kunadian V, et al.** 2024. “Invasive Treatment Strategy for Older Patients with Myocardial Infarction.” *New England Journal of Medicine*, 391(18): 1673–1684.
- Latimer NR, Abrams KR, Lambert PC, Morden JP, Crowther MJ.** 2018. “Assessing methods for dealing with treatment switching in clinical trials: A follow-up simulation study.” *Statistical Methods in Medical Research*, 27(3): 765–784.
- Lin JY, Ten Have TR, Elliott MR.** 2008. “Longitudinal Nested Compliance Class Model in the Presence of Time-varying Noncompliance.” *Journal of the American Statistical Association*, 103(482): 462–473.
- Lundborg P, Plug E, Rasmussen AW.** 2017. “Can Women Have Children and a Career? IV Evidence from IVF Treatments.” *American Economic Review*, 107(6): 1611–37.
- Lundborg P, Plug E, Rasmussen AW.** 2024. “Is There Really a Child Penalty in the Long Run? New Evidence from IVF Treatments.” IZA Discussion Papers 16959.
- Luo J, Ruberg SJ, Qu Y.** 2022. “Estimating the Treatment Effect for Adherers Using Multiple Imputation.” *Pharmaceutical Statistics*, 21(3): 525–534.
- Maron DJ, et al.** 2020. “Initial Invasive or Conservative Strategy for Stable Coronary Disease.” *New England Journal of Medicine*, 382(15): 1395–1407.
- Martinussen T, Vansteelandt S, Tchetgen Tchetgen EJ, Zucker DM.** 2017. “Instrumental variables estimation of exposure effects on a time-to-event endpoint using structural cumulative survival models.” *Biometrics*, 73(4): 1140–1149.
- Masten MA, Poirier A.** 2021. “Salvaging Falsified Instrumental Variable Models.” *Econometrica*, 89(3): 1449–1469.
- Mattei A, Ding P, Ballerini V, Mealli F.** 2024. “Assessing Causal Effects in the Presence of Treatment Switching Through Principal Stratification.” *Bayesian analysis*.
- Michael H, Cui Y, Lorch SA, Tchetgen Tchetgen EJ.** 2024. “Instrumental Variable Estimation of Marginal Structural Mean Models for Time-Varying Treatment.” *Journal of the American Statistical Association*, 119(546): 1240–1251.
- Michiels H, Vandebosch A, Vansteelandt S.** 2024. “Adjusting for Time-Varying Treatment Switches in Randomized Clinical Trials: The Danger of Extrapolation and How to Address It.” *Statistics in Biopharmaceutical Research*, 17(1): 54–66.
- Miller DL.** 2023. “An Introductory Guide to Event Study Models.” *Journal of Economic Perspectives*, 37(2): 203–230.
- Morden JP, Lambert PC, Latimer N, Abrams KR, Wailoo AJ.** 2011. “Assessing methods for dealing with treatment switching in randomised controlled trials: a simulation study.” *BMC Medical Research Methodology*, 11: 4.
- Newcombe R.** 1988. “Explanatory and pragmatic estimates of the treatment effect when deviations from allocated treatment occur.” *Statistics in Medicine*, 7(11): 1179–1186.
- Olarte Parra C, Daniel RM, Bartlett JW.** 2023. “Hypothetical Estimands in Clinical Trials: A Unification of Causal Inference and Missing Data Methods.” *Statistics in Biopharmaceutical Research*, 15(2): 421–432.
- Qu Y, Fu H, Luo J, Ruberg SJ.** 2020. “A General Framework for Treatment Effect Estimators Considering Patient Adherence.” *Statistics in Biopharmaceutical Research*, 12(1): 1–18.
- Robins JM.** 1994. “Correcting for Non-Compliance in Randomized Trials Using Structural Nested Mean Models.” *Communications in Statistics - Theory and Methods*, 23(8): 2379–2412.
- Robins JM.** 1997. “Causal Inference from Complex Longitudinal Data.” In *Latent Variable Modeling and Applica-*


- tions to Causality. eds. Mohamed Berkane. New York: Springer-Verlag.
- Robins JM.** 1998. “Structural Nested Failure Time Models.” In *Encyclopedia of Biostatistics*. Vol. 6, eds. Peter Armitage, and Timothy Colton, 4372–4389. Chichester, UK: John Wiley & Sons.
- Robins JM.** 2000. “Marginal Structural Models Versus Structural Nested Models as Tools for Causal Inference.” In *Statistical Models in Epidemiology: The Environment and Clinical Trials*. Vol. 116 of *IMA Volumes in Mathematics and its Applications*, eds. M. Elizabeth Halloran, and Donald Berry, 95–134. New York, NY, USA: Springer.
- Robins JM.** 2004. “Optimal Structural Nested Models for Optimal Sequential Decisions.” *Proceedings of the Second Seattle Symposium on Biostatistics*.
- Robins JM, Hernán MA.** 2008. “Estimation of the causal effects of time-varying exposures.” In *Longitudinal Data Analysis*. Chapter 23, 553–599. Chapman & Hall/CRC.
- Robins JM, Tsiatis AA.** 1991. “Correcting for Non-Compliance in Randomized Trials Using Rank-Preserving Structural Failure Time Models.” *Communications in Statistics - Theory and Methods*, 20(8): 2609–2631.
- Rose EK, Shem-Tov Y.** 2021. “On Recoding Ordered Treatments as Binary Indicators.” *arXiv 2111.12258*.
- Rytgaard HCW, Gerds TA, van der Laan MJ.** 2022. “Continuous-time targeted minimum loss-based estimation of intervention-specific mean outcomes.” *The Annals of Statistics*, 50(5): 2469–2491.
- Shi J, et al.** 2021. “Instrumental variable estimation for a time-varying treatment and a time-to-event outcome via structural nested cumulative failure time models.” *BMC Medical Research Methodology*, 21(1): 258.
- Smith VA, Coffman CJ, Hudgens MG.** 2021. “Interpreting the Results of Intention-to-Treat, Per-Protocol, and As-Treated Analyses of Clinical Trials.” *JAMA*, 326(5): 433–434.
- Spertus JA, et al.** 2020a. “Health Status after Invasive or Conservative Care in Coronary and Advanced Kidney Disease.” *New England Journal of Medicine*, 382(17): 1619–1628.
- Spertus JA, et al.** 2020b. “Health-Status Outcomes with Invasive or Conservative Care in Coronary Disease.” *New England Journal of Medicine*, 382(15): 1408–1419.
- van der Vaart AW.** 2000. *Asymptotic statistics*. Vol. 3, Cambridge university press.
- Vohra V, Goldin J.** 2024. “Identifying the Cumulative Causal Effect of a Non-Binary Treatment from a Binary Instrument.” *The Review of Economics and Statistics*, 1–20.
- Vytlacil E.** 2002. “Independence, Monotonicity, and Latent Index Models: An Equivalence Result.” *Econometrica*, 70(1): 331–341.
- Yende-Zuma N, Mwambi H, Vansteelandt S.** 2019. “Adjusting the Effect of Integrating Antiretroviral Therapy and Tuberculosis Treatment on Mortality for Noncompliance: A Time-varying Instrumental Variables Analysis.” *Epidemiology*, 30(2): 197–203.
- Ying A, Tchetgen Tchetgen EJ.** 2023. “Structural cumulative survival models for estimation of treatment effects accounting for treatment switching in randomized experiments.” *Biometrics*, 79(3): 1597–1609.

Figure 1: Intention-to-treat Effects by Wave




Notes: This figure reports intention-to-treat (ITT) estimates from the ISCHEMIA trial. These contrast average SAQ scores by treatment assigned. Both outcome variables have means in the range 86-94, with a standard deviation of 13-15. SAQ scores are measured at the time of follow up. These estimates control for baseline angina frequency scores (as in [Spertus et al. \[2020b\]](#)) and for enrollment region. Patients who were deceased or did not complete a follow-up questionnaire are omitted.

Figure 2: Histogram of Revascularization Exposure by Treatment Assigned




Notes: This figure plots the histogram of revascularization exposure by follow-up wave. Exposure $T_w = 0$ for participants not revascularized as of wave w .

Figure 3: Immediate-Exposure Revascularization Effects by Wave




Notes: This figure plots IV (implemented by 2SLS) and OLS estimates of average immediate exposure effects in each wave, identified by IMCO, using equation (2) in Theorem 1. IV estimates use assignment to instrument immediate-exposure, T_1 . OLS estimates are the corresponding (uninstrumented) as-treated T_1 effect. Controls are the same as in Figure 1. Standard errors are clustered at the individual level.

Figure 4: Potential Outcome Means for Immediate-Exposure Groups by Wave



Notes: This figure plots potential outcome means for T_1 compliers, always-takers, and never-takers in each wave, using equations (A.2), (A.3), and (A.9) in Theorem A1. Compliers are treated immediately (i.e., in wave 1) when assigned $Z = 1$ but not otherwise. Always- (never-)takers are treated (untreated) in wave 1 regardless of assignment.

Figure 5: 2SLS Estimates of Revascularization Exposure Effects



Notes: This figure plots 2SLS estimates of average causal effects of 1-5 years of revascularization exposure, relative to never having been revascularized, using equation (8) of Theorem 3. Controls are the same as in Figure 1. Standard errors are clustered at the individual level.

Table 1: Means, Treatment Rates, and ITT Effects

Follow-up wave	Revascularization Rate		Difference in rates	SAQ Summary Score		SAQ Angina Frequency	
	Control Group	Treatment Group		Mean	ITT	Mean	ITT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1	0.116	0.802	0.683 (0.011)	85.7 [15.6]	3.23 (0.455)	92.1 [14.6]	3.69 (0.421)
	N	2173	2146				
2	0.169	0.809	0.637 (0.013)	87.3 [14.6]	1.85 (0.461)	93.3 [13.1]	1.84 (0.415)
	N	1835	1823				
3	0.194	0.808	0.612 (0.016)	87.4 [15.1]	2.57 (0.579)	93.7 [12.9]	1.38 (0.494)
	N	1267	1232				
4	0.244	0.824	0.579 (0.021)	87.7 [15.4]	1.76 (0.779)	93.9 [13.3]	1.69 (0.668)
	N	747	737				
5	0.291	0.834	0.550 (0.032)	88.6 [14.9]	2.00 (1.09)	94.1 [12.9]	1.57 (0.936)
	N	350	320				

Notes: Column 1 reports the revascularization rate by wave for patients randomized to the conservative treatment group. Column 2 reports the corresponding revascularization rate for those assigned invasive. Column 3 reports the first-stage effect of treatment assignment on revascularization by wave: this is column 2 minus column 1. Columns 4 and 6 report sample means. Columns 5 and 7 report ITT estimates for the effect of treatment assigned. SAQ scores are measured at the time of follow up. Patients who were deceased or did not complete a follow-up questionnaire are omitted. Controls are the same as in Figure 1. Standard deviations appear in square brackets and standard errors appear in parentheses.

Table 2: 2SLS and OLS Estimates of Revascularization Exposure Effects

Years of exposure	SAQ Summary Score			SAQ Angina Frequency		
	2SLS	OLS	2SLS-OLS	2SLS	OLS	2SLS-OLS
	(1)	(2)	(3)	(4)	(5)	(6)
1	4.68 (0.665)	1.96 (0.436)	2.72 (0.577)	5.33 (0.615)	2.74 (0.400)	2.59 (0.521)
2	3.05 (0.695)	1.94 (0.453)	1.11 (0.605)	3.06 (0.621)	2.25 (0.397)	0.803 (0.535)
3	4.11 (0.877)	1.96 (0.586)	2.14 (0.756)	2.46 (0.740)	2.05 (0.498)	0.417 (0.651)
4	3.24 (1.20)	2.42 (0.785)	0.818 (1.03)	3.13 (1.02)	2.45 (0.667)	0.682 (0.854)
5	3.77 (1.69)	1.52 (1.13)	2.25 (1.38)	2.94 (1.45)	1.30 (0.980)	1.64 (1.04)
2SLS-OLS joint test		28.0 [0.000]			27.2 [0.000]	

Notes: This table compares 2SLS and OLS estimates of the effects of 1-5 years of revascularization exposure computed using equation (8), stacking data from all waves. Columns 3 and 6 report Hausman [1978]-type t-tests for the difference between 2SLS and OLS estimates, where standard errors are computed using the variance of the difference of estimates. Chi-square statistics at the bottom of the table test 2SLS-OLS joint equality. This statistic has a $\chi^2(5)$ distribution under the null. Controls are the same as in Figure 1. Standard errors, clustered on person, are reported in parentheses. P-values for joint tests appear in brackets in the last row.

Table 3: Complier and Always-Taker Characteristics

Follow-up wave	Sample mean (1)	Complier mean (2)	Always-takers		
			Share among treated (3)	Immediate mean (4)	Marginal mean (5)
Panel A: Baseline SAQ Summary Score					
1	74.3 [18.9]	74.5 (0.589)	0.25	63.6 (1.36)	63.6 (1.36)
2	74.2 [18.9]	73.7 (0.633)	0.24	64.0 (1.51)	64.6 (1.46)
3	74.0 [19.4]	74.0 (0.787)	0.25	62.9 (1.91)	63.7 (1.82)
4	73.8 [20.1]	74.4 (1.08)	0.27	61.3 (2.53)	62.3 (2.34)
5	74.6 [20.1]	75.7 (1.61)	0.28	65.4 (3.68)	67.7 (3.15)
Panel B: Baseline SAQ Angina Frequency Score					
1	81.5 [19.5]	81.2 (0.630)	0.25	71.8 (1.45)	71.8 (1.45)
2	81.4 [19.7]	80.5 (0.683)	0.24	71.5 (1.64)	72.6 (1.56)
3	80.9 [19.7]	80.6 (0.859)	0.25	70.1 (2.07)	71.3 (1.95)
4	80.6 [21.0]	80.7 (1.17)	0.27	69.3 (2.72)	70.9 (2.47)
5	81.5 [20.6]	83.8 (1.68)	0.28	71.1 (4.00)	73.5 (3.35)

Notes: This table reports immediate complier and always-taker means obtained using Theorem 4 for baseline summary scores (Panel A) and baseline angina frequency scores (Panel B). Column 1 shows overall sample means. Column 2 reports immediate complier means computed using equation (9) for each wave. The share of always-takers among the treated appears in column 3; this is the denominator of π_w divided by sample share treated. Column 4 reports immediate always-taker means computed using equation (10); column 5 reports marginal means for always-takers, computed using equation (14). Standard deviations appear in square brackets and robust standard errors appear in parentheses.

Appendix: “Instrumental Variables with Time-Varying Exposure: Dynamic Effects of Revascularization on Quality of Life”

A Proofs

A.1 Proof of Theorem 1

Note that, for any given $w \in \{1, \dots, \bar{w}\}$, Assumption 1 implies

$$\begin{aligned} E[Y_w | Z = 1] &= E \left[\sum_{t=0}^w Y_w(t) \mathbf{1}[T_w(1) = t] \right], \\ E[Y_w | Z = 0] &= E \left[\sum_{t=0}^w Y_w(t) \mathbf{1}[T_w(0) = t] \right]. \end{aligned}$$

Assumption 2 and IMCO then imply

$$\begin{aligned} E[Y_w | Z = 1] - E[Y_w | Z = 0] &= E \left[\sum_{0 \leq t < s \leq w} (Y_w(s) - Y_w(t)) \mathbf{1}[T_w(1) = s, T_w(0) = t] \right] \\ &= E \left[\sum_{t=0}^{w-1} (Y_w(w) - Y_w(t)) \mathbf{1}[T_w(1) = w, T_w(0) = t] \right] \\ &= P[T_w(1) > T_w(0)] \sum_{t=0}^{w-1} \omega_{wt} E[Y_w(w) - Y_w(t) | T_w(1) = w, T_w(0) = t], \end{aligned}$$

where $\omega_{wt} \equiv P[T_w(1) = w, T_w(0) = t | T_w(1) > T_w(0)]$. Note that IMCO and irreversibility of treatment imply

$$P[T_w(1) > T_w(0)] = P[T_1(1) > T_1(0)] = E[T_1 | Z = 1] - E[T_1 | Z = 0].$$

That $\sum_{t=0}^{w-1} \omega_{wt} = 1$ follows from the above and

$$E \left[\sum_{t=0}^{w-1} \mathbf{1}[T_w(1) = w, T_w(0) = t] \right] = E[\mathbf{1}[T_w(1) > T_w(0)]]$$

under IMCO. Moreover, for any $t < w$,

$$E[\mathbf{1}[T_w = t] | Z = 1] - E[\mathbf{1}[T_w = t] | Z = 0] = E[\mathbf{1}[T_w(1) = t] - \mathbf{1}[T_w(0) = t]],$$

which equals $-P[T_w(1) = w, T_w(0) = t]$ under monotonicity and IMCO. To conclude, note that the definition of the IV estimand in the theorem implies

$$\tau_w = \frac{E[Y_w | Z = 1] - E[Y_w | Z = 0]}{E[T_1 | Z = 1] - E[T_1 | Z = 0]}.$$

A.2 Proof of Theorem 2

As in the proof of Theorem 1 in Angrist and Imbens [1995], Assumptions 1-2 can be used to show that the reduced form for wave-specific IV with outcome Y_w and T_w instrumented by Z for given $w \in \{1, \dots, \bar{w}\}$ can be written as

$$\rho_w = \sum_{t=1}^w E[Y_w(t) - Y_w(t-1) | T_w(1) \geq t > T_w(0)] \pi_{wt},$$

where $\pi_{wt} \equiv P[T_w(1) \geq t > T_w(0)]$. Under Assumption 5, this becomes:

$$\rho_w = \sum_{t=1}^w \lambda_t \pi_{wt}.$$

Let ρ denote the vector of ρ_w , with Π being a lower-triangular matrix with non-zero elements π_{wt} (w -th row, t -th column). Finally, let λ be the vector of λ_t . Then, we can write:

$$\rho = \Pi \lambda.$$

Under Assumption 3, Π is invertible since it is lower-triangular with non-zero diagonal elements (absorbing treatment implies $\pi_{ww} = \pi_{11}$ for all w). Hence:

$$\lambda = \Pi^{-1} \rho. \tag{A.1}$$

The 2SLS procedure described in the theorem using equation (6) has reduced form and first stage parameter vectors identical to the stacked wave-by-wave reduced form and first stage described here since the stacked model interacts Z with wave dummies and allows wave-specific intercepts. Hence, these parameters satisfy (A.1).

A.3 Proof of Theorem 3

We first show that $\sum_{i=1}^t \lambda_i = E[Y_w(t) - Y_w(0) | T_w(1) \geq t > T_w(0)]$ under Assumption 6 (strong wave ignorability), for any w and $t \leq w$. Note that for all $i < t$:

$$\{T_i(1) \geq i > T_i(0)\} \iff \{T_t(1) \geq t > T_t(0)\},$$

since, with an absorbing treatment, those with t months of exposure at wave t must have been exposed initially. Hence, for any $i < t$:

$$E[Y_i(i) - Y_i(i-1) | T_i(1) \geq i > T_i(0)] = E[Y_i(i) - Y_i(i-1) | T_t(1) \geq t > T_t(0)].$$

Moreover, for any $i < t$, strong wave ignorability implies:

$$E[Y_i(i) - Y_i(i-1) | T_i(1) \geq i > T_i(0)] = E[Y_w(i) - Y_w(i-1) | T_w(1) \geq i > T_w(0)].$$

and

$$E[Y_i(i) - Y_i(i-1) | T_t(1) \geq t > T_t(0)] = E[Y_w(i) - Y_w(i-1) | T_w(1) \geq t > T_w(0)].$$

Thus:

$$\begin{aligned} \sum_{i=1}^t \lambda_i &= \sum_{i=1}^t E[Y_w(i) - Y_w(i-1) | T_w(1) \geq i > T_w(0)] \\ &= \sum_{i=1}^t E[Y_w(i) - Y_w(i-1) | T_w(1) \geq t > T_w(0)] \\ &= E \left[\sum_{i=1}^t Y_w(i) - Y_w(i-1) \middle| T_w(1) \geq t > T_w(0) \right] \\ &= E[Y_w(t) - Y_w(0) | T_w(1) \geq t > T_w(0)]. \end{aligned}$$

For the second part of the theorem, note that $R_{wt} = \sum_{i=1}^t D_{wi}$. Thus, the 2SLS estimand instrumenting R_{wt} in equation (8) generates the sum of coefficients from the 2SLS estimand instrumenting D_{wt} in equation (6). That is, $\sum_{i=1}^t \lambda_i = \Lambda_t$.

A.4 Marginal Potential Outcome Means

This appendix extends Theorem 4 to identify expectations of potential outcomes for different types of compliers and always-takers. Theorem 4 in the text, which establishes identification of baseline covariate means in our dynamic setup, is a consequence of this more general result. In particular, Theorem 4 follows from the theorem below by setting $Y_w(t) = X$ for all w and t .

Theorem A1. *Suppose Assumptions 1-3 hold. Then:*

i. Immediate complier means are given by

$$E[Y_w(w) | T_1(1) > T_1(0)] = \frac{E[\mathbf{1}[T_1 = 1] \times Y_w | Z = 1] - E[\mathbf{1}[T_1 = 1] \times Y_w | Z = 0]}{E[\mathbf{1}[T_1 = 1] | Z = 1] - E[\mathbf{1}[T_1 = 1] | Z = 0]}. \quad (\text{A.2})$$

ii. Immediate always-taker means are given by

$$E[Y_w(w) | T_1(1) = T_1(0) = 1] = E[Y_w | T_1 = 1, Z = 0]. \quad (\text{A.3})$$

Moreover, if Assumption 4 also holds, we have that:

iii. Complier means equal immediate complier means in each wave:

$$E[Y_w(w) | T_w(1) > T_w(0)] = E[Y_w(w) | T_1(1) > T_1(0)] \quad (\text{A.4})$$

for all $w \in \{1, \dots, \bar{w}\}$;

iv. For each $w > 1$ and $t \in \{0, \dots, w-1\}$ such that

$$E[\mathbf{1}[T_w = t] \mid Z = 1] - E[\mathbf{1}[T_w = t] \mid Z = 0] \neq 0,$$

disaggregated complier means are given by

$$E[Y_w(t) \mid T_w(1) = w, T_w(0) = t] = \frac{E[\mathbf{1}[T_w = t] \times Y_w \mid Z = 1] - E[\mathbf{1}[T_w = t] \times Y_w \mid Z = 0]}{E[\mathbf{1}[T_w = t] \mid Z = 1] - E[\mathbf{1}[T_w = t] \mid Z = 0]}, \quad (\text{A.5})$$

v. For each $w > 1$, later always-taker means are given by

$$E[Y_w(T_w(1)) \mid w > T_w(1) = T_w(0) \geq 1] = E[Y_w \mid w > T_w \geq 1, Z = 1]. \quad (\text{A.6})$$

Marginal always-takers means, which average immediate and later always-takers, can be obtained using

$$\begin{aligned} & E[Y_w(T_w(1)) \mid T_w(1) = T_w(0) \geq 1] \\ &= \pi_w E[Y_w \mid T_w = w, Z = 0] + (1 - \pi_w) E[Y_w \mid w > T_w \geq 1, Z = 1], \end{aligned} \quad (\text{A.7})$$

$$\text{where } \pi_w \equiv \frac{E[\mathbf{1}[T_w = w] \mid Z = 0]}{E[\mathbf{1}[T_w = w] \mid Z = 0] + E[\mathbf{1}[1 \leq T_w < w] \mid Z = 1]}.$$

vi. For each $t \in \{0, \dots, w-1\}$, disaggregated always- and never-taker means are given by

$$E[Y_w(t) \mid T_w(1) = T_w(0) = t] = E[Y_w \mid T_w = t, Z = 1]. \quad (\text{A.8})$$

For never-takers, set $t = 0$. Identification of never-taker means does not require IMCO (Assumption 4). Finally, for any $w \in \{1, \dots, \bar{w}\}$, immediate never-taker means are

$$E[Y_w \mid T_1 = 0, Z = 1] = E[Y_w \mid 0 \leq T_w < w, Z = 1] = \sum_{t=0}^{w-1} \tilde{\omega}_{wt} E[Y_w(t) \mid T_w(1) = T_w(0) = t], \quad (\text{A.9})$$

where $\tilde{\omega}_{wt} \equiv P[T_w(1) = T_w(0) = t \mid 0 \leq T_w < w]$ for all $t \in \{0, \dots, w-1\}$.

Proof. Consider the first two results, (A.2) and (A.3), which do not require Assumption 4. Fix $w \in \{1, \dots, \bar{w}\}$. Because treatment is irreversible, $T_w = w$ if, and only if, $T_1 = 1$. Monotonicity then implies:

$$\begin{aligned} E[\mathbf{1}[T_1 = 1] \times Y_w \mid Z = 1] &= E[\mathbf{1}[T_1(1) = 1] \times Y_w(w)] \\ &= P[T_1(1) = T_1(0) = 1] E[Y_w(w) \mid T_1(1) = T_1(0) = 1] \\ &\quad + P[T_1(1) > T_1(0)] E[Y_w(w) \mid T_1(1) > T_1(0)]. \end{aligned}$$

Analogously, $E[\mathbf{1}[T_1 = 1] \times Y_w \mid Z = 0] = P[T_1(1) = T_1(0) = 1] E[Y_w(w) \mid T_1(1) = T_1(0) = 1]$, which establishes equation (A.2). Equation (A.3) follows from $E[Y_w \mid T_1 = 1, Z = 0] = E[Y_w(w) \mid$

$T_1(0) = 1] = E[Y_w(w) \mid T_1(1) = T_1(0) = 1]$ under monotonicity.

Now, for a given $w \in \{2, \dots, \bar{w}\}$, consider results in the theorem that depend on Assumption 4:

1. To establish (A.4), note that because treatment is irreversible, $T_1(1) > T_1(0)$ implies $T_w(1) > T_w(0)$. Conversely, $T_w(1) > T_w(0)$ implies $T_w(1) = w$ almost surely, which in turn implies $T_1(1) > T_1(0)$ because treatment is irreversible.
2. Fix $t \in \{0, \dots, w-1\}$. Note that when $Z = 1$, $\mathbf{1}[T_w = t] = 1$ if and only if $T_w(1) = t$. Therefore, under Assumptions 2 and 4, $E[\mathbf{1}[T_w = t] \mid Z = 1] = P[T_w(1) = T_w(0) = t \mid Z = 1]$ since compliers have $T_w(1) = w$ almost surely. Moreover, under monotonicity,

$$E[\mathbf{1}[T_w = t] \times Y_w \mid Z = 1] = P[T_w(1) = T_w(0) = t]E[Y_w(t) \mid T_w(1) = T_w(0) = t]. \quad (\text{A.10})$$

When $Z = 0$, in addition to always-takers with $T_w(0) = t$, $\mathbf{1}[T_w = t] = 1$ for compliers with $T_w(0) = t < w = T_w(1)$. Thus,

$$\begin{aligned} E[\mathbf{1}[T_w = t] \times Y_w \mid Z = 0] &= P[T_w(1) = T_w(0) = t]E[Y_w(t) \mid T_w(1) = T_w(0) = t] \\ &\quad + P[T_w(1) = w, T_w(0) = t]E[Y_w(t) \mid T_w(1) = w, T_w(0) = t], \end{aligned}$$

which together with (A.10) establishes (A.5). Equation (A.8) follows from the same argument used to establish (A.10).

3. To show (A.6), note that monotonicity implies

$$\begin{aligned} E[Y_w \mid w > T_w \geq 1, Z = 1] &= E[Y_w(T_w(1)) \mid w > T_w(1) \geq 1] \\ &= E[Y_w(T_w(1)) \mid w > T_w(1) = T_w(0) \geq 1], \end{aligned}$$

since $P[T_w(1) = T_w(0) \mid 1 \leq T_w(1) < w] = 1$ under Assumption 4.

4. To show (A.7), note that

$$\begin{aligned} &E[Y_w(T_w(1)) \mid T_w(1) = T_w(0) \geq 1] \\ &= P[T_w(1) = T_w(0) = w \mid T_w(1) = T_w(0) \geq 1]E[Y_w(T_w(1)) \mid T_w(1) = T_w(0) = w] \\ &\quad + P[w > T_w(1) = T_w(0) \geq 1 \mid T_w(1) = T_w(0) \geq 1]E[Y_w(T_w(1)) \mid w > T_w(1) = T_w(0) \geq 1]. \end{aligned}$$

The fact that $E[Y_w(T_w(1)) \mid T_w(1) = T_w(0) = w] = E[Y_w \mid T_w = w, Z = 0]$ follows from equation (A.3) by noting that $T_1(1) = T_1(0) = 1$ if, and only if, $T_w(1) = T_w(0) = w$ because treatment is irreversible. The result that $E[Y_w(T_w(1)) \mid w > T_w(1) = T_w(0) \geq 1] = E[X_w \mid w > T_w \geq 1, Z = 1]$ follows from (A.6). Finally, monotonicity implies $E[\mathbf{1}[T_w = w] \mid Z = 0] = P[T_w(1) = T_w(0) = w]$ and, under IMCO, $E[\mathbf{1}[1 \leq T_w < w] \mid Z = 1] = P[w > T_w(1) = T_w(0) \geq 1]$. Thus, $\pi_w = P[T_w(1) = T_w(0) = w \mid T_w(1) = T_w(0) \geq 1]$.

Note that equation (14) uses $E[X \mid T_1 = 1, Z = 0]$ while equation (A.7) uses $E[Y_w \mid T_w = w, Z = 0]$. Because treatment is irreversible, conditioning on $T_1 = 1$ is equivalent to condition

on $T_w = w$. In the case of equation (14), because X does not vary across waves, this means that $E[X \mid T_1 = 1, Z = 0]$ can be computed only once even though $E[X \mid T_w(1) = T_w(0) \geq 1]$ varies with w . On the other hand, in equation (A.7), $E[Y_w \mid T_w = w, Z = 0]$ varies with w even though the latent group does not.

5. To conclude, fix $w \in \{1, \dots, \bar{w}\}$. The first equality in (A.9) is immediate from irreversibility of treatment. The second equality then follows from

$$E[Y_w \mid 0 \leq T_w < w, Z = 1] = \sum_{t=0}^{w-1} P[T_w = t \mid 0 \leq T_w < w, Z = 1] E[Y_w \mid T_w = t, Z = 1].$$

From (A.8), $E[Y_w \mid T_w = t, Z = 1]$ equals $E[Y_w(t) \mid T_w(1) = T_w(0) = t]$ for all $t \in \{0, \dots, w-1\}$. Lastly, for any $t \in \{0, \dots, w-1\}$, Assumption 1 implies

$$P[T_w = t \mid 0 \leq T_w < w, Z = 1] = P[T_w(1) = t \mid 0 \leq T_w < w],$$

which equals $P[T_w(1) = T_w(0) = t \mid 0 \leq T_w < w]$ under IMCO.

□

B Supplemental Material

B.1 Comparison to Bowden et al. [2025]

This appendix compares our identification strategy to the use of instruments in Section 2.3 of Bowden et al. [2025]. Bowden et al. [2025] is more restrictive in that it does not allow for control-group crossovers; in our notation, this means $T_w(0) = 0$ with probability one. A second difference is that Bowden et al. [2025] impose a parametric specification for the dynamic incremental effects λ_t . In our notation, this amounts to imposing the restriction:

$$\lambda_t = \beta \alpha^t, \tag{B.1}$$

for $\beta, \alpha \in \mathbb{R}$, with α implicitly assumed to be non-negative (see equation (6) and related discussions in Bowden et al. [2025]).²² This parametric restriction on the incremental treatment effects implies our wave ignorability assumption (Assumption 5) but is much stronger: for example, with $\alpha \geq 0$, it doesn't allow for non-monotonic dynamics of incremental effects or for cumulative treatment effects that are initially positive and then decrease over time. Even with unrestricted α , Equation (B.1) rules out several plausible paths for treatment effects—such as when effects persist for more than one period before decaying (see Appendix Figure B.2).

We now extend the Bowden et al. [2025] framework to allow for control-group crossovers and show how their parametric model for incremental effects can be imposed and tested in our setting.

²²Note that the definition of the causal effects $\beta_k(j)$ in their Section 2.1 restricts potential outcome heterogeneity across IV latent types.

Using the notation of Theorem 2, the estimator Bowden et al. [2025] propose is a generalized method of moments (GMM) estimator of β and α based on the moment condition:

$$E \left[(Z - E[Z]) \begin{bmatrix} Y_1 - \beta D_{11} Z \\ Y_2 - (\beta \alpha D_{22} + \beta D_{21}) Z \\ \vdots \\ Y_{\bar{w}} - \sum_{j=1}^{\bar{w}} D_{wj} \beta \alpha^{j-1} Z \end{bmatrix} \right] = 0. \quad (\text{B.2})$$

We propose a modified estimator which replaces each $D_{wt}Z$ term with just D_{wt} :

$$E \left[(Z - E[Z]) \begin{bmatrix} Y_1 - \beta D_{11} \\ Y_2 - \beta \alpha D_{22} - \beta D_{21} \\ \vdots \\ Y_{\bar{w}} - \sum_{j=1}^{\bar{w}} D_{wj} \beta \alpha^{j-1} \end{bmatrix} \right] = 0. \quad (\text{B.3})$$

This modification is inconsequential in settings without control-group crossovers: $Z = 0$ implies $D_{wt} = 0$ in that case, such that $D_{wt}Z = D_{wt}$. In the more general setting we consider, an estimator based on equation (B.3) can be seen as a restricted version of the 2SLS estimator in Theorem 2 which imposes the parametric restriction (B.1) on incremental effects. That is, if one were to relax the equation (B.1) assumption by replacing $\beta \alpha^t$ with unrestricted λ_t coefficients in equation (B.3), the resulting estimates would coincide exactly with the 2SLS estimates in Theorem 2.²³

Table B.1 reports GMM estimates of β and α for the two primary ISCHEMIA outcomes, using the extended moment condition (B.3) after adding linear controls as in the main analysis. Since β and α are overidentified in this setting, we use the efficiently-weighted two-step GMM estimator and report overidentification test statistics and p-values [Hansen, 1982]. Columns 1-3 of the table estimate a version of the model that is faithful to the original Bowden et al. [2025] specification, with $\alpha \geq 0$, while columns 4-6 allow α to be negative.²⁴ Columns 1 and 4 pool the two ISCHEMIA outcomes in order to increase precision.

The table shows that the augmented Bowden et al. [2025] model of incremental effects does not fit the ISCHEMIA data well under natural parameter values. In columns 1-3 the estimate of α is driven to the lower bound of zero and overidentification tests largely reject—with $p < 0.01$ in the pooled model and the SAQ Angina Frequency outcome. Even when allowing negative α in columns 4-6, the pooled model yields a marginal overidentification test rejection. Broadly, these results highlight the value of our nonparametric approach.

²³Since Z is binary, the moment conditions $E[\eta_w] = E[Z\eta_w] = 0$, $w = 1, \dots, \bar{w}$, define the 2SLS estimand in Theorem 2. For wave w , these imply $\mu + \alpha_w = E[Y_w - \sum_{t=1}^w \lambda_t D_{wt}]$ and $E[(Z - E[Z])\eta_w] = 0$. Therefore, they imply $E[(Z - E[Z])(\eta_w - \mu - \alpha_w)] = 0$; stacking across waves gives (B.3). The relevance condition (Assumption 3) implies that the solutions to the 2SLS moment condition and (B.3) are unique and hence coincide.

²⁴In practice, to impose $\alpha \geq 0$, we first minimize the GMM criterion function allowing α to take only strictly positive values, which drives it to zero. Then, we impose $\alpha = 0$ and choose β to minimize the GMM criterion under such constraint. Columns 1-3 report results from this second step.

B.2 Latent-Index Representation of Monotonicity

This appendix shows that Assumption 2 is implied by a nonparametric latent-index model. The argument here is similar to that in Vytlačil [2002] for a single static treatment. Let

$$S_i(z) = \inf\{s \in \{1, \dots, \bar{w}\} : T_{is}(z) > 0\},$$

with $\inf \emptyset = \infty$. Since revascularization is absorbing,

$$T_{iw}(z) = \sum_{s=1}^w \mathbf{1}\{S_i(z) \leq s\}.$$

Proposition B.1. *The following two statements are equivalent:*

$$T_{iw}(1) \geq T_{iw}(0) \quad \text{for every } w = 1, \dots, \bar{w},$$

and

$$S_i(1) \leq S_i(0).$$

Moreover, this condition is equivalent to the existence of a scalar assignment index $\nu(z)$, normalized as $\nu(0) = 0$ and $\nu(1) = 1$, and patient-specific random thresholds

$$C_{i1} \geq C_{i2} \geq \dots \geq C_{i\bar{w}}$$

such that

$$\mathbf{1}\{S_i(z) \leq s\} = \mathbf{1}\{\nu(z) \geq C_{is}\}, \quad s = 1, \dots, \bar{w}, \quad z \in \{0, 1\}.$$

Equivalently,

$$T_{iw}(z) = \sum_{s=1}^w \mathbf{1}\{\nu(z) \geq C_{is}\}.$$

Proof. First, suppose $S_i(1) \leq S_i(0)$. Then, for every s ,

$$\mathbf{1}\{S_i(1) \leq s\} \geq \mathbf{1}\{S_i(0) \leq s\}.$$

Summing from $s = 1$ to w gives

$$T_{iw}(1) = \sum_{s=1}^w \mathbf{1}\{S_i(1) \leq s\} \geq \sum_{s=1}^w \mathbf{1}\{S_i(0) \leq s\} = T_{iw}(0).$$

Conversely, suppose $S_i(1) > S_i(0)$. At $w = S_i(0)$, we have $T_{iw}(0) > 0$ but $T_{iw}(1) = 0$, contradicting $T_{iw}(1) \geq T_{iw}(0)$. Therefore

$$T_{iw}(1) \geq T_{iw}(0) \text{ for all } w \quad \Longleftrightarrow \quad S_i(1) \leq S_i(0).$$

It remains to establish the latent-index representation. If the latent-index representation holds, monotonicity follows immediately because $\nu(1) \geq \nu(0)$:

$$\mathbf{1}\{\nu(1) \geq C_{is}\} \geq \mathbf{1}\{\nu(0) \geq C_{is}\}$$

for every s . Summing over $s \leq w$ yields $T_{iw}(1) \geq T_{iw}(0)$.

For the converse, assume monotonicity, or equivalently $S_i(1) \leq S_i(0)$. Normalize $\nu(0) = 0$ and $\nu(1) = 1$. Define

$$C_{is} = \begin{cases} 2, & s < S_i(1), \\ 1/2, & S_i(1) \leq s < S_i(0), \\ -1, & s \geq S_i(0). \end{cases}$$

These thresholds are weakly decreasing in s . Moreover,

$$\mathbf{1}\{1 \geq C_{is}\} = \mathbf{1}\{S_i(1) \leq s\},$$

and

$$\mathbf{1}\{0 \geq C_{is}\} = \mathbf{1}\{S_i(0) \leq s\}.$$

Thus,

$$\mathbf{1}\{S_i(z) \leq s\} = \mathbf{1}\{\nu(z) \geq C_{is}\}, \quad z \in \{0, 1\}.$$

This establishes the result. □

B.3 Semiparametric Efficiency

This appendix shows that the sample-analog 2SLS estimator is semiparametrically efficient for the 2SLS estimand in Theorem 2. Proofs of analogous results for Theorems 1 and 3 follow by the same steps. The results stem from the fact that the estimands in Theorem 1-3 are defined by moment conditions whose number equals the number of parameters, resulting in a just-identified Generalized Method of Moments (GMM) model for the distribution of data. [Chen and Santos \[2018\]](#) show that when the Jacobian matrix in GMM estimation is full-rank the notion of just-identification in GMM coincides, under mild conditions, with a broader notion of just-identification that implies all regular estimators are asymptotically equivalent (see their GMM Illustration Example and Theorem 3.1). What follows specializes this more general efficiency result to our setting.

Let X be the random vector stacking instrument, outcomes, and treatment status for all waves:

$$X \equiv (Z, Y_1, \dots, Y_{\bar{w}}, T_1, \dots, T_{\bar{w}})'$$

For simplicity, when dealing with derivatives of functionals over probability measures, we assume that the distribution of $Y \equiv (Y_1, \dots, Y_{\bar{w}})'$ has bounded support (which also guarantees 2SLS is regular, per [Chen and Santos \[2018\]](#), Remark 3.1). This and the other support restrictions for the

instrument (binary) and treatment (integer-valued and irreversible) are encoded in the convention that X takes values in a set $\mathcal{X}$. We then let $\mathcal{P}$ be the set of probability measures over $\mathcal{X}$ such that $0 < E_P[Z] < 1$ and $E_P[T_1 | Z = 1] - E_P[T_1 | Z = 0] > 0$, where, in this Appendix, we index expectations with the probability measure.

For $w \in \{1, \dots, \bar{w}\}$ and $t \leq w$, define $D_{wt} \equiv \mathbf{1}[T_w \geq t]$. Let $\mathbf{D}$ denote the random $\bar{w} \times \bar{w}$ lower-triangular matrix where entry (w, t) is equal to D_{wt} for $t \leq w$ and zero for $t > w$. Finally, for a probability measure $P \in \mathcal{P}$, let

$$\rho(P) \equiv E_P[Y | Z = 1] - E_P[Y | Z = 0], \quad \Pi(P) \equiv E_P[\mathbf{D} | Z = 1] - E_P[\mathbf{D} | Z = 0].$$

Since $\Pi(P)$ is lower triangular with a non-zero diagonal, it is invertible for any $P \in \mathcal{P}$. By Appendix A.2, the functional

$$P \mapsto \lambda(P) \equiv \Pi(P)^{-1} \rho(P)$$

identifies the target parameter in Theorem 2.

We consider efficient estimation of this functional when defined on $\mathcal{P}$. The next proposition establishes that the asymptotic variance of any regular estimator of $P \mapsto \lambda(P)$ is no smaller than that of the 2SLS estimator.

Proposition B.2. *Consider the set of probability measures $\mathcal{P}$ and the functional $\lambda : \mathcal{P} \rightarrow \mathbb{R}^{\bar{w}}$ defined above. Then the asymptotic variance matrix of any regular sequence of estimators of λ (as defined in van der Vaart [2000], Chapter 25) is bounded from below, in the positive semi-definite order, by the asymptotic variance matrix of the 2SLS estimator.*

Proof. Fix $P \in \mathcal{P}$. The tangent space of $\mathcal{P}$ at P is the set of real-valued functions that are mean-zero and square-integrable with respect to P , $L_0^2(P)$. To see this, proceed as in Example 25.16 of van der Vaart [2000]: fix $g \in L_0^2(P)$ to construct a one-dimensional parametric submodel $(P_t)_{t \in [0,1]}$ such that

$$t \mapsto \frac{dP_t}{dP} = \frac{\psi(tg)}{\int \psi(tg(x))dP(x)}, \quad \mathbb{R} \ni u \mapsto \psi(u) \equiv \frac{2}{1 + \exp(-2u)},$$

where dP_t/dP is the Radon-Nykodim derivative of P_t with respect to P . This parametric submodel is in $\mathcal{P}$ for sufficiently small $t \in [0, 1]$: since $u \mapsto \psi(u)$ is bounded, the Dominated Convergence Theorem implies that $\int |dP_t - dP| \rightarrow 0$ as $t \downarrow 0$ and, therefore, that

$$|E_{P_t}[\mathbf{1}[Z = z]] - E_P[\mathbf{1}[Z = z]]| \rightarrow 0, \quad |E_{P_t}[\mathbf{1}[Z = z] T_1] - E_P[\mathbf{1}[Z = z] T_1]| \rightarrow 0$$

for $z \in \{0, 1\}$ as $t \downarrow 0$. Moreover, the submodel $(P_t)_{t \in [0,1]}$ is differentiable in quadratic mean at $P = P_0$ (i.e., it satisfies (25.13) in van der Vaart [2000]) as per Example 25.16 in van der Vaart [2000]. This shows that the tangent space of $\mathcal{P}$ at P is indeed $L_0^2(P)$.

Now fix an arbitrary one-dimensional parametric submodel $(P_t)_{t \in [0,1]}$ in $\mathcal{P}$ such that $P_0 = P$ and $(P_t)_{t \in [0,1]}$ is differentiable in quadratic mean at P with score function $g \in L_0^2(P)$. Differentiability

in quadratic mean and the Cauchy-Schwarz inequality imply

$$\begin{aligned} \left\| \frac{dP_t - dP}{t} - g dP \right\|_1 &\leq \left\| \frac{g}{2} dP^{1/2} \right\|_2 \left\| dP_t^{1/2} - dP^{1/2} \right\|_2 \\ &\quad + \left\| \frac{dP_t^{1/2} - dP^{1/2}}{t} - \frac{g}{2} dP^{1/2} \right\|_2 \left\| dP_t^{1/2} + dP^{1/2} \right\|_2 = o(1) \end{aligned}$$

as $t \downarrow 0$. Since Y is bounded, the above display implies

$$\lim_{t \downarrow 0} \frac{\lambda(P_t) - \lambda(P)}{t} = E_P[\phi_\lambda(X)g(X)],$$

where

$$\begin{aligned} \phi_\lambda(X) &\equiv \Pi(P)^{-1} \left[\frac{Z}{p} \{Y - \mu_1^Y - (\mathbf{D} - \mu_1^D)\lambda(P)\} - \frac{1-Z}{1-p} \{Y - \mu_0^Y - (\mathbf{D} - \mu_0^D)\lambda(P)\} \right] \\ &= \text{Cov}(Z, \mathbf{D})^{-1} [(Z - p)(Y - \mathbf{D}\lambda(P) - E_P[Y - \mathbf{D}\lambda(P)])]. \end{aligned}$$

and

$$p \equiv P[Z = 1], \quad \mu_z^Y \equiv E_P[Y \mid Z = z], \quad \mu_z^D \equiv E_P[\mathbf{D} \mid Z = z].$$

Because $\phi_\lambda \in L^2(P)$, the map $g \mapsto E_P[\phi_\lambda(X)g(X)]$ defined on $L^2(P)$ is linear and continuous. Therefore, Theorem 25.20 in [van der Vaart \[2000\]](#) implies that $\text{Var}_P[\phi_\lambda(X)]$ is the lower bound for the variance matrix of the asymptotic variance of every regular sequence of estimators.

To conclude the proof, note that for an i.i.d. sample $(X_i)_{i=1}^n$ from the probability measure P , the 2SLS estimator has the asymptotically linear expansion

$$\sqrt{n}[\hat{\lambda} - \lambda(P)] = \frac{1}{\sqrt{n}} \sum_{i=1}^n \phi_\lambda(X_i) + o_p(1)$$

as $n \rightarrow \infty$, since Y is bounded. □

B.4 Finite-sample Behavior of Alternative Estimators

This appendix reports results of a small-scale simulation study, tailored to the ISCHEMIA data, that illustrates the finite-sample performance of the 2SLS estimator in Theorem 3 and compares it to OLS and a g-method estimator. The target parameters are cumulative causal effects, set equal to main text estimates reported in Figure 5. For simplicity and to ease comparison across estimators, the simulation sets these as constant across observations.

We simulate an unbalanced panel with 5 waves, mirroring the data structure of ISCHEMIA. We vary the number of initial observations across 500, 1,500, and 5,000. Inclusion in the sample at wave w is determined by

$$M_w = M_{w-1}Q_w, \quad Q_w \sim \text{Bernoulli}(q_w/q_{w-1}), \quad w = 2, \dots, 5,$$

for $M_1 \equiv 1$, $q_1 \equiv 1$, and $(q_w)_{w=2}^5$ chosen to mimic drop rates in ISCHEMIA.

Each observation has a baseline outcome, Y_0 , drawn with replacement from the sample of baseline SAQ summary scores in the ISCHEMIA data. We also construct a standard normal counterpart to it, denoted y_0 .²⁵ The untreated potential outcome path follows:

$$Y_w(0) = Y_0 + a_w + u_w, \quad w = 1, \dots, 5,$$

where $(a_w)_{w=1}^5 = (-0.5, -1, -1.5, -2, -2.5)$ is a downward deterministic trend. To induce serially correlated health shocks, the error term u_w follows an autoregressive model:

$$u_w = \varphi u_{w-1} + \sigma \sqrt{1 - \varphi^2} \tilde{u}_w, \quad u_1 = \sigma \tilde{u}_1, \quad \tilde{u}_w \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

Here $\varphi = 0.6$ is the persistence parameter, and $\sigma = 20$ is the marginal standard deviation of u_{iw} .

Next, define a latent “prognosis index” that determines nonrandom selection into the treatment:

$$H = -\frac{\sqrt{5}}{5} \sum_{w=1}^5 \tilde{u}_w.$$

In particular, for $\beta_H = 0.25$, $\beta_y = 0.5$, and

$$U = \Phi \left(\beta_H H + \beta_y y_0 + [1 - \beta_H^2 - \beta_y^2]^{\frac{1}{2}} \eta \right), \quad \eta \sim \mathcal{N}(0, 1) \implies U \sim \text{Unif}[0, 1],$$

the potential wave of treatment takeup when assigned to $z \in \{0, 1\}$, $S(z) \in \{1, 2, 3, 4, 5, \infty\}$, is:

$$S(z) = \inf \{w \in \{1, \dots, 5\} : U \geq 1 - p_{zw}\},$$

where $\inf \emptyset \equiv \infty$ and p_{0w}, p_{1w} approximate empirical cumulative revascularization rates in the conservative and invasive arms in wave w , respectively. Exposure at wave w equals $T_w = \max\{0, w - S + 1\}$ and write $R_{wt} \equiv \mathbf{1}[T_w = t]$. We then set

$$Y_w(t) = Y_w(0) + \Lambda_t,$$

where the Λ_t are the estimates in Figure 5. Assignment Z is drawn as a Bernoulli with probability 0.5, independently of all potential outcomes and potential treatment variables.

We consider three estimation methods for the Λ_t : the 2SLS estimator used in Figure 5, OLS for the same outcome equation, and a g-method estimator. The g-method estimates come from an additive structural nested mean model, following Robins [1994]. Specifically, we use a GMM

²⁵Specifically, in each simulation, each observation is drawn with a latent baseline rank $V \sim \text{Unif}[0, 1]$. We generate the baseline outcome as $Y_0 = \hat{F}^{-1}(V)$, where $\hat{F}^{-1}$ denotes the empirical quantile function of baseline SAQ summary score in ISCHEMIA. We then define $y_0 = \Phi^{-1}(V)$ where $\Phi(\cdot)$ is the standard normal CDF.

procedure with the moment condition

$$E \left[\mathbf{1}[w > s] I_{w-s} (A_{w-s} - E[A_{w-s} \mid I_{w-s} = 1, X_{w-s}]) \left(Y_w - \sum_{t=1}^5 R_{wt} \Lambda_t \right) \right] = 0,$$

stacked across waves w and $s = 0, \dots, 4$, where A_w is a Bernoulli variable indicating wave w is the period of revascularization, I_w indicates that the period of revascularization is greater or equal to w , and X_w is a set of covariates. For covariates, we consider assignment Z , the baseline variable determining selection y_0 , one-wave lagged outcome, and difference in lagged outcomes. We specify $E[A_{w-s} \mid I_{w-s} = 1, X_{w-s}]$ as a third degree polynomial in y_0 , linear on lags and differenced lags, with intercepts and coefficients for Z that vary with $w - s$.²⁶ To estimate the Λ_t using this method, we stack the above moment conditions and the OLS moment conditions for estimating $E[A_{w-s} \mid I_{w-s} = 1, X_{w-s}]$; we cluster standard errors by individual.

Table B.2 reports absolute bias (relative to true effects), average standard errors, 95% coverage (i.e., the share of replications in which confidence intervals cover the true effects at the 95% level), and power (i.e., the share of replications in which we reject a zero effect at the 5% level) from 1,000 replications. Panel A averages over waves, while panels B and C are for waves 1 and 5, respectively. IV performs well across sample sizes and waves, with relative absolute bias no larger than 5% and coverage around 95% throughout. Nonetheless, it has relatively low precision for small sample sizes and, therefore, small power. OLS, in contrast, is severely biased despite high precision—yielding poor coverage. The smaller bias of the g-method estimator relative to OLS reflects that it partially accounts for selection. However, because selection into treatment depends on the “prognosis index”, it is substantially more biased than IV. While precise in large samples, the g-method standard errors are larger than those of IV in smaller samples.

B.5 Sensitivity to Wave Ignorability

This appendix explores the sensitivity of the cumulative effects estimates in Figure 5 to violations of our strong wave ignorability assumption (Assumption 6). The target causal parameter is the average cumulative effect of immediate exposure for immediate compliers: for wave w ,

$$\Lambda_w \equiv E[Y_w(w) - Y_w(0) \mid T_1(1) > T_1(0)], \quad (\text{B.4})$$

with slight abuse of notation relative to the main text. Under Assumptions 1-3 and strong wave ignorability, Theorem 3 describes a 2SLS estimand that identifies the set of Λ_w (see also the discussion in Footnote 15). The 2SLS framework is useful for sensitivity analysis because violations of the identification assumptions map to a correlation between the instrument and the error term. The next proposition characterizes that mapping for violations of strong wave ignorability.

²⁶For the first decision period (i.e., $w - s = 1$), the lagged outcome is set to baseline and the lagged outcome difference is set to zero; for the second period, the lagged difference is the wave-1 outcome minus baseline; from the third period onward, the lagged difference is the previous outcome minus the two-period lagged outcome.

Proposition B.3. Suppose Assumptions 1-3 hold and let $R_{wt} \equiv \mathbf{1}[T_w = t]$ for each $t \in \{1, \dots, \bar{w}\}$. For each $w \in \{1, \dots, \bar{w}\}$, consider the linear model

$$Y_w = \tilde{\phi} + \sum_{j=2}^{\bar{w}} \tilde{\delta}_j \mathbf{1}[j = w] + \sum_{t=1}^{\bar{w}} \Lambda_t R_{wt} + \gamma_w Z + \tilde{\varepsilon}_w, \quad (\text{B.5})$$

where the set of Λ_t is as in (B.4), $\Lambda_0 \equiv 0 \equiv \gamma_1$,

$$\gamma_w \equiv \sum_{0 \leq t' < t \leq w} E[(Y_w(t) - Y_w(t') - (\Lambda_t - \Lambda_{t'})) \mathbf{1}[T_w(1) = t, T_w(0) = t']], \quad w \in \{2, \dots, \bar{w}\},$$

$$\tilde{\phi} \equiv E[Y_1 - \Lambda_1 R_{11}],$$

$$\tilde{\delta}_j \equiv E\left[Y_j - \sum_{t=1}^{\bar{w}} \Lambda_t R_{jt} - \gamma_j Z\right] - E[Y_1 - \Lambda_1 R_{11}], \quad j \in \{2, \dots, \bar{w}\}.$$

Then, $E[\tilde{\varepsilon}_w | Z] = 0$ for all $w \in \{1, \dots, \bar{w}\}$.

Proof. For $w = 1$, the model simplifies to $Y_1 = \tilde{\phi} + \Lambda_1 R_{11} + \gamma_1 Z + \tilde{\varepsilon}_1$. Because Λ_1 is the average effect for immediate compliers in the first wave and R_{11} is the Bernoulli treatment indicator, the Imbens and Angrist [1994] LATE Theorem implies $E[\tilde{\varepsilon}_w | Z] = 0$ for $\tilde{\phi} = E[Y_1 - \Lambda_1 R_{11}]$ and $\gamma_1 = 0$.

Now fix $w \in \{2, \dots, \bar{w}\}$. By definition of $\tilde{\phi}$ and $\tilde{\delta}_w$, $E[\tilde{\varepsilon}_w] = 0$. Because Z is binary, it suffices to show that $E[Z\tilde{\varepsilon}_w] = 0$, which is equivalent to

$$\gamma_w E[Z] = E\left[Z \left(Y_w - \tilde{\phi} - \tilde{\delta}_w - \sum_{t=1}^{\bar{w}} \Lambda_t R_{wt}\right)\right].$$

It then follows from Assumptions 1 and 3, $Y_w = Y_w(0) + \sum_{t=1}^w (Y_w(t) - Y_w(0)) R_{wt}$, and the definition of $\tilde{\phi}, \tilde{\delta}_w$ that $E[Z\tilde{\varepsilon}_w] = 0$ if, and only if,

$$\begin{aligned} \gamma_w \text{Var}[Z] &= E\left[(Z - E[Z]) \left(\sum_{t=1}^w (Y_w(t) - Y_w(0) - \Lambda_t) R_{wt}\right)\right] \\ &= E[(Z - E[Z])Z] E\left[\sum_{t=1}^w (Y_w(t) - Y_w(0) - \Lambda_t) \mathbf{1}[T_w(1) = t]\right] \\ &\quad + E[(Z - E[Z])(1 - Z)] E\left[\sum_{t=1}^w (Y_w(t) - Y_w(0) - \Lambda_t) \mathbf{1}[T_w(0) = t]\right]. \end{aligned}$$

Therefore, Assumption 2 implies that $E[Z\tilde{\varepsilon}_w] = 0$ if, and only if,

$$\begin{aligned} \gamma_w &= E\left[\sum_{t=1}^w (Y_w(t) - Y_w(0) - \Lambda_t) \mathbf{1}[T_w(1) = t]\right] - E\left[\sum_{t=1}^w (Y_w(t) - Y_w(0) - \Lambda_t) \mathbf{1}[T_w(0) = t]\right] \\ &= \sum_{0 \leq t' < t \leq w} E[(Y_w(t) - Y_w(t') - (\Lambda_t - \Lambda_{t'})) \mathbf{1}[T_w(1) = t, T_w(0) = t']]. \end{aligned}$$

□

Proposition B.3 parameterizes violations of strong wave ignorability in terms of the γ_w coefficients, in line with earlier sensitivity methods for linear instrumental variable models (e.g., Cinelli and Hazlett [2025]; Conley, Hansen and Rossi [2012]; Masten and Poirier [2021]). Specifically, γ_w is generally nonzero when wave- w compliers moved from some treatment margin t' to another margin t (with $T_w(1) = t$ and $T_w(0) = t'$) have systematically different incremental effects than that of immediate compliers, $\Lambda_t - \Lambda_{t'}$. The proposition shows that these deviation parameters arise as direct effects of the instrument Z in the stacked model from Theorem 3.

We apply this result to the general sensitivity framework of Conley, Hansen and Rossi [2012] (henceforth, CHR) to explore robustness to nonzero γ_w . The CHR approach builds on the observation that if the vector of γ_w was known, a 2SLS of the modified outcome $Y_w - \gamma_w Z$ on the set of R_{wt} instrumented by Z and $Z \times \mathbf{1}[j = w], j \in \{2, \dots, \bar{w}\}$ would estimate the set of Λ_t . In practice, however, the set of γ_w is unknown. We follow CHR in postulating a bounded set of possible values for the vector of γ_w , $\Gamma \subseteq \mathbb{R}^{\bar{w}}$, with $\gamma_1 \equiv 0$. Then, for each point in Γ , we can construct confidence intervals for each Λ_t . Taking the union of the confidence intervals over Γ guarantees the prescribed coverage as long as the true vector of γ_w lies on Γ . CHR approximate such a union of confidence intervals by discretizing Γ and taking unions over the grid points.

We consider two approaches to specify Γ , both parametrized by bounds on the possible incremental effect deviations underlying each γ_w :

$$\Delta_w^{t,t'} \equiv E[Y_w(t) - Y_w(t') \mid T_w(1) = t, T_w(0) = t'] - (\Lambda_t - \Lambda_{t'}).$$

Specifically, we assume that $|\Delta_w^{t,t'}|$ is at most a $\rho \in [0, 1]$ fraction of the absolute difference in cumulative effects:

$$|\Delta_w^{t,t'}| \leq \rho |\Lambda_t - \Lambda_{t'}|. \quad (\text{B.6})$$

Our first approach is conservative. Because the probabilities $P[T_w(1) = t, T_w(0) = t']$ are not identified, it takes the worst-case by setting $|\gamma_w| \leq \max\{\rho |\Lambda_t - \Lambda_{t'}| : 0 \leq t' < t \leq w\}$. The second approach instead sets

$$|\gamma_w| \leq \sum_{0 \leq t < w} (E[\mathbf{1}[T_w = t] \mid Z = 0] - E[\mathbf{1}[T_w = t] \mid Z = 1]) \rho |\Lambda_w - \Lambda_t|,$$

which are valid bounds on γ_w under IMCO (Assumption 4) and (B.6) since then


$$\begin{aligned} \gamma_w &= \sum_{0 \leq t < w} P[T_w(1) = w, T_w(0) = t] \Delta_w^{w,t}, \\ P[T_w(1) = w, T_w(0) = t'] &= E[\mathbf{1}[T_w = t] \mid Z = 0] - E[\mathbf{1}[T_w = t] \mid Z = 1]. \end{aligned}$$

In practice, because the target parameters Λ_t are not identified when strong wave ignorability doesn't hold, we use the estimates in Figure 5 to bound $|\Delta_w^{t,t'}|$. This yields bounds that are

compatible with the magnitude of effects in the ISCHEMIA setting.


Figure B.3 reports union-of-confidence-interval estimates under the two sensitivity calibrations, for $\rho \in \{0, 0.25, 0.5\}$ (the $\rho = 0$ results replicate the baseline Figure 5 estimates). The results suggest that the estimated average cumulative effects of revascularization are not overly sensitive to violations of wave ignorability, especially for the SAQ summary score. Under the IMCO-based calibration, the data supports the main results for violations up to $\rho = 0.25$. At this value, the null of no effect is rejected at the 5 percent level for both outcomes in all follow-up waves except the last. For the SAQ summary score, the first three follow-up effects remain statistically different from zero even when $\rho = 0.5$. The more conservative max-bound calibration leads to wider intervals, as expected, while still suggesting positive effects on the SAQ summary score at intermediate follow-up horizons. Taken together, these results build confidence in the robustness of our baseline findings.

Figure B.1: Differences in Exposure CDF by Treatment Assignment Status




Notes: This figure plots estimated differences in CDFs between the treatment and control groups for exposure levels T_w . For each exposure year t , the estimate comes from the regression of the indicator $1[T_w \leq t]$ on treatment assignment status Z . Vertical bars denote the 95 % robust confidence intervals. This figure is based on observations from wave 5.

Figure B.2: Permitted and Ruled-Out Paths for Incremental Effects in [Bowden et al. \[2025\]](#)



Notes: The first three series in this figure illustrate possible paths for incremental effects λ_t under the [Bowden et al. \[2025\]](#) parameterization, (B.1), for $\beta = 5$. The first series sets $\alpha = 0.5$, the second series sets $\alpha = -0.5$ and the third series sets $\alpha = 0$. The fourth series in this figure is ruled out by Equation (B.1).

Figure B.3: Sensitivity Analysis for Average Cumulative Causal Effects



Notes: This figure reports sensitivity results for the average cumulative effect of immediate exposure on SAQ Summary Score and SAQ Angina Frequency for immediate compliers. The top panel shows the IMCO-based results for $\rho \in \{0, 0.25, 0.5\}$, while the bottom panel shows results based on max-bounds for the same values of ρ . For each value of $\rho > 0$, the set of possible values for the vector of γ_w is discretized using 5 grid points for each γ_w , except γ_1 which is set to zero. The point estimates reported here equal the ones in the main text (Figure 5). Appendix B.5 describes the sensitivity analysis procedure in detail.

Table B.1: Comparison with [Bowden et al. \[2025\]](#)

	Restricted $\alpha \geq 0$			Unrestricted α		
	Pooled	SAQ SS	SAQ AFS	Pooled	SAQ SS	SAQ AFS
	(1)	(2)	(3)	(4)	(5)	(6)
α	0	0	0	-0.39	-0.37	-0.39
	—	—	—	(0.109)	(0.119)	(0.111)
	—	—	—	[0.0003]	[0.0018]	[0.0005]
β	3.71	3.99	3.83	4.69	4.77	4.92
	(0.484)	(0.579)	(0.493)	(0.563)	(0.651)	(0.585)
Overid. Tests	26.4	6.95	17.3	15.1	0.57	4.94
	[0.0018]	[0.138]	[0.0017]	[0.0573]	[0.904]	[0.177]

Notes: This table reports efficiently-weighted two-step GMM estimates based on the moment condition in Equation (B.3). Column 1 reports the results from a model imposing $\alpha = 0$ when the outcomes, SAQ Summary Score and SAQ Angina Frequency Score, are pooled. Columns 2 and 3 report results using the same model for each outcome separately. Columns 4, 5, and 6 report results for the same outcomes, respectively, for a model where α is unrestricted. Estimates were computed with controls for baseline angina frequency scores and enrollment region. Individual-clustered standard errors are reported in parentheses, and p-values for overidentification tests are reported in square brackets.

Table B.2: Simulation Results Comparing IV, OLS, and G-method

Estimator	N	Abs.	Standard	Coverage	Power
		Relative	Error		
		Bias			
		(1)	(2)	(3)	(4)
Panel A: Average					
IV	5000	0.029	1.810	0.953	0.585
	1500	0.027	3.306	0.954	0.281
	500	0.047	5.775	0.950	0.126
OLS	5000	1.667	1.186	0.035	1.000
	1500	1.668	2.163	0.220	0.938
	500	1.687	3.745	0.517	0.753
G-est	5000	1.127	3.433	0.527	0.096
	1500	1.172	6.308	0.713	0.062
	500	1.325	11.074	0.850	0.055
Panel B: Wave 1					
IV	5000	0.009	1.134	0.955	0.985
	1500	0.011	2.071	0.958	0.619
	500	0.039	3.599	0.950	0.249
OLS	5000	1.436	0.712	0.000	1.000
	1500	1.433	1.301	0.000	1.000
	500	1.423	2.258	0.150	1.000
G-est	5000	0.859	0.658	0.000	0.179
	1500	0.868	1.199	0.076	0.079
	500	0.862	2.079	0.515	0.052
Panel C: Wave 5					
IV	5000	0.064	3.017	0.953	0.231
	1500	0.042	5.511	0.955	0.125
	500	0.046	9.689	0.947	0.075
OLS	5000	1.575	2.032	0.168	0.998
	1500	1.589	3.701	0.619	0.733
	500	1.672	6.391	0.837	0.345
G-est	5000	0.902	7.567	0.946	0.066
	1500	1.084	13.987	0.950	0.063
	500	1.600	24.825	0.948	0.068

Notes: This table reports simulation results comparing the finite sample performance of IV, OLS, and g-method for 1,000 replications. Sample sizes considered are 500, 1,500, and 5,000. Column 1 reports absolute bias relative to true causal effects. Column 2 reports standard errors averaged across replications. Column 3 reports 95% nominal coverage averaged across replications (i.e., the share of replications in which confidence intervals cover the true effects at the 95% level). Column 4 reports the share of replications in which the null of no effect is rejected at the 5% level. Panel A reports results averaged across waves, while panels B and C report results for waves 1 and 5, respectively. Appendix B.4 details the data-generating process and estimators.