Definition and Evaluation of the Monotonicity Condition for Preference-Based Instruments

Sonja A. Swanson,
Department of Epidemiology, Harvard School of Public Health, Boston, Massachusetts

Matthew Miller,
Department of Epidemiology, Harvard School of Public Health, Boston, Massachusetts;
Department of Health Science, Northeastern University, Boston, Massachusetts

James M. Robins, and
Department of Epidemiology, Harvard School of Public Health, Boston, Massachusetts;
Department of Biostatistics, Harvard School of Public Health, Boston, Massachusetts

Miguel A. Hernán
Department of Epidemiology, Harvard School of Public Health, Boston, Massachusetts;
Department of Biostatistics, Harvard School of Public Health, Boston, Massachusetts; Harvard-MIT Division of Health Sciences and Technology, Boston, Massachusetts

Abstract

Background—Preference-based instrumental variable methods are often used in comparative effectiveness research. Many instrumental variable studies estimate the local average treatment effect (i.e., the effect in the "compliers") under the assumption of monotonicity, i.e., no "defiers," and well-defined compliance types. However, the monotonicity assumption has not been empirically tested and the meaning of monotonicity itself is unclear.

Methods—Here we clarify the definition of local and global monotonicity and propose a novel study design to assess the monotonicity assumption empirically. Our design requires surveying physicians about their treatment plans and prescribing preferences for the same set of patients. We also discuss measures of monotonicity that can be calculated from this survey data. As an illustration, we conducted a pilot study in a survey of 53 physicians who reported treatment plans and prescribing preferences for hypothetical patients who were candidates for antipsychotic treatment.

Results—In our study, nearly all patients exhibited some degree of monotonicity violations. In addition, patients could not be cleanly classified as compliers, defiers, always-takers, or never-takers.

Conclusions—We conclude that preference-based instrumental variable estimates should be interpreted cautiously because bias due to monotonicity violations is likely and because the

Corresponding Author: Sonja A. Swanson, Department of Epidemiology, Harvard School of Public Health, 677 Huntington Avenue, Kresge 9th Floor, Boston, Massachusetts 02115, sswanson@hsph.harvard.edu, 651.270.9984.

subpopulation to which the estimate applies may not be well-defined. Investigators using preference-based instruments may consider supplementing their study with a survey to empirically assess the magnitude and direction of bias due to violations of monotonicity.

Keywords

compliance type; complier; defier; instrumental variable; local average treatment effect; monotonicity

Comparative effectiveness research frequently uses instrumental variable methods to estimate causal effects.[1–3] Such methods require investigators to propose a variable (the "instrument") that meets the three instrumental conditions: it is associated with treatment, it causes the outcome only through treatment, and its effect on the outcome is not confounded. Commonly proposed instruments are based on facility or physician prescribing preferences,[1,2] e.g., a physician's preference for Treatment A over Treatment B. Even if preference met the instrumental conditions, an additional condition is necessary to identify a causal effect. Most instrumental variable applications assume the additional condition of monotonicity.[3,4]

The monotonicity assumption implies there are no "defiers", i.e., no patients who would be prescribed Treatment A when seen by a physician who usually prefers B *and* would be prescribed Treatment B when seen by a physician who usually prefers A (a rigorous and more general definition is provided below). Under this assumption, instrumental variable methods can be used to identify the average causal effect in the subset of patients who would be prescribed the treatment preferred by any treating physician, i.e., the "compliers". Monotonicity generally cannot be verified because we cannot observe what would have happened had the same patient been treated by another physician.

Monotonicity might be reasonably assumed if patient characteristics naturally collapsed into a single dimension related to treatment decisions (e.g., a propensity score), and clinicians had different cut points along this continuum for deciding when to prescribe which treatment. However, given the complexity of information physicians integrate into their prescribing decisions, preferences are unlikely to be so cleanly ordered. As a simplified example, consider a physician who generally prefers Treatment A, but prescribes Treatment B for more physically active patients (e.g., because Treatment A is associated with risk of motor-skill impairment), and another physician who generally prefers Treatment B, but makes exceptions for patients with a family history of diabetes (e.g., because a new study suggests such patients might respond better to Treatment A). Any physically active patient with a family history of diabetes who could potentially have seen either of these providers would "defy" both preferences and thus violate the monotonicity assumption, even conditional on covariates other than physical activity and family history.

Despite many opportunities for monotonicity violations, the possible bias introduced from such violations in instrumental variable analyses has not been previously explored. Moreover, the meaning of monotonicity itself is unclear in realistic applications of instrumental variable analyses with preference-based instruments, something rarely discussed in the literature. Here we (1) define monotonicity and the interpretation of

instrumental variable estimates in the context of preference-based instruments and a dichotomous prescribing decision (e.g., Treatment A vs. B), (2) describe a novel study design to assess deviations from monotonicity empirically by surveying physicians about their prescribing preferences and the treatment decisions they would make for a set of hypothetical patients, and (3) implement a pilot study to demonstrate the feasibility of our design to detect monotonicity violations when studying the effects of atypical versus conventional antipsychotic medication on risk of death in the elderly.

DEFINITION OF MONOTONICITY AND INTERPRETATION OF ESTIMATES

For each patient, let Z be the instrument (Z=1 indicates the patient's physician prefers Treatment A, Z=0 indicates the patient's physician prefers Treatment B), X be the treatment (X=1 indicates being prescribed Treatment A, X=0 indicates being prescribed Treatment B), and X^z the counterfactual treatment under a given preference z. Throughout, we refer to X as "treatment" as shorthand for "being prescribed treatment." If the counterfactuals are deterministic,[5] all patients in the study population can be classified into one of four mutually exclusive compliance types:

- "Always-takers": patients who would be prescribed Treatment A by any physician, i.e., patients with $X^{z=1}=X^{z=0}=1$.

- "Never-takers": patients who would not be prescribed Treatment A by any physician, i.e., patients with $X^{z=1}=X^{z=0}=0$.

- "Compliers": patients for whom a physician who prefers Treatment A would prescribe Treatment A *and* a physician who prefers Treatment B would prescribe Treatment B, i.e., patients with $X^{z=1}=1$ and $X^{z=0}=0$.

- "Defiers": patients for whom a physician who prefers Treatment A would have prescribed Treatment B *and* a physician who prefers Treatment B would have prescribed Treatment A, i.e., patients with $X^{z=1}=0$ and $X^{z=0}=1$.

These compliance types are illustrated in Figure 1. Note compliance types are study-specific and instrument-dependent: a patient is not inherently a complier but is defined as such only in the context of a particular study with respect to a particular proposed instrument.[4,6]

Monotonicity means there are no defiers. When monotonicity and the three instrumental conditions hold, the average treatment effect for the subset of compliers is identifiable (known as the local average treatment effect [LATE]).[4] The effect in the compliers is a problematic quantity because we generally do not know a patient's compliance type, as we usually observe their treatment under only one physician's preference, and therefore the subset of compliers is unknown.

Moreover, the above compliance types are generally not well defined. To see this, consider a patient for whom some (but not all) physicians who prefer Treatment A would prescribe Treatment B and some (but not all) physicians who prefer Treatment B would prescribe Treatment A. Depending on which pair of physicians' treatment decisions we consider, this patient could be a defier, complier, always-taker, or never-taker. That compliance types are

generally ill defined has been rarely[7] mentioned in the instrumental variable literature. We now describe more precise definitions of monotonicity and compliance types.

For a dichotomous or non-dichotomous instrument Z, monotonicity means the counterfactual treatment is a non-decreasing function of the instrument:

$$z_1 < z_2 \Rightarrow X_i^{z_1} \le X_i^{z_2} \text{ for all } i,$$

where, in our example, X_i^z denotes the treatment that subject i would receive if treated by a physician with preference z. However, the above definition is incomplete because physicians with identical preferences z may treat the same patient differently: if the physician is not explicitly specified, the counterfactuals X^z are not well defined. Let us then further index the counterfactuals by physician P so that $X_i^{z_1,p}$ denotes the treatment that subject i would have received if treated by physician p who has preference value z_1. On the one hand, this notation is unnecessary because $X_i^{z_1,p} = X_i^p$. On the other hand, this notation makes clear that $X_i^{z_2,p}$ is ill-defined because physician p has preference z_1, not z_2 (i.e., it is unclear what is meant by "the treatment subject i would have received if treated by physician p who has preference value z_1 had physician p had, counter to the fact, preference z_2"). Monotonicity cannot be satisfactorily defined in terms of either $X_i^{z_1}$ $X_i^{z_2}$ or $X_i^{z_1,p}$ $X_i^{z_2,p}$. Rather, we need to define monotonicity as

$$z_1 < z_2 \Rightarrow X_i^{z_1,p} \le X_i^{z_2,p'} \iff X_i^p \le X_i^{p'}$$

where p and p' represent two different physicians. The above definition is still vague because it does not specify who p and p' are. One possibility is to require that the inequality applies to all possible (p, p') pairs: global monotonicity. Another possibility is to specify p as the physician who actually treated patient i and p' as the physician with preference z_2 who would have treated patient i if no physician with z_1 would have been available: local monotonicity. Global monotonicity implies local monotonicity.

The discussion above illustrates how the LATE estimated with a dichotomous instrument under (global or local) monotonicity is not a well-defined parameter for preference-based instruments. The LATE traditionally is said to estimate the effect in the subpopulation of compliers, but this subpopulation is ill defined for the same reasons that the subpopulation of defiers is ill defined. In order to provide a precise counterfactual definition of compliers, one would need to specify the physician p' with preference z_2 who would have treated patient i if no physician with z_1 was available. Then patient i would be a complier if $X_i^{z_1,p} < X_i^{z_2,p'}$ for $z_1 < z_2$. If p or p' cannot be specified, the ill-definition of compliance types could alternatively be viewed as multiple versions of the instrument.[8]

STUDY DESIGN AND MEASURES OF MONOTONICITY

Suppose we had administrative data on a cohort of patients prescribed some treatment(s) of interest and are planning to use physician's preference as a proposed instrument to assess the effect of treatment on one or more outcomes using this dataset. In this section, we describe a

survey to be completed by the prescribing physicians from the cohort. This supplemental survey allows empirical assessment of the monotonicity condition.

Survey Design

The survey includes two components. The first component presents hypothetical patients and asks physicians for their treatment plans. Hypothetical patients should be descpiribed with sufficient information for the physicians to make relatively well-informed decisions, which could be provided in a number of formats. In our pilot study, described in eAppendix section 1, the information is presented as case histories, i.e., as short vignettes describing the reason for visit and relevant patient characteristics. Other formats for presenting the information could also be used such as x-rays or other clinical measures.[9] Decisions about the format and information included depend on the particular study question and the patient characteristics suspected to be most relevant to the treatment decision (including characteristics that may not be measured in the administrative data). The patients should also reasonably represent the original patient population in order to emulate their counterfactual treatment distribution. Perfect representation will not be feasible, however. For example, in the pilot study described below, we present only a small number of hypothetical patients that are loosely representative of published studies with respect to univariate distributions of measured patient characteristics.

The other component of the study design is an assessment of the physicians' prescribing preferences. One option is to use self-report on prior prescribing history; other options are described in the online supplemental materials (eAppendix section 2). If the study could be linked to administrative data, the measure of preference may come from the original data and/or could be the same measure used in the primary study. The implications of measurement error in the context of assessing monotonicity,[10] including the various measures used in our pilot study, are discussed more fully in our online supplemental materials. The analytic strategy in the main text relies on the assumption that preference is dichotomous and measured without error.

The key to the survey design is that we observe the counterfactual treatments for all of the hypothetical patients had they seen any of the physicians completing the survey. Coupled with measures of preference, this allows us to assess monotonicity directly, as described below.

Measures of Monotonicity

By observing all counterfactual treatments made by all physicians for all patients, we can assess global monotonicity by assessing whether, for any patient, there is a physician pair such that the physicians have different levels of preferences and would both prescribe against their preference. Global monotonicity is not directly relevant to understanding the magnitude of bias; however, we consider it a first step to understanding violations of monotonicity, particularly because assessment of local monotonicity and bias require additional assumptions regarding which physicians could see which patients. Global monotonicity also aligns with the definition of monotonicity presented in previous literature describing preference-based instruments.[10]

We begin by assuming all patients could be seen by all physicians who completed the survey with equal probability and that the physicians are independent from one another. In this case, the distribution of prescribing decisions for all possible physician pairs could be used to estimate the probability a patient i is a particular compliance type:

$$\Pr[X_i^{z=0}=x_0, X_i^{z=1}=x_1]=\Pr[X_i^{z=0}=x_0]\Pr[X_i^{z=1}=x_1]$$
$$=\Pr[X_i=x_0|Z=0]\Pr[X_i=x_1|Z=1]$$
$$=\left\{\frac{\sum_{j=1}^m I(X_{i,j}=x_0, Z_j=0)}{\sum_{j=1}^m I(Z_j=0)}\right\}\left\{\frac{\sum_{j=1}^m I(X_{i,j}=x_1, Z_j=1)}{\sum_{j=1}^m I(Z_j=1)}\right\}$$

where j indexes the m physicians and $I()$ is an indicator function. We can then estimate the probability of each compliance type across all n patients:

$$\Pr[X^{z=0}=x_0, X^{z=1}=x_1]=E[\Pr[X_i^{z=0}=x_0, X_i^{z=1}=x_1]]$$
$$=\frac{1}{n}\sum_{i=1}^n\left\{\frac{\sum_{j=1}^m I(X_{i,j}=x_0, Z_j=0)}{\sum_{j=1}^m I(Z_j=0)}\right\}\left\{\frac{\sum_{j=1}^m I(X_{i,j}=x_1, Z_j=1)}{\sum_{j=1}^m I(Z_j=1)}\right\}$$

Again, our compliance types will be ill-defined if physicians with the same preference disagree; under our assumption that all patients could be seen by all physicians with equal probability, we can view this as being multiple versions of the instrument.[8] As such, the estimate could be interpreted as follows: if we were to randomly draw a patient from the population and then randomly assign the patient to a relevant physician pair, this is the probability we draw a specific compliance type. For well-defined causal dichotomous instruments, it has been previously demonstrated that bias in the LATE is a function of the relative proportion of defiers to compliers and the difference between the effects in the compliers and defiers.[4] This is because the estimate is a "weighted" average of the effects in the compliers and defiers, but the "weights" for the defiers are negative. Our survey design could potentially inform the probabilities of each compliance type, while we would rely on subject matter knowledge to propose a reasonable range of suspected effect heterogeneity across the compliers and defiers. See eAppendix 3 for consideration of a related estimand.

Richardson and Robins[11] identified bounds for the LATE, the average effect within other compliance types, and the average treatment effect in the full study population under the three instrumental conditions plus an assumed feasible distribution of compliance types. When coupled with data from a follow-up study, one could potentially use this survey design to inform the distribution of compliance types, and then compute bounds for these treatment effects using the observed distributions in the cohort data. In practice, the validity of this approach would depend on the same assumptions described above, namely our assumptions regarding which physicians could see which patients and that the hypothetical patients reasonably represented the patients in the cohort study. It is also possible that the estimated distribution of compliance types from the survey design would be incompatible with the cohort data, indicating that one or more of the assumptions are ill-placed.

The assumption that all patients could be seen by all physicians with equal probability is unrealistic: for example, patients typically have geographic and insurance restrictions on the physicians they can see. If we had information on such restrictions, the analytic strategy

described above could be adapted to include such restrictions by excluding or down-weighting the relevant patient–physician pairs. Arguably, only two counterfactual treatments would be relevant, if we thought they were well-defined: the treatment given by the physician the patient would have actually seen, and the treatment given by the physician the patient would have seen had they been forced to see a physician of a different preference level.[12,13]

All of the above definitions describe monotonicity when there are only two possible treatments (e.g., Treatment A or B), when in fact many treatment decisions include other possibilities (e.g., Treatment C, or no treatment). To be consistent with the analytic strategy employed by previously published instrumental variable studies, we exclude observations from analyses in our pilot study when the physicians choose alternative options. However, selecting on treatment in instrumental variable analyses can lead to substantial biases.[14]

PILOT STUDY

We now describe a pilot study using the proposed survey design to assess the monotonicity condition in the context of estimating the effects of atypical and conventional antipsychotic medications. The current implementation has several limitations, most notably a small sample size and low response rate. The results we present are demonstrations of the usefulness and feasibility of such a survey, but we stress that the estimates obtained should be viewed with due skepticism.

Physician Study Population

IMS Health provided data to identify physicians with a relevant medical specialty (family medicine, internal medicine, psychiatry), active history of prescribing antipsychotic medications, and valid email address.[15,16] Details of the eligibility criteria and the data used to identify these physicians are described in the online supplemental materials (eAppendix 2). We identified 17,665 eligible physicians and contacted a random subsample of 4,800. Physicians were twice emailed information about the study, including a hyperlink to the online questionnaire. Fifty-three (1%) completed the questionnaire. Because survey responses were confidential, we have limited information on the representativeness of the physicians who completed the survey: respondents appeared reasonably representative with respect to medical specialty, but were more likely to be retired or semi-retired. Details of the demographic and medical practice characteristics are described in eTable 1.

Questionnaire

The hypothetical patient population was informed by four studies[17–20] with instrumental variable analyses used to estimate the effect of atypical versus conventional antipsychotic medication on death in the elderly. Psychiatrists experienced in prescribing antipsychotic medications were consulted and their input helped assure that each presented scenario was realistic and provided sufficient information to make relatively well-informed prescribing decisions. The full questionnaire is included as an online supplement (eAppendix section 3).

Physicians were presented with vignettes describing hypothetical patients, with information including patient characteristics likely to inform treatment plans: reason for visit, treatment

history, relevant comorbidities, general health measures, psychosocial factors, and results from pertinent medical tests. Physicians were asked to consider this information and indicate whether their most likely treatment approach would include prescribing an antipsychotic medication and, if so, whether they would choose a conventional or atypical. Following each of these (index) vignettes, physicians were asked for their treatment plans given additional information about the hypothetical patient presented in four scenarios (e.g., suppose the index patient also had diabetes). These additional scenarios varied personal history, family history, and experiences with prior antipsychotic treatment greater than a year ago. The hypothetical patient population for primary analyses included elderly patients featured in twenty scenarios with a distribution of age, sex, likely indications, and comorbidities approximating the distributions in previous studies (eTable 2).[17–20] We also included five scenarios regarding a younger patient with schizophrenia (data not shown).

We present results using the analytic strategies described above for this study design. The results in the main text are based on only one measure of preference: self-report of the class of antipsychotic medication prescribed to the most recent patient initiating antipsychotic medication. Descriptions of the other proxies assessed, and accompanying results, can be found in the online supplemental materials (eAppendix section 2, eTables 3–4, eFigures 1–2).

Results

Six physicians (12%) last prescribed a conventional antipsychotic while 47 (88%) prescribed an atypical antipsychotic. Across all patients and physicians, the majority (73%) of antipsychotic prescriptions were for an atypical antipsychotic. Global monotonicity was violated, in that there was at least one pair of physicians for at least one patient who prescribed opposite of their preferences: in fact, such violations were apparent for 17 patients (85%).

Across physician pairs, some patients were never labeled a defier and others were defiers nearly a quarter of the time (Figure 2). Patient #8, for example, never exhibited a monotonicity violation because no physician prescribed her conventional antipsychotic medication. In contrast, physicians often treated Patients #19 and #20 against their prescribing preference. To illustrate this, the distribution of prescribing decisions by physicians' preference for Patients #8 and #19 are provided in Table 1. The estimated proportion of defiers was 10%; the proportion of compliers was 34% (Figure 2).

By comparing the probabilities of being a defier for the four base vignettes with their additional scenarios, we found that comorbidities and prior adverse reactions to antipsychotic medications were associated with an increased probability of being a defier. For example, Patient #1 had an estimated 0.04 probability of being a defier, but when the physicians were instructed that this patient also had a recent myocardial infarction and a recent hip fracture (Patient #5), his estimated probability was 0.16. Rather than indicating that comorbidities are associated with an increased probability of being a defier, however, this observation could be an artifact of the question structure, which was primarily used to shorten survey time but which may have encouraged the physicians to switch their treatments for the modified vignettes.

The four previously published instrumental variable studies reported a median effect in the compliers of 8.1 deaths per 100 patients for conventional versus atypical antipsychotic medication. This estimate would be biased if defiers exist and we expected effect heterogeneity (Figure 3). Some evidence for such effect heterogeneity is apparent in previous studies: the effect estimates differ between community-dwelling patients and nursing home patients (e.g., a difference in risk differences of up to 8 deaths per 100 patients).[18–20] Because the probability of being a defier was associated with having multiple medical complications, this suggests the difference between the effect in compliers and defiers may be of similar magnitude. If the proportion of defiers in these four studies had been 10%, and that of compliers 34%, and we expected differences in effects between the compliers and defiers between 5 and 10 deaths per 100 patients, the corrected effect in the compliers would be between 3.9 and 6.0 deaths per 100 patients (% bias = 35–108%). The effect in the defiers would be closer to the null or of opposite sign (−6.1 to 1.0 deaths per 100 patients).

DISCUSSION

We described and demonstrated a study design for empirically assessing the monotonicity assumption in the context of preference-based instruments. Our findings raise concerns of potential bias for all instrumental variable analyses proposing preference-based instruments to estimate the LATE. This finding of widespread monotonicity violations is not unexpected. As discussed in the introduction, monotonicity may be reasonable for instruments that have only one dimension encouraging treatment (e.g., as it was initially proposed in the context of a randomized design), but questionable for any instrument, such as physician's preference, that has more than one dimension of encouragement.[9,21]

We used our pilot study results to propose corrected estimates of previously published studies of antipsychotic medication prescribing decisions. As previously presented, the instrumental variable analyses from these studies suggest that atypical antipsychotic medication is safer for one quarter to one third of the study population, with no evidence for the rest of the population. With our bias adjustments, we may instead conclude that atypical antipsychotic medication would be slightly preferable for 34% of the population, while for 10% of the population conventional antipsychotics would likely be preferable, and for 56% of the population we remain uncertain. However, even this conclusion is incomplete as it does not capture the fact that the 34% and 10% are not just unidentifiable subsets, but are only well defined if we know the relevant physician pairs that would have treated each patient. Moreover, such a conclusion rests on several strong assumptions, e.g.: (1) that our patients could be seen with equal probability by all survey respondents, (2) that preference is dichotomous and measured without error, (3) that the prescribing decisions of the physicians who responded to the survey adequately represented the prescribing decisions of all physicians who could have seen these patients, and (4) that our hypothetical patient population would have been treated similarly to the prior study populations. We do not believe that any of these assumptions hold perfectly or even approximately in our pilot study, and view the bias corrections as a demonstration of the type of results such a survey can provide. Adequate understanding of compliance types and correction for bias due to monotonicity violations in a particular study demands that investigators conduct their own

empirical assessment. That said, the limitations notwithstanding, our pilot study provides valid empirical evidence against global monotonicity, and highlights the possibility of large biases and issues with ill-definition in preference-based instrumental variable analyses.

Causal inference requires untestable assumptions, whether using instrumental variable or other methods. The onus is on the investigators to demonstrate why any assumption made appears reasonable and the extent to which a violation of the assumption may affect their conclusions.[3,4] The current study underscores skepticism about the use of preference-based instrumental variable methods to estimate the LATE, indicating that at best it applies to a not just unidentifiable but ill-defined subset of the population, and offers an approach to understanding the magnitude of potential bias when the monotonicity assumption does not hold.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Abbreviations

LATE   local average treatment effect

References

1. Chen Y, Briesacher BA. Use of instrumental variable in prescription drug research with observational data: a systematic review. J Clin Epidemiol. 2011; 64:687–700. [PubMed: 21163621]

2. Davies NM, Smith GD, Windmeijer F, Martin RM. Issues in the reporting and conduct of instrumental variable studies: a systematic review. Epidemiology (Cambridge, Mass). 2013; 24:363–9.

3. Swanson SA, Hernán MA. Commentary: how to report instrumental variable analyses (suggestions welcome). Epidemiology (Cambridge, Mass). 2013; 24:370–4.

4. Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. Journal of the American Statistical Association. 1996; 91:444–55.

5. Hernán, MA.; Robins, JM. Causal Inference. Chapman & Hall/CRC; 2015.

6. Pearl J. Principal stratification—a goal or a tool? The international journal of biostatistics. 2011;7.

7. Brookhart MA, Schneeweiss S. Preference-based instrumental variable methods for the estimation of treatment effects: assessing validity and interpreting results. The international journal of biostatistics. 2007; 3:14.

8. Hernán MA, VanderWeele TJ. Compound treatments and transportability of causal inference. Epidemiology (Cambridge, Mass). 2011; 22:368–77.

9. Korn EL, Baumrind S. Clinician preferences and the estimation of causal treatment differences. Statistical Science. 1998; 13:209–35.

10. Hernán MA, Robins JM. Instruments for causal inference: an epidemiologist's dream? Epidemiology (Cambridge, Mass). 2006; 17:360–72.

11. Richardson, T.; Robins, JM. Analysis of the binary instrumental variable model. In: Dechter, R.; Geffner, H.; Halpern, JY., editors. Heuristics, Probability, and Causality: A Tribute to Judea Pearl. 2010. p. 415-44.

12. Stalnaker, RC. A theory of conditionals. Ifs. Springer; 1980. p. 41-55.

13. Lewis D. Causation. The Journal of Philosophy. 1973; 70:556–67.

14. Swanson SA, Robins JM, Miller M, Hernán MA. Selecting on treatment: A pervasive form of bias in instrumental variable analyses. American journal of epidemiology. e-published 2015.

15. AMA. [Accessed 2013] Physician Masterfile. 2013. at http://www.ama-assn.org/ama/pub/about-ama/physician-data-resources/physician-masterfile.page?

16. [Accessed 2013] Xponent Prescription Database. 2013. at http://www.imshealth.com/portal/site/imshealth/menuitem.c76283e8bf81e98f53c753c71ad8c22a/?vgnextoid=39d15dba9648e310VgnVCM10000076192ca2RCRD&vgnextchannel=abb6e590cb4dc310VgnVCM100000a48d2ca2RCRD

17. Huybrechts KF, Brookhart MA, Rothman KJ, et al. Comparison of different approaches to confounding adjustment in a study on the association of antipsychotic medication with mortality in older nursing home patients. American journal of epidemiology. 2011; 174:1089–99. [PubMed: 21934095]

18. Pratt N, Roughead EE, Ryan P, Salter A. Antipsychotics and the risk of death in the elderly: an instrumental variable analysis using two preference based instruments. Pharmacoepidemiol Drug Saf. 2010; 19:699–707. [PubMed: 20583208]

19. Schneeweiss S, Setoguchi S, Brookhart A, Dormuth C, Wang PS. Risk of death associated with the use of conventional versus atypical antipsychotic drugs among elderly patients. CMAJ : Canadian Medical Association journal = journal de l'Association medicale canadienne. 2007; 176:627–32.

20. Wang PS, Schneeweiss S, Avorn J, et al. Risk of death in elderly users of conventional vs. atypical antipsychotic medications. New England Journal of Medicine. 2005; 353:2335–41. [PubMed: 16319382]

21. Swanson SA, Hernán MA. Think Globally, Act Globally: An Epidemiologist's Perspective on Instrumental Variable Estimation. Statistical Science. 2014; 29:371–4. [PubMed: 25580054]

22. Small DS, Tan Z, Lorch SA, Brookhart A. Working paper: instrumental variable estimation when compliance is not deterministic: the stochastic monotonicity assumption. 2014 arXiv: 1407.7308.

Figure 1.
Illustration of deterministic compliance types using physician's preference as an instrument
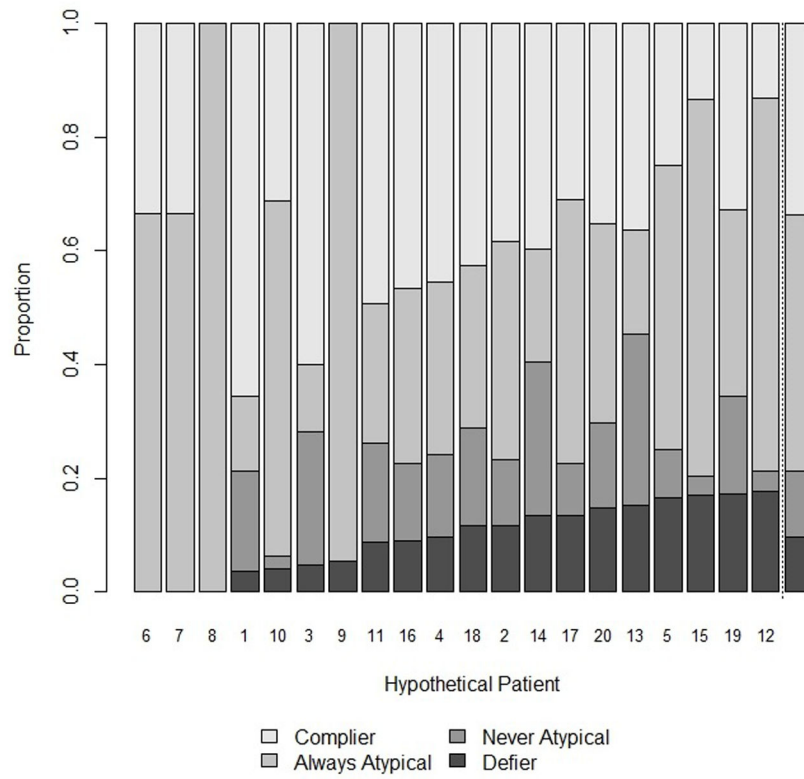
Figure 2.
Distribution of compliance types for each of the 20 hypothetical patients by the distribution of physician preferences and prescribing decision.
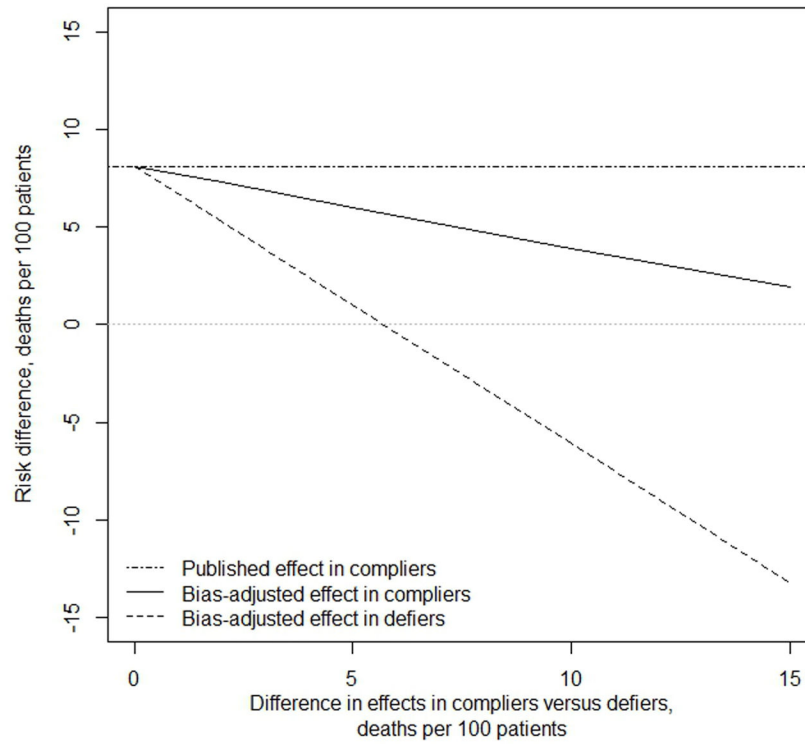
Figure 3.
Effect estimates, as published and bias-adjusted, based on the median of four prior studies comparing conventional to atypical antipsychotic and risk of mortality in the elderly 1. The estimated proportion of defiers was 10%; the estimated proportion of compliers was 34%. Two of the four studies provided information to estimate the proportion of compliers minus the proportion of defiers assuming the proposed instrument was valid: in one study, this was 14%, and in the other it was 30%.

Table 1

Distribution of Prescribing Decisions for Patients #8 and #19 by Physicians' Preference

Preference[a]	Patient #8		Patient #19	
	No. Prescribed Conventional Antipsychotic	No. Prescribed Atypical Antipsychotic	No. Prescribed Conventional Antipsychotic	No. Prescribed Atypical Antipsychotic
Conventional Antipsychotic	0	3	2	2
Atypical Antipsychotic	0	15	11	21

[a] Total number of physicians does not sum to N=53 because physicians were allowed to choose to prescribe no antipsychotic medication.