

NBER WORKING PAPER SERIES

HOW TO EXAMINE EXTERNAL VALIDITY WITHIN AN EXPERIMENT

Amanda E. Kowalski

Working Paper 24834

<http://www.nber.org/papers/w24834>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

July 2018, Revised October 2020

I thank Neil Christy, Simon Essig Aberg, Pauline Mourot, Srajal Nayak, Ljubica Ristovska, Sukanya Sravasti, Rae Staben, and Matthew Tauzer for excellent research assistance. NSF CAREER Award 1350132 provided support. I thank Magne Mogstad, Jeffrey Smith, Edward Vytlačil, and anonymous referees for helpful feedback. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Amanda E. Kowalski. 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.

How to Examine External Validity Within an Experiment
Amanda E. Kowalski
NBER Working Paper No. 24834
July 2018, Revised October 2020
JEL No. C9,C93,H0

ABSTRACT

A fundamental concern for researchers who analyze and design experiments is that the estimate obtained from the experiment might not be externally valid for other policies of interest. Researchers often attempt to assess external validity by comparing data from an experiment to external data. In this paper, I discuss approaches from the treatment effects literature that researchers can use to begin the examination of external validity internally, within the data from a single experiment. I focus on presenting the approaches simply using stylized examples.

Amanda E. Kowalski
Department of Economics
University of Michigan
611 Tappan Ave.
Lorch Hall 213
Ann Arbor, MI 48109-1220
and NBER
aekowals@umich.edu

1 Introduction

The traditional reason that a researcher runs an experiment is to address selection into treatment. For example, a researcher might be worried that individuals with better outcomes regardless of treatment are more likely to select into treatment, so the simple comparison of treated to untreated individuals will reflect a selection effect as well as a treatment effect. By running an experiment, the reasoning goes, a researcher isolates a single treatment effect by eliminating selection.

However, there is still room for selection within an experiment. In many experiments, some lottery losers receive treatment and some lottery winners forgo treatment. In this paper, I mainly consider experiments with “two-sided noncompliance,” in which both occur. In these experiments, individuals participate in a lottery. The lottery winners are in the intervention arm. They receive an intervention that affects selection into treatment. The lottery losers are in the control arm, and they do not receive the intervention. However, all individuals can select to receive or forgo the treatment.

Some researchers view this type of selection as immaterial, and they discard information on which individuals select into treatment by focusing on the comparison of all lottery winners to all lottery losers. Other researchers view this type of selection as a nuisance, and they alter information on which individuals select into treatment by encouraging all individuals to comply with random assignment. I view this type of selection as a useful source of information that can be combined with assumptions to learn about the external validity of an experiment, in the tradition of Heckman et al. (2000). The ability to learn from information on selection gives a researcher new reasons to run an experiment. An experiment is no longer a tool that eliminates selection; it is a tool that identifies selection. Furthermore, under ancillary assumptions, an experiment is no longer a tool that isolates a single treatment effect; it is a tool that identifies a range of heterogeneous treatment effects.

An experiment reconceived as a tool that identifies heterogeneous treatment effects can inform external validity in two senses. In a global sense, if treatment effects vary across groups within an experiment, then there is no single treatment effect that is valid for all policies. In a local sense, if treatment effects vary across groups within an experiment, then there could be a treatment effect for one group that is valid for a different policy. Information on selection, institutional features of the experiment, and data on covariates can inform assumptions about *which* treatment effect within an experiment might be valid for a specific policy of interest. The specific policy of interest could be within the context of the experiment or within an entirely different context, although external validity in other contexts requires stronger assumptions.

In this paper, I discuss recent advances from the treatment effects literature that re-

searchers can use to begin examination of external validity within an experiment. I do not break new ground in terms of methodology or substantive application, and I do not aim to be comprehensive. Rather, I aim to present some existing methods simply using stylized examples, making them accessible to researchers who analyze and design experiments.

One of the virtues of experiments is that standard analysis is straightforward and relies on well-known assumptions. I proceed under the well-known local average treatment effect (LATE) assumptions of independence and monotonicity proposed by Imbens and Angrist (1994). Vytlacil (2002) constructs a model of selection into treatment that assumes no more than the LATE assumptions, and I use it as the foundation for my analysis. The model can be interpreted as a generalized Roy (1951) model of the marginal treatment effect (MTE) introduced by Björklund and Moffitt (1987), in the tradition of Heckman and Vytlacil (1999, 2001b, 2005), Carneiro et al. (2011), Brinch et al. (2017), and Cornelissen et al. (2018). Therefore, the model that serves as the foundation for my analysis also serves as the foundation for the LATE and MTE approaches. I do not present the model here, and I instead focus on depicting its implications graphically. I refer interested readers to my exposition of the model in Kowalski (2020b), which contains an updated version of almost all content from Kowalski (2016) as applied to the Oregon Health Insurance Experiment. I have included the remaining content in Kowalski (2020a) as applied to the Canadian National Breast Screening Study. I also provide a Stata command (Kowalski et al., 2018).

In Section 2, I depict information from a hypothetical experiment run by a health insurer. I begin by presenting information required for standard analysis of an experiment. Next, I present additional information available under the LATE assumptions that is often unreported. This additional information consists of shares and outcomes of always takers, compliers, and never takers, using the terminology of Angrist et al. (1996), obtained following Imbens and Rubin (1997), Katz et al. (2001), Abadie (2002), and Abadie (2003). Although there can be many types of heterogeneity within an experiment, I focus on heterogeneity across always takers, compliers, and never takers and discuss how it can inform external validity with respect to particular policies of interest. For this purpose, it is useful to think about the intervention within the experiment as a particular policy that expands the fraction treated in the intervention arm and maintains the fraction treated in the control arm. Policies that would contract the fraction treated below the level of treatment in the control arm would induce treatment effects on always takers, and policies that would expand the fraction treated above the level of treatment in the intervention arm would induce treatment effects on never takers. In this way, heterogeneity across always takers, compliers, and never takers informs external validity with respect to policies that expand and contract the fraction treated.

In Section 3, I depict a test for heterogeneous selection as the fraction treated expands and contracts that uses a subset of the information available on compliers and never takers under the LATE assumptions. In Kowalski (2020b), I refer to this test as the “untreated outcome test,” and my innovation is in the interpretation—I show that under the LATE assumptions alone, it identifies one specific instance of how selection into treatment varies as the fraction treated varies, a concept that generalizes the notion of “selection bias” (Angrist (1998); Heckman et al. (1998)), which is not identified under the LATE assumptions alone. This test is equivalent to tests proposed by Guo et al. (2014) and Black et al. (2017), and generalized by Mogstad et al. (2018). It is also similar to the Bertanha and Imbens (2014) test proposed for the regression discontinuity context and to the Einav et al. (2010) test in the insurance literature. This test for heterogeneous selection is a natural precursor to a test for treatment effect heterogeneity because if outcomes do not differ across groups due to heterogeneous selection effects, then differences could reflect heterogeneous treatment effects.

In Section 4, I depict a test for treatment effect heterogeneity equivalent to a test proposed by Brinch et al. (2017) and applied in Kowalski (2020a). Brinch et al. (2017) conduct this test under two ancillary assumptions. As I show in Kowalski (2020a), it is possible to conduct the test under only one of their ancillary assumptions; either one will suffice. I implement the test using the more justifiable assumption, and I discuss how data on covariates can be used to assess its plausibility. The assumption implies an upper or lower bound on the average treatment effect for always takers, and an additional assumption can imply an upper or lower bound on the average treatment effect for never takers. Using the insurance experiment, I discuss how the implied bounds on always and never takers yield specific implications for how, and how not, to scale up the treatment.

In Section 5, I demonstrate how stronger assumptions of Hausman (1978); Heckman (1979); Willis and Rosen (1979); Angrist (2004); Huber (2013); Bertanha and Imbens (2014); Guo et al. (2014); Black et al. (2017) and Brinch et al. (2017) yield estimates of treatment effects in lieu of bounds. I also discuss how data can inform the plausibility of the assumptions. Because I do not need stronger assumptions to draw meaningful conclusions about external validity from the insurance experiment, I introduce a stylized clinical trial for hip replacement surgery to illustrate the implications of stronger assumptions for external validity.

The approaches that I discuss here do not supplant other approaches to examine external validity such as subgroup analysis and LATE-reweighting (Hotz et al., 2005; Angrist and Fernandez-Val, 2013). However, I demonstrate that subgroup analysis alone would not be sufficient to reach the same conclusions in the surgery trial. I conclude by discussing implications for experimental design in Section 6.

2 An Experiment under the LATE Assumptions

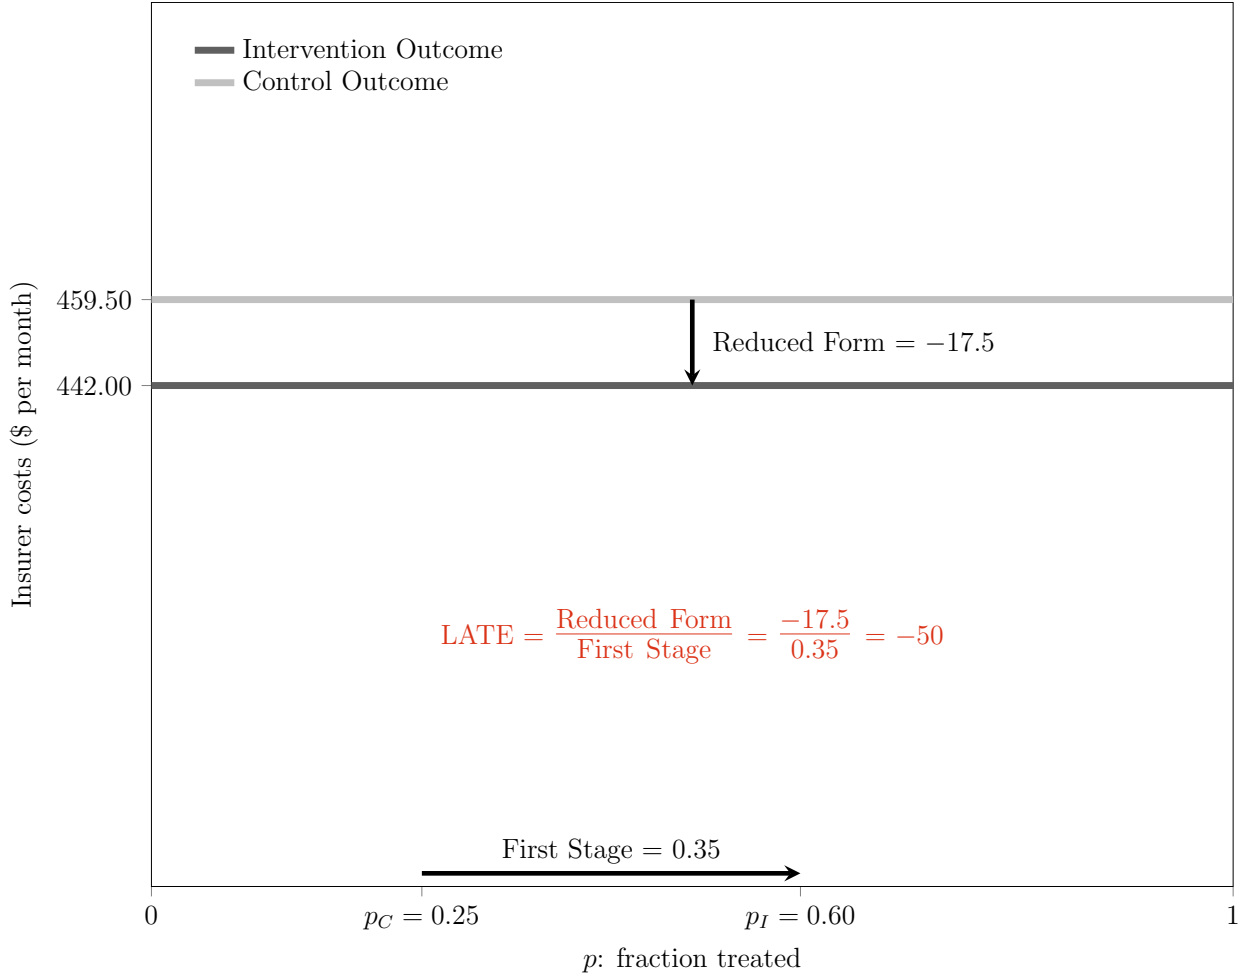
Consider the following stylized example. Last year, a health insurer rolled out a new wellness plan at one pilot site. At that site, for the same premium, beneficiaries could choose between a traditional plan and a wellness plan that offered the same benefits plus gym membership reimbursement. Based on last year’s financials from the pilot site, the insurer thinks that the wellness plan lowers the average costs that it pays on behalf of beneficiaries. On this basis, it is considering rolling out the wellness plan as an option for beneficiaries at all its sites. It is also considering eliminating the traditional plan entirely and enrolling all beneficiaries in the wellness plan at all sites. However, only 25% of beneficiaries at the pilot site chose the wellness plan, and the insurer recognizes that those enrollees might not be representative of all beneficiaries at all sites.

Before pursuing either wellness plan expansion option, the insurer decides to gather more evidence by running a randomized experiment at the pilot site. It mails an informational brochure that promotes the wellness plan to lottery winners but does not mail a brochure to lottery losers. The treatment is enrollment in the wellness plan. There is two-sided noncompliance, as some lottery winners enroll in the traditional plan and some lottery losers enroll in the wellness plan. The insurer observes whether each individual wins the lottery and whether each individual enrolls in the wellness plan. It also observes an outcome for each individual: the average monthly health care costs that it pays on behalf of each beneficiary.

Standard analysis of an experiment begins by comparing average outcomes in the intervention and control arms. I depict these outcomes in Figure 1. As shown, the intervention seems to decrease cost because the average outcome (cost per month) is \$17.5 lower in the intervention arm than it is in the control arm. This difference in average outcomes is often called the “reduced form,” as labeled along the vertical axis, or the “intent to treat (ITT).” It gives an estimate of the impact of the intervention (the mailing of the informational brochure) on the outcome. In experiments with two-sided noncompliance, lottery status does not perfectly determine treatment, so the reduced form does not give an estimate of the impact of the treatment (enrollment in the wellness plan) on the outcome. Calculation of the reduced form does not even require data on treatment. Some researchers report only the reduced form.

Standard analysis of an experiment next compares the probability of treatment in the control and intervention arms. By the LATE independence assumption, lottery status is independent of treatment, so I can depict the probability of treatment in the control and intervention arms along the same horizontal axis in Figure 1. As shown, p_C represents the probability of treatment in the control arm, and p_I represents the probability of treatment in the intervention arm. The difference $p_I - p_C$ is often called the “first stage.” It gives an

Figure 1: Average Outcomes in Intervention and Control Arms Under LATE Assumptions



estimate of the impact of winning the lottery on the fraction treated p . In experiments with two-sided noncompliance, the first stage is less than one. In Figure 1, 25% of lottery losers and 60% of lottery winners receive treatment, so the first stage implies that winning the lottery increases the fraction treated by 35 percentage points. That is, the mailing of the informational brochure increases enrollment in the wellness plan by 35 percentage points.

To obtain an estimate of the impact of the treatment on the outcome, standard analysis of an experiment divides the reduced form by the first stage. This quotient gives the local average treatment effect (LATE) of Imbens and Angrist (1994). Many researchers report the LATE as the single treatment effect that the experiment isolates. Under the LATE assumptions, the LATE gives the average treatment effect for “compliers,” individuals whose treatment status is determined by their random assignment, in the terminology of Angrist et al. (1996). In Figure 1, the LATE of -50 implies that the wellness plan reduces health

care costs by \$50 per month on average among individuals who take up the plan if and only if the insurer mails them the informational brochure.

The evidence from the standard analysis of the experiment looks promising. The wellness plan decreases costs for compliers. But does it decrease costs for all individuals? Is the LATE useful to the insurer in deciding between the two wellness plan expansion options?

To inform the external validity of the LATE with respect to the expansion options, I consider two groups of individuals to which the LATE need not apply: “always takers” who take up treatment regardless of random assignment and “never takers” who do not take up treatment regardless of random assignment, in the terminology of Angrist et al. (1996). Under this terminology, the LATE monotonicity assumption rules out “defiers” who take up treatment if and only if they lose the lottery (Balke and Pearl, 1993; Angrist et al., 1996), so experiments with two-sided noncompliance involve only always takers, compliers, and never takers. Researchers cannot label each individual as an always taker, complier, or never taker: lottery winners who take up treatment could be always takers or compliers; lottery losers who do not take up treatment could be compliers or never takers. However, researchers can label lottery losers who take up treatment as always takers and lottery winners who do not take up treatment as never takers.

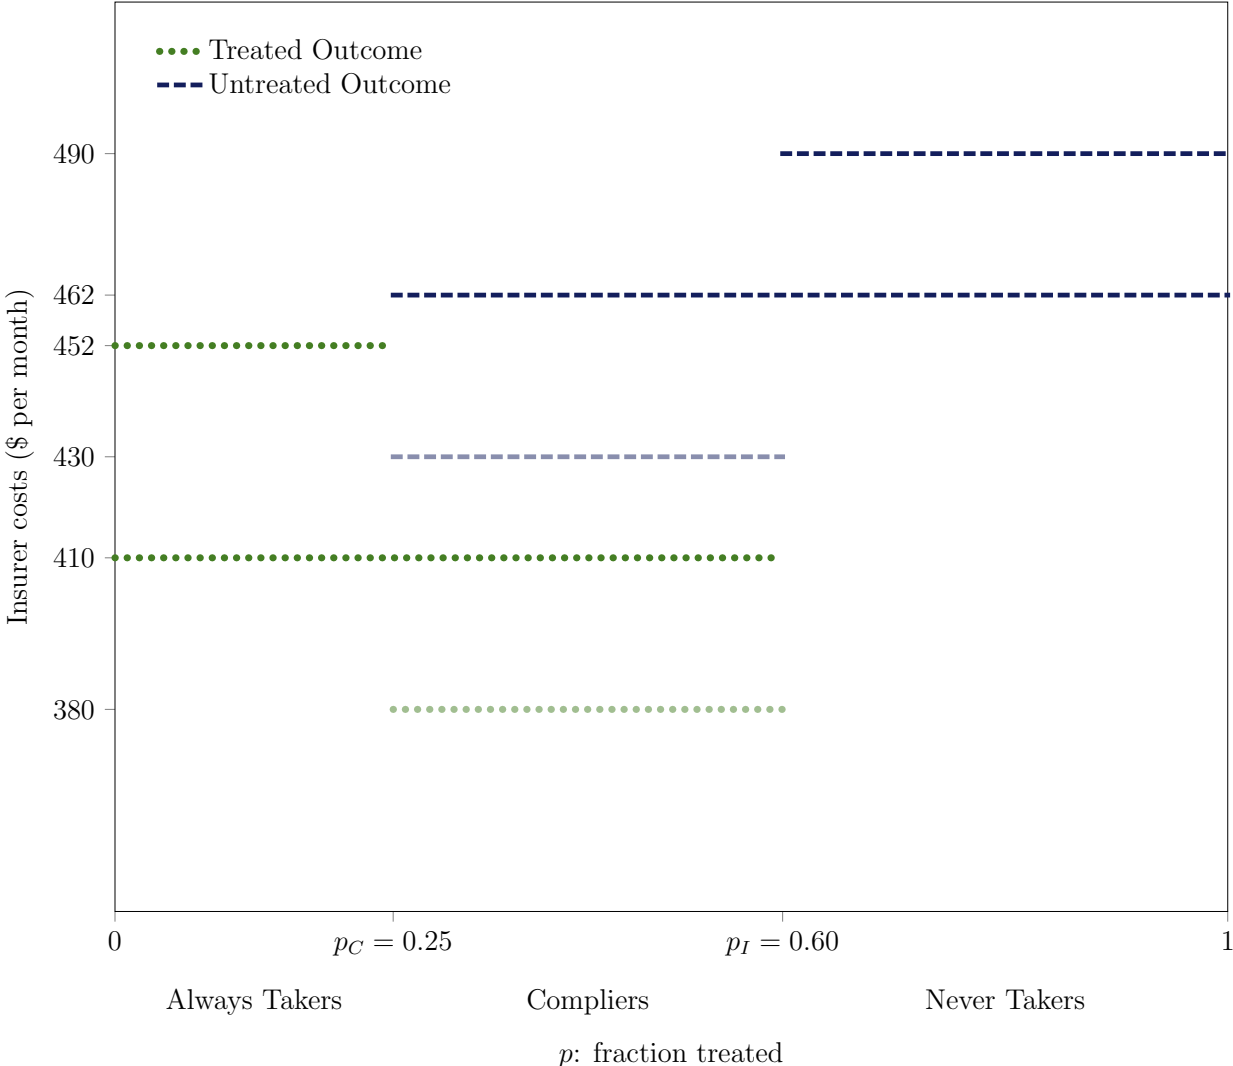
The ability to identify some individuals as always or never takers allows researchers to learn more about compliers. The LATE independence assumption implies that lottery status is independent of whether an individual is an always taker, complier, or never taker. Therefore, the observed share of treated lottery losers yields an estimate of the share of always takers in the full sample, and the observed share of untreated lottery winners yields an estimate of the share of never takers in the full sample. Furthermore, because always and never takers do not change their treatment status based on their lottery status, their average outcomes should not depend on their lottery status. Using the shares and average outcomes of always takers and never takers, researchers can estimate the average outcomes of treated and untreated compliers, as demonstrated by Imbens and Rubin (1997), Katz et al. (2001), Abadie (2002), and Abadie (2003).¹

To illustrate the calculation of the average outcomes of always takers, compliers, and never takers graphically, I continue the insurance example in Figure 2. As shown by Imbens and Rubin (1997) and Vytlacil (2002), the LATE assumptions imply an ordering from always takers to compliers to never takers. Consistent with this ordering, I label ranges of the horizontal axis that correspond to the shares of each group. On the left, the fraction p_C of

¹These approaches also allow researchers to estimate the *distributions* of outcomes of treated and untreated compliers, which paves the way for examination of treatment effect heterogeneity *within* compliers, as in Heckman et al. (1997). Here, I focus on treatment effect heterogeneity across always takers, compliers, and never takers.

individuals who receive treatment regardless of their lottery status are always takers. In the middle, the fraction $(p_I - p_C)$ of individuals who receive treatment if and only if they win the lottery are compliers. On the right, the remaining fraction $(1 - p_I)$ of individuals who do not receive treatment regardless of their lottery status are never takers. The intuition behind the ordering is clear if we interpret the randomized intervention as a policy change that occurs within the intervention arm. In the experiment as a whole, always takers are the individuals who receive treatment under the existing policy, compliers are the new individuals who can be induced to receive treatment by the policy change, and never takers are the remaining individuals who could be induced to receive treatment by a future policy change. In the insurance example, the policy change is the mailing of the informational brochure and the treatment is enrollment in the wellness plan.

Figure 2: Average Treated and Untreated Outcomes in Intervention and Control Arms and Average Treated and Untreated Outcomes of Compliers Under LATE Assumptions




Along the vertical axis of Figure 2, I plot the average treated and untreated outcomes of the intervention and control arms over the relevant ranges of the horizontal axis. As shown, the average treated outcome in the intervention arm is \$410, which represents a weighted average of the treated outcomes of always takers and compliers. The average treated outcome in the control arm is \$452, which represents the average treated outcome of always takers. Because always takers make up 25% of the full sample and always takers combined with compliers make up 60% of the full sample, the average treated outcome of compliers is \$380 ($= (0.6/(0.6-0.25))*410 - (0.25/(0.6-0.25))*452$), as depicted in light shading. Similar logic using the untreated outcomes implies that the average untreated outcome of never takers is \$490 and that the average untreated outcome of compliers is \$430 ($= ((1-0.25)/(0.6-0.25))*462 - ((1-0.60)/(0.6-0.25))*490$), as depicted in light shading. Researchers who would like to replicate these calculations can use the Stata command *mtebinary* (Kowalski et al., 2018).

As shown by Imbens and Rubin (1997), the LATE is equal to the difference in the average treated and untreated outcomes of compliers. Accordingly, in Figure 3, I depict an arrow that gives the sign and magnitude of the LATE. However, I could have obtained the LATE using Figure 1 alone, even if my data would not allow me to construct Figures 2 and 3. Construction of Figures 2 and 3 requires data on outcomes by lottery status *and* treatment. In contrast, construction of Figure 1 only requires data on outcomes by lottery status (for the reduced form) and data on treatment by lottery status (for the first stage). As shown by Angrist (1990) and Angrist and Krueger (1992), it is possible to obtain the LATE via the Wald (1940) approach using separate datasets for the reduced form and first stage. Because the LATE can be obtained using limited data, it stands to reason that it does not capture all available information. Thus, Figure 3 provides additional information relative to Figure 1.

Using the additional information depicted in Figure 3, I emphasize that always and never takers are distinct groups to which the LATE need not apply. In the insurance example, these groups are sizeable. Furthermore, the average treated outcome of always takers and the average untreated outcome of never takers are known. The average untreated outcome of always takers and the average treated outcome of never takers are not known. If they could be identified, then it would be possible to estimate the average treatment effect for each group as the difference between their average treated and untreated outcomes. Similarly, if they could be bounded, then it would be possible to bound the average treatment effect for each group, as discussed by Imbens and Rubin (1997). Such bounds could be implied by natural bounds on the range of outcomes in the tradition of Robins (1989), Manski (1990), and Balke and Pearl (1997), or they could be implied by ancillary assumptions.

Even in the absence of ancillary assumptions, a researcher examining Figure 3 might

Figure 3: Average Outcomes of Always Takers, Compliers, and Never Takers Under LATE Assumptions



conjecture that the LATE is not equal to the average treatment effect for always and never takers, and is thus not externally valid in a global sense. If the treatment effect were the same for everyone, then why do some individuals select into treatment while others do not, even within the same arm of the experiment? Also, why do the average treated and untreated outcomes differ across always takers, compliers, and never takers? Is it because their average treatment effects differ? In the next sections, I interpret the implications of these differences for external validity, first under the LATE assumptions, and then under ancillary assumptions.

3 Test for Heterogeneous Selection under the LATE Assumptions

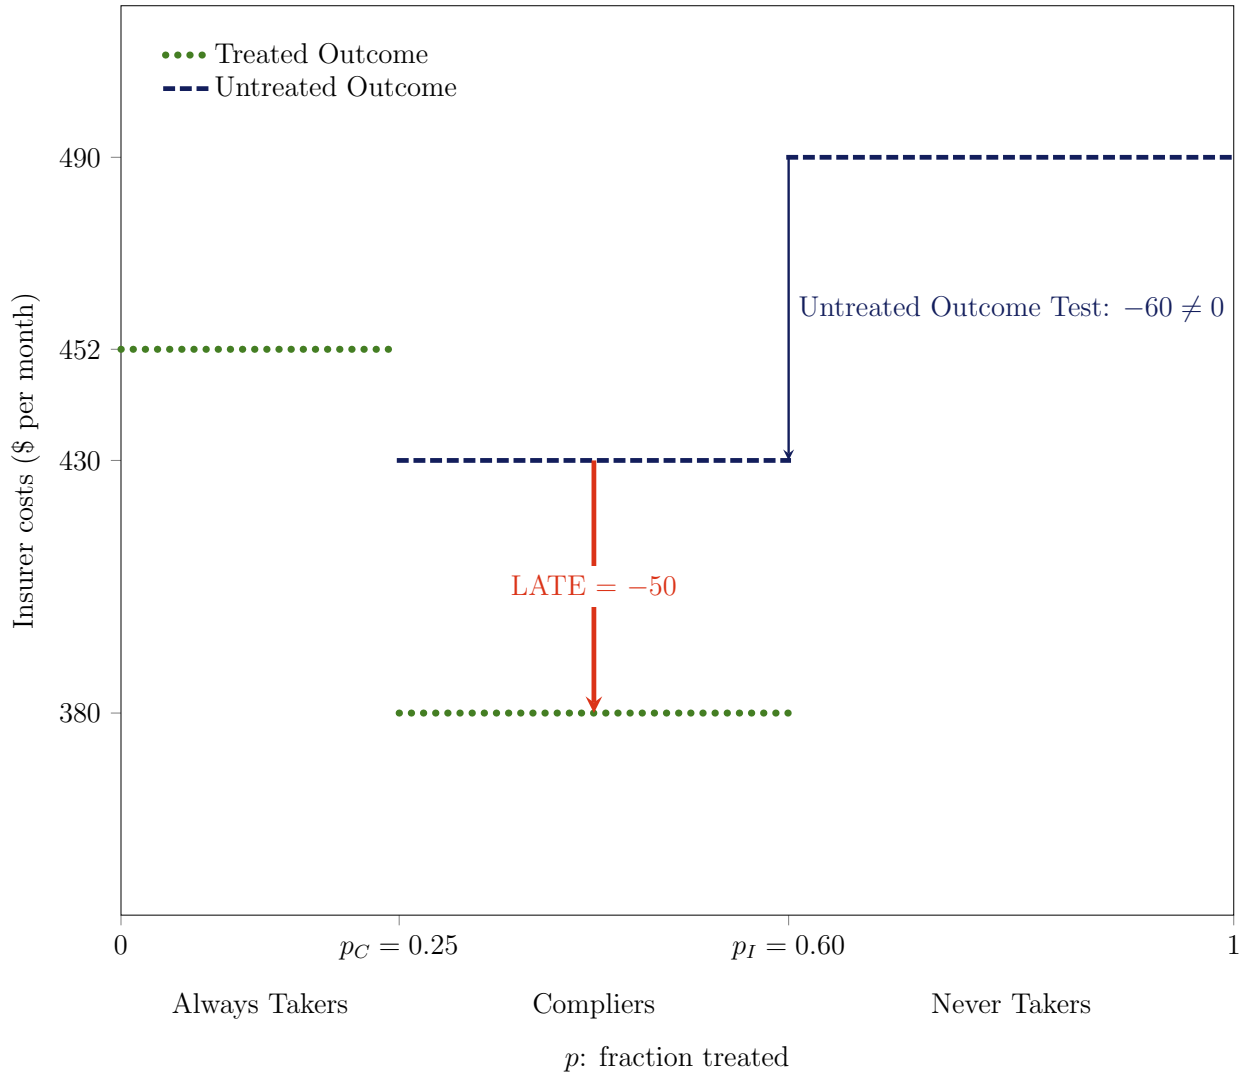
Consider the test of the null hypothesis that the difference between the average untreated outcomes of compliers and never takers is equal to zero. This test is equivalent or similar to tests proposed by Bertanha and Imbens (2014); Guo et al. (2014); Black et al. (2017), and generalized by Mogstad et al. (2018). It is also related to the “cost curve” test of Einav et al. (2010) from the insurance literature when the untreated outcome is uninsured costs (or, in my stylized example, costs under the traditional plan). In Kowalski (2020b), I refer to this as the “untreated outcome test,” and I provide a novel interpretation for it. I show that under the LATE assumptions, it identifies a specific instance of heterogeneous selection as the fraction treated p varies. The logic behind this interpretation is simple. Untreated compliers and never takers do not receive treatment. Therefore, a difference in their outcomes cannot reflect a difference in the treatment effect. It can only reflect a difference in selection.

Continuing the insurance example, Figure 4 shows that the average untreated outcome of compliers is \$60 lower than the average outcome of never takers. If this difference is statistically different from zero, then the test rejects selection homogeneity. Statistical significance can be obtained via a variety of approaches, including bootstrapping, which can account for estimation of the average outcomes as well as their difference.

The sign of the untreated outcome test statistic, the difference in average untreated outcomes between compliers and never takers, indicates whether selection is positive or negative. In Figure 4, a negative untreated outcome test statistic indicates negative selection, such that individuals with higher average costs (never takers) select into treatment after individuals with lower average costs (compliers). When the treatment represents enrollment in an insurance plan, it is customary to refer to negative selection as “advantageous selection” and positive selection as “adverse selection.” Although insurers often face adverse selection, there is a rationale for advantageous selection in the stylized insurance example: individuals with lower health care costs in the traditional plan have better health behaviors and are therefore more likely to enroll in the wellness plan that offers gym membership reimbursement.

There can be positive selection on some outcomes and negative selection on others within the same experiment. In the insurance example, consider a measure of how often beneficiaries go to the gym. Even though there is negative selection on health care costs, there could be positive selection on how often beneficiaries go to the gym. This positive selection on gym behavior could help to explain negative selection on health care costs. It could also offer a rationale for why the treatment effect on health care costs could vary with the fraction treated p . For example, there might be less scope for the wellness plan to reduce health care

Figure 4: Untreated Outcome Test Rejects
 Test Statistic Shows Negative Selection
 Under LATE Assumptions



costs for beneficiaries who go to the gym more frequently.


The analogous *treated* outcome test, which tests the null hypothesis that the difference between the average treated outcomes of always takers and compliers is equal to zero, has also been considered by the literature that examines tests similar or equivalent to the untreated outcome test (Bertanha and Imbens, 2014; Guo et al., 2014; Black et al., 2017). In the insurance literature, the treated outcome test is related to the “cost curve” test of Einav et al. (2010) when the treated outcome is *insured* costs (or, in my stylized example, costs under the wellness plan). In Kowalski (2020b), I emphasize that the treated outcome test does not isolate heterogeneous selection. For this reason, I do not recommend running the treated outcome test, but I discuss it here to illustrate why it does not isolate heterogeneous selection.

Continuing the insurance example, consider the treated outcome test depicted in Figure 5. It shows that the average outcome of always takers is \$72 higher than the average treated outcome of compliers. Assume that this difference is statistically different from zero, so the treated outcome test rejects. This result could be entirely due to heterogeneous selection from always takers to compliers, which would be the case if the average treatment effects for both groups were equal. In that case, the average treatment effect for always takers would be equal to the LATE of -\$50, and the average untreated outcome of always takers would be \$402 ($= 452 - 50$). Alternatively, the result of the treated outcome test could be entirely due to treatment effect heterogeneity from always takers to compliers, which would be the case if there were no selection heterogeneity across the two groups. In that case, the average untreated outcome of always takers would be equal to the average untreated outcome of compliers of \$430, and the average treatment effect for always takers would be an increase of \$22 ($= 452 - 430$). It is also possible that the treated outcome test could reflect a combination of nonzero selection heterogeneity and nonzero treatment effect heterogeneity, which would be the case if the average untreated outcome of always takers were anything other than \$402 or \$430. As shown, the treated outcome test reflects selection heterogeneity *plus* treatment effect heterogeneity, while the untreated outcome test can only reflect selection heterogeneity.

It is tempting to think that the treated outcome test should have the same implications as the untreated outcome test because the distinction between treated and untreated should be immaterial. However, as I discuss in Kowalski (2020b), the distinction between treated and untreated is material to the definition of the treatment and thus to the definitions of the treatment effect, treatment effect heterogeneity, and selection heterogeneity. The treatment effect is defined as the treated outcome minus the untreated outcome, not the untreated outcome minus the treated outcome. Therefore, the treatment effect has magnitude *and* direction, which is why I depict the local average treatment effect with an arrow in the figures. Renaming the treated the untreated and vice versa might seem inconsequential, but such a swap would change the direction of the arrow. The swap would also change the definitions of selection and treatment effect heterogeneity. The treated outcome test would detect only selection, and the untreated outcome test would detect a combination of selection and treatment effect heterogeneity, creating a different but no less material distinction between the tests.

The distinction between treated and untreated outcomes forms the foundation for the tests for heterogeneous treatment effects that I present in the next sections. The tests use differences in untreated outcomes to motivate ancillary assumptions. The assumptions purge selection heterogeneity from treated outcomes to isolate treatment effect heterogeneity.

Figure 5: Treated Outcome Test Rejects
 Test Statistic Shows Positive Selection and/or Treatment Effect Heterogeneity
 Under LATE Assumptions



4 Test for Heterogeneous Treatment Effects under Ancillary Assumptions

Brinch et al. (2017) propose a test for treatment effect heterogeneity that relies on two ancillary assumptions beyond the LATE assumptions: 1) weak monotonicity of untreated outcomes in the fraction treated p , and 2) weak monotonicity of treated outcomes in the fraction treated p . As I show in Kowalski (2020a) and demonstrate here, the test only requires one of their ancillary assumptions; either one is sufficient. It is harder to justify the assumption of weak monotonicity of treated outcomes because variation in treated outcomes can reflect treatment effect heterogeneity as well as selection heterogeneity. I prefer not to make an opaque assumption about treatment effect heterogeneity to test for treatment effect

heterogeneity. I therefore proceed under the assumption of weak monotonicity of untreated outcomes, which is easier to justify because variation in untreated outcomes only reflects selection heterogeneity.

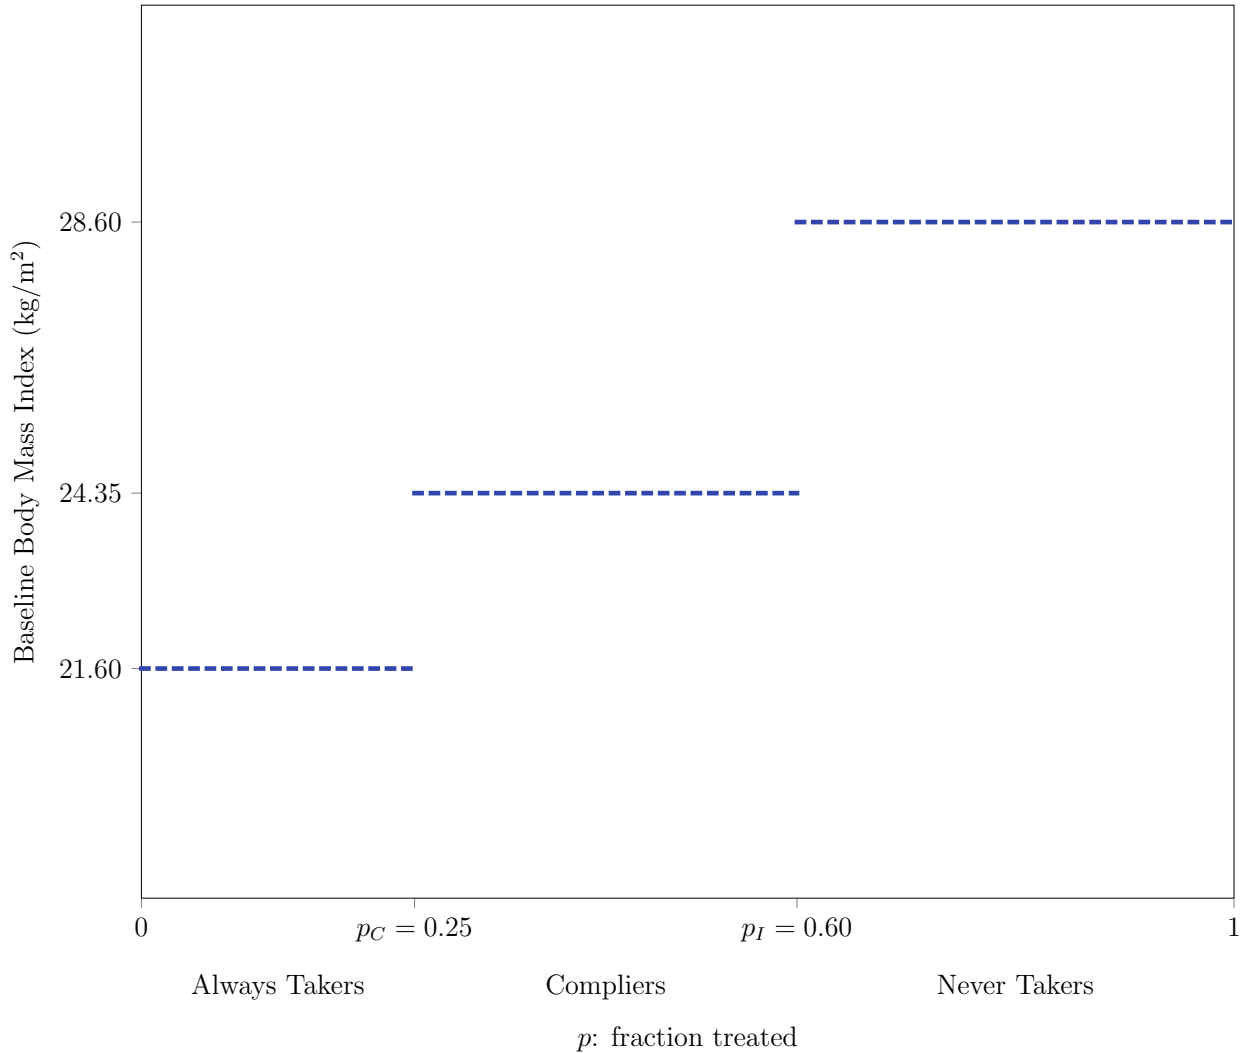
Researchers can justify the assumption of weak monotonicity of untreated outcomes using institutional features of their experiments. The assumption is easiest to justify if there is a plausible mechanism for differential gain from selection that is correlated with untreated outcomes. Such a gain need not be measured in terms of the main outcome of interest. In the insurance example, it is plausible that individuals with lower costs under the traditional plan are healthier and more likely to enroll in the wellness plan because they will be more likely to gain financially from the gym membership reimbursement. This gain to the enrollee is a loss to the insurer in terms of the main outcome of interest: its monthly costs.

Covariates collected at baseline before the experiment can also justify or call into question the assumption of weak monotonicity of untreated outcomes. Unlike untreated outcomes, which are not observed for always takers, baseline covariates can be observed for always takers, compliers, and never takers. Baseline covariates can thus serve as a proxy for untreated outcomes for all individuals. In the insurance example, baseline body mass index (BMI) can serve as a proxy for costs under the traditional plan because both are likely correlated with underlying health. The monotonic relationship in baseline BMI shown in Figure 6 lends support to the assumption of weak monotonicity of untreated outcomes. It shows that in terms of baseline BMI, healthier people are more likely to enroll in the wellness plan.

Figure 7 depicts the implications of the assumption of weak monotonicity of untreated outcomes in the insurance example. As discussed previously, the LATE monotonicity assumption implies an ordering from always takers to compliers to never takers along the horizontal axis. The ancillary assumption of weak monotonicity of untreated outcomes implies the same ordering along the vertical axis. Because the average untreated outcome of compliers is smaller than the average untreated outcome of never takers, yielding a negative untreated outcome test statistic, the ancillary assumption of weak monotonicity of untreated outcomes implies an upper bound on the average untreated outcome of always takers.

As depicted, the upper bound on the average untreated outcome of always takers implies a lower bound on the average treatment effect for always takers, which indicates that the wellness plan increases their average costs by at least \$22. In contrast, the LATE indicates that the wellness plan *decreases* average costs by \$50 for compliers. If the product of the untreated outcome test statistic and the difference between the bound on the average treatment effect for always takers and the LATE is negative, the test rejects the null hypothesis of treatment effect homogeneity under the single assumption that untreated outcomes are

Figure 6: Monotonicity of Baseline Body Mass Index Lends Support to the Ancillary Assumption of Weak Monotonicity of Untreated Outcomes

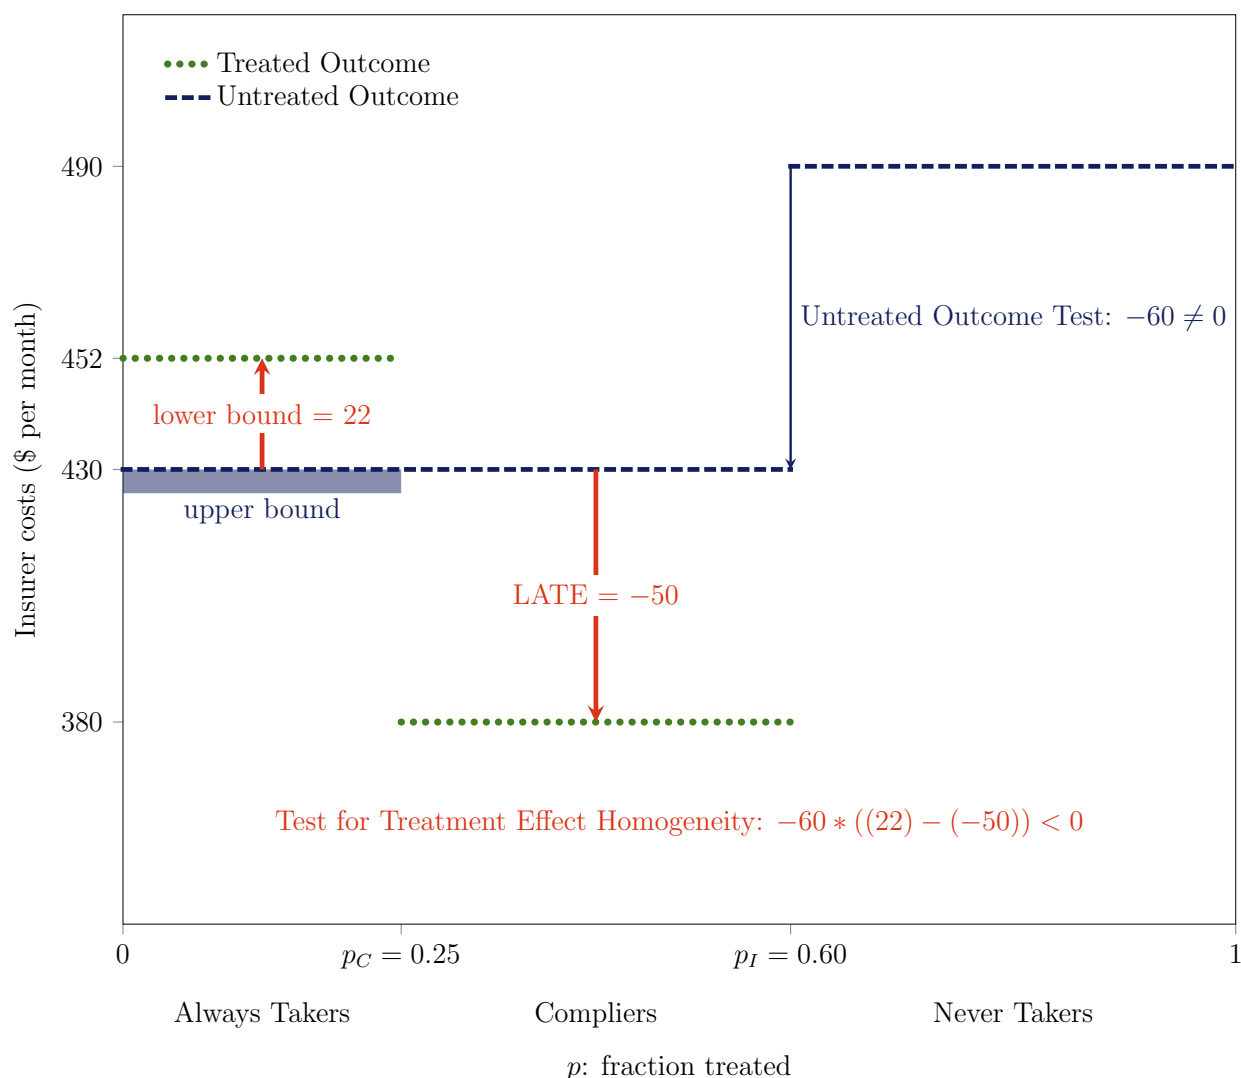


weakly monotonic in the fraction treated p .² It is straightforward to visualize that the test would also reject this null hypothesis under the single assumption that treated outcomes are weakly monotonic in the fraction treated p , demonstrating that either of the Brinch et al. (2017) ancillary assumptions is sufficient to reject treatment effect homogeneity.

What could explain the pattern of treatment effect heterogeneity shown in Figure 7? Baseline covariates show that always takers have lower body mass index than compliers, and

²In Kowalski (2020a), I propose to implement the test in this way to allow the bound on the average treatment effect for always takers to be an upper or lower bound. In the insurance example, if the untreated outcome test statistic had been positive, indicating adverse selection, then monotonicity of untreated outcomes would have implied an *upper* bound on the average treatment effect for always takers. It is straightforward to implement the test with a bootstrap procedure. Brinch et al. (2017) take a different approach. They test the signs of the treated and untreated outcome test statistics separately and then test if they are equal by accounting for multiple hypothesis testing using a Bonferroni correction, which could decrease power relative to the implementation that I propose.

Figure 7: Test for Treatment Effect Homogeneity Rejects
 Under Ancillary Assumption of Weak Monotonicity of Untreated Outcomes:
 Average Treatment Effect Bound for Always Takers



the assumption of weak monotonicity of untreated outcomes in the fraction treated p implies that always takers would have lower insurer costs in the traditional plan. Why then, would the wellness plan increase insurer costs for always takers but decrease them for compliers? Recall that the wellness plan offers gym membership reimbursement. It is plausible that the wellness plan increases insurer costs for always takers because the gym membership reimbursement crowds out an expense that they were already making. Therefore, insurer costs increase through the reimbursement but do not decrease on other dimensions through improved health. In contrast, the wellness plan decreases insurer costs for compliers because it induces them to take up a new gym membership, which decreases insurer costs on net.

The treatment effect heterogeneity between always takers and compliers depicted in Figure 7 has important implications for external validity. The LATE implies that the wellness

plan decreases costs by \$50 for compliers, but costs must increase by at least \$22 for always takers, so the LATE cannot be externally valid in a global sense for all policies. However, the insurer is interested in whether the LATE is externally valid in a local sense with respect to two specific options for expanding the wellness plan. The first option is to roll out the wellness plan at all sites just as it did last year at the pilot site, without the informational mailing. Figure 7 implies that the rollout of the wellness plan increased costs at the pilot site because the only enrollees were always takers. Based on the narrative explanation for the patterns of selection and treatment effect heterogeneity observed at the pilot site, it could be reasonable for the insurer to assume that those patterns would also apply at other sites: individuals with lower insurer costs in the traditional plan and lower body mass index would enroll, and insurer costs would increase because of the gym reimbursement. Therefore, it would not make sense for the insurer to roll out the wellness plan as an option at other sites.

However, the insurer might want to consider rolling out the wellness plan at other sites in conjunction with the informational mailing. At the pilot site, costs increased by at least \$22 for always takers, who represent 25% of beneficiaries, decreased by \$50 for compliers, who represent 35% of beneficiaries, and remained unchanged for never takers, who represent 40% of beneficiaries. Therefore, assuming negligible costs for the informational mailing, as long as the insurer believes that costs did not increase by more than \$100 ($=(50 \cdot 35) / .25$) for always takers, then the combination of the wellness plan and the informational mailing decreased its costs at the pilot site, and it could decrease costs at other sites.

The insurer also has the second wellness plan expansion option to consider: would it make more sense to enroll all beneficiaries at all sites in the wellness plan? To understand the impact of this option, the insurer needs information on the average treatment effect for never takers. Weak monotonicity of untreated outcomes is not sufficient to yield any meaningful conclusions about the average treatment effect for never takers because it is their *treated* outcomes that are unobserved. Although weak monotonicity of treated outcomes would yield meaningful conclusions, such an assumption is harder to justify as discussed. An alternative is to impose an additional ancillary assumption of weak monotonicity of treatment effects in conjunction with weak monotonicity of untreated outcomes. This approach is transparent because it makes assumptions about selection and treatment effect heterogeneity separately.

Researchers can again turn to institutional features of their experiments and data on covariates to justify an additional ancillary assumption of weak monotonicity of treatment effects. Suppose that before the experiment began, the insurer asked all beneficiaries at the pilot site how likely they would be to join a gym if offered gym membership reimbursement. It found that on average, always takers report that they would be the most likely to join a gym and never takers report that they would be the least likely to do so. However, given

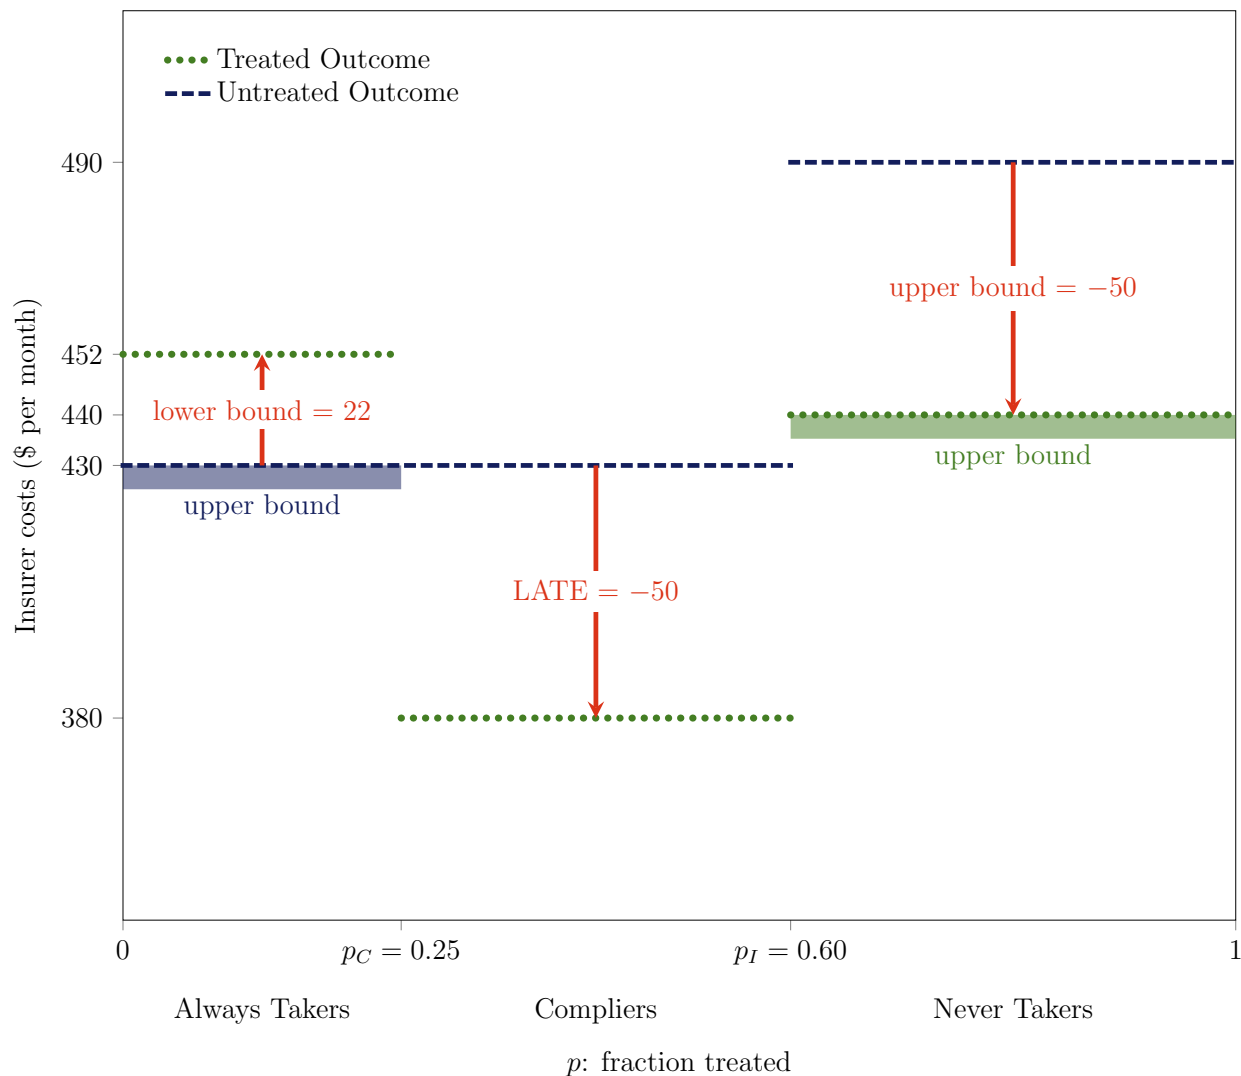
that the wellness plan decreases average insurer costs by more for compliers than for always takers, it is plausible that average costs would decrease by even more for never takers. Recall that never takers are the least healthy as measured by their baseline body mass index, so they have the biggest potential for insurer cost reductions through improved health. Thus, the variation in the likelihood of joining a gym observed across always takers, compliers, and never takers can support an additional ancillary assumption of weak monotonicity of treatment effects in the fraction treated p .

Figure 8 depicts the implications of an additional ancillary assumption of weak monotonicity of treatment effects. Under the assumption of weak monotonicity of untreated outcomes alone, the average treatment effect for always takers must be weakly greater than the average treatment effect for compliers. Weak monotonicity of treatment effects then implies that the average treatment effect for compliers is weakly greater than the average treatment effect for never takers. Therefore, the LATE of $-\$50$ gives an upper bound on the average treatment effect for never takers.

This upper bound helps to inform external validity in a local sense with respect to the second expansion option of enrolling all beneficiaries in the wellness plan. It implies that the wellness plan decreases insurer costs by at least $\$50$ per person for the 40% of beneficiaries that are never takers. At the same time, the wellness plan increases insurer costs by at least $\$22$ per person for the 25% of beneficiaries that are always takers and decreases costs by $\$50$ per person for the 35% of beneficiaries that are compliers. These bounds do not imply a bound on the average treatment effect across all beneficiaries. However, it is possible to calculate the average cost increase for always takers that would make the insurer indifferent between enrolling all beneficiaries in the wellness plan and enrolling all beneficiaries in the traditional plan at the pilot site. As long as the average cost does not increase by more than $\$150$ ($= (50 \cdot .4 + 50 \cdot .35) / .25$) per month for always takers, the insurer is better off enrolling all beneficiaries in the wellness plan. An extra $\$150$ per month for always takers would bring their average cost to $\$580$, which is much higher than the average cost for never takers, who have the highest costs, so it is plausible that the best wellness expansion option for the insurer is to enroll all beneficiaries in the wellness plan. This option strictly dominates the option of simply keeping the wellness plan as a choice at the pilot site.

But how should the insurer think about the merits of the expansion options beyond the pilot site? One traditional approach is to estimate LATEs within subgroups determined by covariates available at the pilot site and the site of interest and then re-weight those LATEs. Such an approach only informs the impact of the wellness plan for compliers who respond to the informational mailing, which is not necessarily relevant for either expansion option. To augment such an approach, the insurer could examine covariates at the additional sites that

Figure 8: Treatment Effect Bounds for Always Takers and Never Takers Under Ancillary Assumptions of Weak Monotonicity of Untreated Outcomes and Weak Monotonicity of Treatment Effects



vary with the fraction treated p at the pilot site. For example, if a site of interest has an average body mass index that is comparable to that of the always takers, then the insurer might be wary of expanding the wellness program to all beneficiaries at that site.

5 Tests for Heterogeneous Treatment Effects under Stronger Ancillary Assumptions

Thus far, I have demonstrated how researchers can reject the null of treatment effect homogeneity under the ancillary assumption of weak monotonicity of untreated outcomes in the fraction treated p . They can impose stronger assumptions to generate more powerful tests


and to obtain estimates of average treatment effects for always and never takers instead of bounds. Although it is natural to progress from weaker assumptions to stronger assumptions in empirical work, stronger assumptions were proposed first. To illustrate stronger assumptions, I present a second stylized example in which I cannot reject treatment effect heterogeneity under weak monotonicity of untreated outcomes and treatment effects, but I can under stronger assumptions.

Suppose researchers at a large hospital want to evaluate the impact of outpatient hip replacement surgery as compared to inpatient hip replacement surgery. Unlike inpatient surgery, outpatient surgery aims to send patients home on the same day. Potential benefits include reduced infection risk due to decreased time at the hospital, increased patient satisfaction from rehabilitation at home, and reduced healthcare costs. Data from the hospital records show that the rate of subsequent hip fractures is lower among patients who undergo outpatient surgery, consistent with higher quality and reduced cost. However, the researchers are concerned that patients who receive outpatient surgery might be systematically different from patients who receive inpatient surgery. To gather more evidence, they conduct a randomized trial. They choose trial participants from the population of patients scheduled to have hip replacement surgery and randomly assign them to a default of outpatient surgery (intervention) or inpatient surgery (control). Patients can consult with their doctors before choosing whether to proceed with their default surgery type, which leads to two-sided non-compliance with the treatment, receipt of outpatient surgery. The main outcome of interest is hip fracture within six months after surgery.

Results from the trial show that the average subsequent hip fracture rate is 2.75 percentage points lower in the intervention arm than it is in the control arm. The corresponding LATE depicted in Figure 9 implies that the treatment reduces the subsequent hip fracture rate for compliers by 5 percentage points on average. Moreover, the researchers find that the LATE is negative within all subgroups that they define by age, sex, and comorbidity score. Based on this promising evidence, should they go beyond recommending outpatient surgery as the default and mandate it for everyone? Or should they instead provide a subsidy to patients who undergo outpatient surgery?

To answer these questions, the researchers can begin by testing for heterogeneous selection within the trial under the LATE assumptions. The untreated outcome test statistic in Figure 9 is negative, implying selection heterogeneity such that the average subsequent hip fracture rate is 3 percentage points lower for untreated compliers than it is for never takers. This selection heterogeneity is consistent with the explanation that patients who are sicker are more likely to receive inpatient surgery because it allows for more direct supervision, which could justify an ancillary assumption of monotonicity of untreated outcomes. Baseline

Figure 9: Test for Treatment Effect Homogeneity Does Not Reject
 Under Ancillary Assumption of Weak Monotonicity of Untreated Outcomes:
 Bound on Average Treatment Effect for Always Takers



comorbidity score, which serves as a proxy for underlying health, also provides support for such an assumption. The average baseline comorbidity score increases from always takers to compliers to never takers. This evidence can also provide support for an assumption of weak monotonicity of treatment effects if there is reason to believe that patients with fewer comorbidities will respond better to outpatient surgery.

As shown in Figure 9, the ancillary assumption of weak monotonicity of untreated outcomes implies a lower bound on the average treatment effect for always takers. The product of the untreated outcome test statistic and the difference between the lower bound on the average treatment effect for always taker and the LATE is positive, so the test fails to reject the null of treatment effect homogeneity. Moreover, the additional ancillary assumption of weak monotonicity of treatment effects does not provide an informative bound on the aver-

age treatment effect for never takers in the sense that the average treatment effect for never takers could be greater or less than the LATE.


Even though the researchers cannot reject the null hypothesis of treatment effect homogeneity under weak monotonicity assumptions, they might be able to do so under stronger assumptions. One such set of stronger assumptions is linearity of untreated outcomes and treatment effects in the fraction treated p . Linearity assumptions are established in the literature. Olsen (1980) imposes linearity of treated outcomes, while Brinch et al. (2017) impose linearity of untreated and treated outcomes to test the null hypothesis of treatment effect homogeneity. Hausman (1978); Angrist (2004); Huber (2013); Bertanha and Imbens (2014); Guo et al. (2014) and Black et al. (2017) propose equivalent or similar tests, but they do not all explicitly state linearity assumptions. Linearity of untreated and treated outcomes is equivalent to linearity of untreated outcomes and treatment effects. However, I prefer to make assumptions on untreated outcomes and treatment effects because they are more transparent to motivate and assess.

Data on covariates can lend support to linearity assumptions, as I demonstrate in Kowalski (2020b). Continuing the surgery example, suppose the observed monotonic relationship in the average baseline comorbidity score is actually linear across always takers, compliers, and never takers, as depicted in Figure 10. Linear variation in the comorbidity score can support ancillary assumptions of linearity of untreated outcomes and treatment effects.

Figure 11 demonstrates the implications of the ancillary linearity assumptions in the surgery example. As in Kowalski (2020b), I refer to the function that specifies how untreated outcomes vary with the fraction treated p as the marginal untreated outcome function $\text{MUO}(p)$ for consistency with the corresponding function for treatment effects, which is commonly called the marginal treatment effect function $\text{MTE}(p)$. These functions are marginal functions in the sense that it is possible to obtain average quantities from them using weights over the fraction treated p following Heckman and Vytlačil (1999, 2001b, 2005), Carneiro et al. (2011), Brinch et al. (2017), and Kowalski (2020b). Linearity of untreated outcomes and treatment effects in the fraction treated p implies that the MUO and MTE functions are linear. The sum of the MUO and MTE functions yields the marginal treated outcome function $\text{MTO}(p)$. The MTO function is also linear in Figure 11 because the MUO and MTE functions are linear.

The ancillary linearity assumptions preserve the LATE of -5 while also yielding an estimate of the treatment effect at every fraction treated p , as depicted by the marginal treatment effect function $\text{MTE}(p)$ in Figure 11. As shown, the MTE function has a positive slope. Statistical significance of the slope can be obtained via bootstrap as in Brinch et al. (2017) and Kowalski (2020b). I do not recommend plotting confidence intervals for the MTE


Figure 10: Linearity of Baseline Comorbidity Score Lends Support to the Ancillary Assumptions of Linearity of Untreated Outcomes and Linearity of Treatment Effects



function, as doing so would convey statistical significance at a particular fraction treated p , which could be misleading about the statistical significance of quantities of interest derived from the MTE function, such as its slope. Instead, I recommend reporting confidence intervals on those quantities directly. If the slope of the MTE function is statistically significant, the test for treatment effect homogeneity rejects under the ancillary linearity assumptions.

The positively sloped MTE function in Figure 11 implies that the LATE cannot be externally valid in a global sense for all policies. However, the researchers are interested in whether the LATE is externally valid in a local sense with respect to two specific options. The first option is to mandate outpatient surgery for everyone. This option would enroll all never takers. The average treatment effect obtained by averaging the MTE function over the support of the fraction treated p for never takers shows that an outpatient surgery mandate would increase the average subsequent hip fracture rate for never takers by approximately 1.6 percentage points. Therefore, it is not sensible to choose a mandate over the policy of

Figure 11: Test for Treatment Effect Homogeneity Rejects Under Stronger Ancillary Assumptions of Linearity of Untreated Outcomes and Linearity of Treatment Effects: Treatment Effect Estimates



recommending outpatient surgery as the default.

However, instead of a mandate, should the researchers provide a subsidy to patients who undergo outpatient surgery? To understand the impact of this option, the researchers assume that all compliers who receive outpatient surgery under the default recommendation will also receive it under the subsidy, and the subsidy will also induce some never takers to receive outpatient surgery. They also assume that the ordering of patients along the horizontal axis does not change due to the subsidy. Given the slope of the MTE function, this assumption implies that the patients most likely to benefit from outpatient surgery are the most likely to receive it under the subsidy. Under these assumptions, the researchers want to set the subsidy just large enough to expand the fraction treated to the optimal level of approximately 78%, denoted by p^* in Figure 11, so that only the patients who benefit from outpatient surgery are induced to receive it. In practice, they can adjust the subsidy

over time if they overshoot or undershoot.

Covariates can also offer guidance on how to set the subsidy. To determine which covariates explain treatment effect heterogeneity, researchers can incorporate them into an MTE function under shape restrictions, as demonstrated in Kowalski (2020b). In the surgery example, linearity of baseline comorbidity score in the fraction treated p implies that incorporating it into the linear MTE function under an additive separability restriction will make the MTE function flat, indicating that this covariate explains all treatment effect heterogeneity. This result could be useful to the researchers pursuing the subsidy option because instead of guessing the right subsidy level, they could offer outpatient surgery only to those patients whose baseline comorbidity scores indicate that they would benefit from it.

Subgroup analysis does not offer similar insights into the optimal fraction treated p . Of course, it does not require ancillary assumptions or shape restrictions, but it delivers more limited results. In the surgery example, although the researchers estimated a negative LATE within all subgroups that they defined by age, sex, and comorbidity score, these LATEs only apply to compliers. The LATEs can possibly be negative within all subgroups even if there are positive treatment effects for always or never takers. Furthermore, subgroup analysis and LATE-reweighting are limited by which covariates the researchers have and how they use them to form subgroups. Researchers relying solely on subgroup analysis might erroneously conclude that treatment effects are homogeneous if they lack the right covariates. In contrast, if the linear MTE function has a nonzero slope and none of the available covariates make it flat, then they will know to keep looking for other covariates that can explain treatment effect heterogeneity. Even in the extreme case where values of the comorbidity score perfectly distinguish between always takers, compliers, and never takers, subgroup analysis need not completely uncover treatment effect heterogeneity because the LATEs cannot be estimated within subgroups in which all participants are treated or untreated. If researchers only report subgroup analysis among subgroups in which they can obtain estimates, they can erroneously conclude that all treatment effects are negative, even if treatment effects are positive within subgroups determined by a covariate that completely explains treatment effect heterogeneity.

Alternative ancillary assumptions that identify the MTE function at every fraction treated p also allow for tests of treatment effect heterogeneity and estimates of average treatment effects for always and never takers. For example, Kline and Walters (2019) show that the distributional assumptions made by the “Heckit” estimator of Heckman (1979) and the estimator used by Mroz (1987) identify the MTE function at every fraction treated p . The assumptions made by Willis and Rosen (1979) also identify the MTE function at every fraction treated p . As another example, Brinch et al. (2017) propose that MUO and MTO functions are quadratic and monotonic over the fraction treated from 0 to 1, and those as-

assumptions identify the MTE function at every fraction treated p . If covariates are available, it is also possible to use them to estimate more flexible forms for the MTE function, as in Carneiro and Lee (2009); Carneiro et al. (2011); Maestas et al. (2013); Kline and Walters (2016); Brinch et al. (2017) and Kowalski (2016) under additional shape restrictions.

While I consider expansion options within the hospital in the surgery example, researchers can also use the MUO and MTE functions from one hospital to examine potential expansion options at other hospitals. Evaluating whether a treatment effect estimated in one context is externally valid to policies in another context requires an additional assumption that both contexts have the same underlying MUO and MTE functions. I refer interested readers to Kowalski (2020b), where I demonstrate how examination of covariates, institutional details, and related outcomes can motivate such an assumption.

6 Implications for Experimental Design

To strengthen the case for external validity, researchers should design the intervention within an experiment to be as similar as possible to the policy of interest. Heckman and Vytlacil (2001a, 2007) make this insight clear with the concept of “policy-relevant treatment effects.” For any treatment, there can be multiple policies that affect takeup, and each policy can generate different sets of always takers, compliers, and never takers. For example, never takers in an experiment that involves a mailing intervention might be compliers in an experiment that involves a phone call intervention. Therefore, the average treatment effects for always takers, compliers, and never takers—and more generally, the shape of the MTE function—depend on the intervention.

Researchers risk weakening the case for external validity if they design the experimental intervention to force perfect compliance when the policy of interest would not force all individuals to receive treatment. Some researchers do so to increase power. Others do so with the goal of estimating a LATE that can be interpreted as the average treatment effect in a given population. However, unless the policy of interest would also force all individuals to receive treatment, an experiment with perfect compliance is not superior to an experiment with noncompliance for purposes of external validity.

Researchers also risk weakening the case for external validity if they design the experimental intervention to generate always and never takers that would not arise under the policy of interest, either because the policy of interest would involve *different* always or never takers or because it would not involve *any* always or never takers. Some policies naturally do not involve any always takers, especially if they introduce interventions that are not available otherwise. If there are no always takers or no never takers, the tests for

heterogeneous treatment effects discussed in this paper cannot be applied without further assumptions. However, the untreated outcome test can still identify heterogeneous selection when there are no always takers because it compares compliers and never takers.

Moreover, researchers should consider designing experiments with a range of interventions instead of a single intervention if they are interested in external validity to a range of policies. Several experimental designs involve a range of interventions, including those discussed in Burtless and Hausman (1978) Ashraf et al. (2010), Chassang et al. (2012), Basu (2015), Narita (2018), and Berry et al. (2020), among others. In the insurance example, researchers could implement such a design by offering a range of subsidies for enrollment in the wellness plan.

Designs that involve a range of interventions potentially involve a loss of power, but they have important advantages for the examination of external validity. Specifically, they can identify selection and treatment effect heterogeneity even if there are no always and never takers. If the range of interventions induces a continuous support over the fraction treated, then researchers can identify selection and treatment effect heterogeneity over that support without any ancillary assumptions, as long as individuals take up treatment in the same order as they would under the policy of interest. Such experimental variation in the fraction treated can also allow researchers to investigate the concern that treatment effects change as the fraction treated changes because of general equilibrium factors, especially if the variation is across experiments as in Lee et al. (2020).

Beyond designing their interventions thoughtfully, researchers should collect data to facilitate examination of external validity. To apply the approaches discussed in this paper, it is imperative to collect data that allow tabulations of outcomes by lottery status *and* treatment. It is also useful to collect data that allow similar tabulations of covariates. Data on covariates can also facilitate comparisons across experiments (Hotz et al. (2005); Angrist and Fernandez-Val (2013)). Approaches to assess external validity across experiments are even more powerful when used in concert with approaches to assess external validity within experiments.

References

- Alberto Abadie. Bootstrap tests for distributional treatment effects in instrumental variable models. *Journal of the American statistical Association*, 97(457):284–292, 2002.
- Alberto Abadie. Semiparametric instrumental variable estimation of treatment response models. *Journal of econometrics*, 113(2):231–263, 2003.

- Joshua D Angrist. Lifetime earnings and the vietnam era draft lottery: evidence from social security administrative records. *The American Economic Review*, pages 313–336, 1990.
- Joshua D Angrist. Estimating the labor market impact of voluntary military service using social security data on military applicants. *Econometrica*, 66(2):249–288, 1998.
- Joshua D Angrist. Treatment effect heterogeneity in theory and practice. *The Economic Journal*, 114(494):C52–C83, 2004.
- Joshua D Angrist and Ivan Fernandez-Val. ExtrapoLATE-ing: External validity and overidentification in the LATE framework. In *Advances in Economics and Econometrics: Volume 3, Econometrics: Tenth World Congress*, volume 51, page 401. Cambridge University Press, 2013.
- Joshua D Angrist and Alan B Krueger. The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples. *Journal of the American statistical Association*, 87(418):328–336, 1992.
- Joshua D Angrist, Guido W Imbens, and Donald B Rubin. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455, 1996.
- Nava Ashraf, James Berry, and Jesse M Shapiro. Can higher prices stimulate product use? evidence from a field experiment in zambia. *The American economic review*, 100(5):2383–2413, 2010.
- Alexander Balke and Judea Pearl. Nonparametric bounds on causal effects from partial compliance data. Working paper, 1993. URL <https://escholarship.org/uc/item/2rn0420q>.
- Alexander Balke and Judea Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.
- Anirban Basu. Welfare implications of learning through solicitation versus diversification in health care. *Journal of health economics*, 42:165–173, 2015.
- James Berry, Greg Fischer, and Raymond Guiteras. Eliciting and utilizing willingness to pay: Evidence from field trials in northern ghana. *Journal of Political Economy*, 128(4):1436–1473, 2020.

- Marinho Bertanha and Guido W. Imbens. External validity in fuzzy regression discontinuity designs. Working Paper 20773, National Bureau of Economic Research, December 2014. URL <https://www.nber.org/papers/w20773>.
- Anders Björklund and Robert Moffitt. The estimation of wage gains and welfare gains in self-selection models. *The Review of Economics and Statistics*, pages 42–49, 1987.
- Dan A Black, Joonhwi Joo, Robert LaLonde, Jeffrey A Smith, and Evan J Taylor. Simple tests for selection bias: Learning more from instrumental variables. Working Paper 6932, CESifo, March 2017. URL https://www.cesifo-group.de/DocDL/cesifo1_wp6392.pdf.
- Christian N Brinch, Magne Mogstad, and Matthew Wiswall. Beyond LATE with a discrete instrument. *Journal of Political Economy*, 125(4):985–1039, 2017.
- Gary Burtless and Jerry A. Hausman. The effect of taxation on labor supply: Evaluating the gary negative income tax experiment. *Journal of Political Economy*, 86(6):1103–1130, 1978.
- Pedro Carneiro and Sokbae Lee. Estimating distributions of potential outcomes using local instrumental variables with an application to changes in college enrollment and wage inequality. *Journal of Econometrics*, 149(2):191–208, 2009.
- Pedro Carneiro, James J. Heckman, and Edward J. Vytlacil. Estimating marginal returns to education. *American Economic Review*, 101(6):2754–81, October 2011.
- Sylvain Chassang, Gerard Padro I Miquel, and Erik Snowberg. Selective trials: A principal-agent approach to randomized controlled experiments. *American Economic Review*, 102(4):1279–1309, 2012.
- Thomas Cornelissen, Christian Dustmann, Anna Raute, and Uta Schönberg. Who benefits from universal child care? estimating marginal returns to early child care attendance. *Journal of Political Economy*, 126(6):2356–2409, 2018.
- Liran Einav, Amy Finkelstein, and Mark R Cullen. Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125(3):877, 2010.
- Zijian Guo, Jing Cheng, Scott A Lorch, and Dylan S Small. Using an instrumental variable to test for unmeasured confounding. *Statistics in medicine*, 33(20):3528–3546, 2014.
- Jerry A Hausman. Specification tests in econometrics. *Econometrica*, 46(6):1251–1271, 1978.

- James Heckman, Neil Hohmann, Jeffrey Smith, and Michael Khoo. Substitution and dropout bias in social experiments: A study of an influential social experiment. *The Quarterly Journal of Economics*, 115(2):651–694, 2000.
- James J Heckman and Edward Vytlacil. Policy-relevant treatment effects. *American Economic Review*, 91(2):107–111, 2001a.
- James J. Heckman and Edward Vytlacil. Structural Equations, Treatment Effects, and Econometric Policy Evaluation. *Econometrica*, 73(3):669–738, 05 2005.
- James J Heckman and Edward J Vytlacil. Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences*, 96(8):4730–4734, 1999.
- James J. Heckman and Edward J. Vytlacil. Local instrumental variables. In Cheng Hsiao, Kimio Morimune, and James L. Powell, editors, *Nonlinear Statistical Modeling: Proceedings of the Thirteenth International Symposium in Economic Theory and Econometrics: Essays in Honor of Takeshi Amemiya*, pages 1–46. Cambridge University Press, 2001b.
- James J Heckman and Edward J Vytlacil. Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. *Handbook of econometrics*, 6:4875–5143, 2007.
- James J Heckman, Jeffrey Smith, and Nancy Clements. Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies*, 64(4):487–535, 1997.
- James J. Heckman, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd. Characterizing selection bias using experimental data. *Econometrica*, 66(5):1017–1098, 1998.
- JJ Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–162, 1979.
- V Joseph Hotz, Guido W Imbens, and Julie H Mortimer. Predicting the efficacy of future training programs using past experiences at other locations. *Journal of Econometrics*, 125(1):241–270, 2005.
- Martin Huber. A simple test for the ignorability of non-compliance in experiments. *Economics Letters*, 120(3):389–391, 2013.

- Guido W. Imbens and Joshua D. Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–75, 1994.
- Guido W. Imbens and Donald B. Rubin. Estimating outcome distributions for compliers in instrumental variables models. *The Review of Economic Studies*, 64(4):555–574, 1997.
- Lawrence F Katz, Jeffrey R Kling, Jeffrey B Liebman, et al. Moving to opportunity in boston: Early results of a randomized mobility experiment. *The Quarterly Journal of Economics*, 116(2):607–654, 2001.
- Patrick Kline and Christopher R Walters. Evaluating public programs with close substitutes: The case of head start. *The Quarterly Journal of Economics*, 131(4):1795–1848, 2016.
- Patrick Kline and Christopher R Walters. On heckits, LATE, and numerical equivalence. *Econometrica*, 87(2):677–696, 2019.
- Amanda Kowalski. Doing more when you’re running LATE: Applying marginal treatment effect methods to examine treatment effect heterogeneity in experiments. Working Paper 22362, National Bureau of Economic Research, June 2016. URL <http://www.nber.org/papers/w22362>.
- Amanda Kowalski, Yen Tran, and Ljubica Ristovska. MTEBINARY: Stata module to compute Marginal Treatment Effects (MTE) With a Binary Instrument. Statistical Software Components, Boston College Department of Economics, 2018. URL <https://ideas.repec.org/c/boc/bocode/s458285.html>.
- Amanda E Kowalski. Behavior within a clinical trial and implications for mammography guidelines. Working Paper 25049, National Bureau of Economic Research, November 2020a. URL <http://www.nber.org/papers/w25049>.
- Amanda E. Kowalski. Reconciling seemingly contradictory results from the Oregon health insurance experiment and the Massachusetts health reform. Working Paper 24647, National Bureau of Economic Research, November 2020b. URL <http://www.nber.org/papers/w24647>.
- Kenneth Lee, Edward Miguel, and Catherine Wolfram. Experimental evidence on the economics of rural electrification. *Journal of Political Economy*, 128(4):1523–1565, 2020.
- Nicole Maestas, Kathleen J Mullen, and Alexander Strand. Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt. *The American Economic Review*, 103(5):1797–1829, 2013.

- Charles F Manski. Nonparametric bounds on treatment effects. *The American Economic Review*, 80(2):319–323, 1990.
- Magne Mogstad, Andres Santos, and Alexander Torgovitsky. Using instrumental variables for inference about policy relevant treatment effects. *Econometrica*, 86(5):1589–1619, 2018.
- Thomas A. Mroz. The sensitivity of an empirical model of married women’s hours of work to economic and statistical assumptions. *Econometrica*, 55(4):765–799, 1987.
- Yusuke Narita. Toward an ethical experiment. Working Paper 2127, Cowles Foundation, 2018. URL <https://cowles.yale.edu/sites/default/files/files/pub/d21/d2127.pdf>.
- Randall J. Olsen. A least squares correction for selectivity bias. *Econometrica*, 48(7):1815–1820, 1980.
- James M Robins. The analysis of randomized and non-randomized aids treatment trials using a new approach to causal inference in longitudinal studies. *Health service research methodology: a focus on AIDS*, pages 113–159, 1989.
- Andrew Donald Roy. Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2):135–146, 1951.
- Edward Vytlacil. Independence, monotonicity, and latent index models: An equivalence result. *Econometrica*, 70(1):331–341, 2002.
- Abraham Wald. The fitting of straight lines if both variables are subject to error. *Ann. Math. Statist.*, 11(3):284–300, 09 1940.
- Robert J Willis and Sherwin Rosen. Education and self-selection. *Journal of Political Economy*, 87(5, Part 2):S7–S36, 1979.