

We wish to complement the interesting paper by Angrist, Imbens, and Rubin (AIR) by offering several alternative analytic strategies. We focus on randomized drug treatment trials. In their discussion of noncompliance, AIR focuses on estimating the local average treatment effect (LATE), which is the average effect of treatment in the compliers. In contrast, Robins (1989) focused on estimation of the global average treatment effect (ATE) in the entire study population. Both LATE and ATE differ from the intent-to-treat (ITT) parameter, which is the average effect of treatment assignment. We show that in a typical placebo-controlled trial, all three parameters will equal zero under the sharp null hypothesis of no treatment effect. We argue that under the alternative, the ATE parameter can be of greater public health interest than the LATE or ITT parameter. We review results of Robins, Manski, and Balke and Pearl on the estimation of the ATE parameter. We show that in trials comparing a new therapy to a standard therapy, the null hypothesis of bioequivalence does not imply the ITT parameter is zero, and thus the ITT parameter is often of no public health interest. We review results on the estimation of the ATE parameter in bioequivalence trials.

Following AIR, for subject $i = 1, \dots, n$, Z_i denotes the dichotomous randomization indicator (i.e. treatment arm); $D_i(z)$ denotes the actual treatment when randomized to arm z , $z = 0, 1$; $D_i = D_i(Z_i)$ denotes the observed treatment; $Y_i(z, d)$ denotes the outcome that would be observed if randomized to arm z and treatment d were taken, $z = 0, 1, d = 0, 1$; $Y_i = Y_i(Z_i, D_i(Z_i))$ denotes the observed outcome; and expectations are sample averages. This notation incorporates Rubin's stable unit treatment value assumption (SUTVA) assumption. Like AIR, we shall ignore sampling variability by restricting attention to estimands—the large-sample limits of estimators. The foregoing notation is sufficient to describe a trial in which each subject is either on or off a single active treatment; for example, a placebo-controlled trial. However, it is not sufficient to describe a bioequivalence trial that compares a new therapy to a standard therapy, because in a bioequivalence trial, non-compliers may choose to take no drug at all. Thus d needs to be at least trichotomous, corresponding to standard therapy, new therapy, and no therapy.

1. TRIALS WITH A SINGLE ACTIVE TREATMENT

Robins (1989) and AIR considered the analysis of randomized trials of a single active treatment with noncompliance under (1) the exclusion restriction $Y_i(1, d) = Y_i(0, d) \equiv Y_i(d)$, for all i and d , (2) the monotonicity assumption

that $D_i(1) \geq D_i(0)$ for all i , and (3) the random assignment assumption that $D_i(z), Y_i(z, d), z = 0, 1, d = 0, 1$ are jointly independent of Z_i which we write as $\{D_i(z), Y_i(z, d); z = (0, 1), d = (0, 1)\} \perp\!\!\!\perp Z_i$. Robins and AIR also investigated the sensitivity of inferences to violations of these assumptions. (The three assumptions correspond exactly to assumptions (1)–(3) in Robins 1989, p. 123.) Robins studied the average treatment effect (ATE_z) controlling for treatment assignment z ; that is, $E[Y_i(z, 1) - Y_i(z, 0)]$. Under the exclusion restriction, $ATE_1 = ATE_0 \equiv ATE \equiv E[Y_i(1) - Y_i(0)]$. AIR studied the local average treatment effect among the compliers, which is $E[Y_i(1) - Y_i(0) | D_i(1) - D_i(0) = 1]$ under the exclusion restriction. In contrast, the ITT parameter is $E[Y_i(1, D_i(1)) - Y_i(0, D_i(0))]$. Under random assignment, the ITT parameter equals the difference in treatment arm-specific means $E(Y_i | Z_i = 1) - E(Y_i | Z_i = 0)$.

Assuming the exclusion restriction, all three parameters will be zero under the sharp null hypothesis of no causal effect of D on Y ; that is, $Y_i(1) = Y_i(0)$ for all i . Even with noncompliance, the large investment in randomized trials is considered worthwhile because, in contrast to a nonrandomized study, valid tests of the sharp null hypothesis can be obtained from the observed data by comparing treatment arm-specific means. Debates over which of the three parameters should represent the causal parameter of interest arise under alternatives to the sharp null.

A common argument in favor of the ITT parameter is that it corresponds to the overall treatment effect that would be realized if the treatment were actually adopted in the community. But this argument assumes that the noncompliance rate observed in the trial would equal the subsequent rate in the community, which may often not be the case. For example, once the treatment is proven to be efficacious in a trial, then nearly all individuals in the community may be willing to stringently comply with the treatment protocol (Robins 1989). In such a case, if the study subjects are representative of the community, then the ATE parameter, rather than the ITT or LATE parameter, would correspond to the public health parameter of interest. An advantage of the LATE parameter is that it is identifiable under monotonicity, whereas the ATE parameter is not. But unless no subject in the control arm takes active treatment, the subset of the study population for whom the LATE parameter is the treatment effect (i.e., the compliers) is itself nonidentifiable (AIR 1995). As discussed later, the LATE parameter is not identifiable under more complex noncompliance patterns, even if monotonicity holds.

The distribution of the observed data only determines bounds for the ATE parameter. Assuming Y_i dichotomous,

James M. Robins is Professor of Epidemiology and Biostatistics, Harvard School of Public Health, Boston, MA 02115. Sander Greenland is Professor of Epidemiology, UCLA School of Public Health, Los Angeles, CA 90095.

Robins (1989) calculated bounds for ATE_z under the $2^3 = 8$ combinations of the truth or falsity of the exclusion restriction, monotonicity, and random assignment. Related results were independently obtained by Manski (1990, 1994). Interestingly, the bounds for the ATE_z parameter do not depend on the monotonicity assumption. Some argue against reporting bounds for nonidentifiable parameters, because bounds are often so wide as to be useless for making public health decisions. But we view the latter problem as a reason for reporting bounds in conjunction with other analyses: Wide bounds make clear that the degree to which public health decisions are dependent on merging the data with strong prior beliefs. Even when the ITT null hypothesis of equality of treatment arm-specific means is rejected, the bounds may appropriately include zero. If treatment benefits some subjects and harms others, the ATE parameter may be zero even though both the sharp and ITT null hypotheses are false. Conversely, the ATE parameter may be nonzero under the ITT null, seriously complicating the interpretation of tests of the ITT null in trials with substantial noncompliance. But there are times that bounds can be quite informative. For example, Balke and Pearl (1993) reanalyzed data from the Lipid Research Clinic's Coronary Primary Prevention Trial and showed that the Robins-Manski bounds are quite informative, with the lower bound lying far above the null value of zero.

Henceforth we assume that both the exclusion restriction and the random assignment assumption hold. Balke and Pearl (1993) showed that for certain distributions of the observed data, the Robins-Manski bounds, $-1 + \max_z \{ \text{pr}(Y_i = 1, D_i = 1 | Z_i = z) \} + \max_z \{ \text{pr}(Y_i = 0, D_i = 0 | Z_i = z) \} \leq ATE \leq 1 - \max_z \{ \text{pr}(Y_i = 0, D_i = 1 | Z_i = z) \} - \max_z \{ \text{pr}(Y_i = 1, D_i = 0 | Z_i = z) \}$, are not sharp. They derived narrower sharp bounds for these distributions. Specifically, when the bounds do not coincide, the Robins-Manski bounds are sharp under the weak randomization assumption $Y_i(d) \perp\!\!\!\perp Z_i, d = 0, 1$ (Manski 1994), whereas the Balke-Pearl bounds are sharp under the strong randomization assumption $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp Z_i$. The strong randomization assumption is satisfied in a truly randomized study. The even stronger randomization assumption $\{Y_i(0), Y_i(1), D_i(0), D_i(1)\} \perp\!\!\!\perp Z_i$, although appropriate in a randomized trial, does not change the ATE bounds. Balke and Pearl (1993) also showed that there are distributions of the observed data that are incompatible with jointly assuming the exclusion restriction and random assignment of Z .

If one wishes to identify the ATE parameter, further strong nonidentifiable assumptions must be added. Robins (1989, p. 122, assumptions 5-8) provided four different identifying assumptions and computed the ATE parameter under each as a form of sensitivity analysis. Assumption 6 of Robins (1989) implies that the ATE parameter equals the instrumental variable (IV) estimand $\{E(Y_i | Z_i = 1) - E(Y_i | Z_i = 0)\} / \{E(D_i | Z_i = 1) - E(D_i | Z_i = 0)\}$. This result assumes neither monotonicity nor that the subject-specific treatment effects $Y_i(1) - Y_i(0)$ are constant. Assumption 6 of Robins (1989) is the assumption that

both (a) within each treatment arm z , the average treatment effect is the same for the treated ($D_i = 1$) as for the untreated ($D_i = 0$), i.e., $E\{Y_i(1) - Y_i(0) | Z_i = z, D_i = 1\} = E\{Y_i(1) - Y_i(0) | Z_i = z, D_i = 0\}$, and (b) the average treatment effect among the treated is the same for both treatment arms; that is, $E\{Y_i(1) - Y_i(0) | Z_i = 1, D_i = 1\} = E\{Y_i(1) - Y_i(0) | Z_i = 0, D_i = 1\}$. Assumptions (a) and (b) are always true under the sharp null. Robins (1989, sec. 16, 1994) introduced the class of structural nested mean models (SNMM's) for the average effect of treatment on the treated. Assumption (b) is equivalent to assuming a simple SNMM. Robins (1994) proved that the SNMM (b) alone implies that the average treatment effect in the treated $E\{Y_i(1) - Y_i(0) | D_i = 1\}$ is the IV estimand. The additional assumption (a) guarantees that the average treatment effect in the untreated ($D_i = 0$) equals that in the treated ($D_i = 1$), and hence that the ATE parameter equals the IV estimand. An estimand or parameter that is zero if and only if the ITT parameter is zero is called ITT null consistent. Because, when defined, the IV estimand is ITT null consistent, the ATE parameter is also ITT null consistent under assumptions (a) and (b).

2. BIOEQUIVALENCE TRIALS

A critical difference between a trial with a single active therapy and a bio equivalence trial is that in the presence of noncompliance, the sharp null hypothesis of the bioequivalence of the two therapies does not imply equality of treatment arm-specific mean outcomes. Consider a randomized bioequivalence trial in which a new therapy ($D = 1$) is compared to standard proven therapy ($D = 0$). Suppose that all subjects are initially compliant, but 50% of subjects assigned to standard therapy ($Z = 0$) and 20% of subjects assigned to the new therapy ($Z = 1$) later become noncompliant and stop all therapy due to mild, easily palliated side effects. Even if the ITT parameter $E\{Y_i | Z_i = 1\} - E\{Y_i | Z_i = 0\}$ demonstrated a beneficial effect of assignment to the new therapy, the benefit might be wholly attributable to the high noncompliance rate in the standard therapy arm. Thus the ITT test of equality of treatment arm-specific means may not be of regulatory or public health interest. To formalize our point, we consider four treatments: always remain on standard therapy ($D = 0$); always remain on the new therapy ($D = 1$); begin standard therapy, then stop all therapy ($D = 2$); and begin the new therapy, then stop all therapy ($D = 3$). If all therapy terminations were due to mild, easily palliated side effects, then the sharp bioequivalence null hypothesis of medical interest is $Y_i(1) = Y_i(0)$ for all i . But this null hypothesis does not imply the ITT null $E\{Y_i | Z = 1\} = E\{Y_i | Z = 0\}$. The ITT null is implied by the sharp null hypothesis that $Y_i(d) = Y_i$, for all i and $d = 0, 1, 2, 3$; this latter hypothesis is not of interest because it implies that being off therapy was as efficacious as being on the standard proven therapy, which is already known to be false.

In a bioequivalence trial, possible parameters of interest would be the ATE parameter $E\{Y_i(1) - Y_i(0)\}$ or the LATE

parameter $E[Y_i(1) - Y_i(0)|D_i(1) = 1, D_i(0) = 0]$. Neither of these parameters is identifiable from bioequivalence trial data, even under a monotonicity assumption. Sharp bounds for the ATE parameter are $-1 + \text{pr}(Y_i = 1, D_i = 1|Z_i = 1) + \text{pr}(Y_i = 0, D_i = 0|Z_i = 0) \leq \text{ATE} \leq 1 - \text{pr}(Y_i = 0, D_i = 1|Z_i = 1) - \text{pr}(Y_i = 1, D_i = 0|Z_i = 0)$. If one wishes to identify the ATE parameter, then further strong nonidentifiable assumptions must be added. Heyting, Tolboom, and Essers (1992), Robins (1987), Robins and Rotnitzky (1992), and Robins, Rotnitzky, and Zhao (1995) have studied identifying assumptions for ATE in this setting. If the decision to quit therapy is essentially a second randomization (i.e., $Y_i(d) \perp\!\!\!\perp D_i|Z_i$ for $d = 0, 1$) we say that the noncompliance is random. In that case, $E[Y_i(d)]$ is identifiable and equal to $E[Y_i|Z_i = d, D_i = d]$ for $d = 0, 1$. If, as is usually the case, one does not believe compliance is random given Z_i , then one can try to collect data on additional pre- or post-randomization covariates L such that compliance is random conditional on the covariates; that is, $Y_i(d) \perp\!\!\!\perp D_i|Z_i, L_i, d = 0, 1$. Under this assumption, $E[Y_i(d)]$ is identifiable and, for discrete L , equals $\sum_l E[Y_i|Z_i = d, D_i = d, L_i = l] \text{pr}[L_i = l|Z_i = d]$ for $d = 0, 1$. Robins (1987) called this formula the G computation algorithm formula given covariates L . This formula can also be written as the inverse probability of censoring weighted (IPCW) estimand $E[Y_i I(Z_i = d, D_i = d)/\pi_{1d}\pi_{2d}(L_i)]$, where $\pi_{1d} = \text{pr}[Z_i = d]$ and $\pi_{2d}(L_i) \equiv \text{pr}[D_i = d|Z_i = d, L_i]$.

The IPCW estimand is easily generalized to allow investigation of the sensitivity of the estimate of $E[Y_i(d)]$ to the assumption $Y_i(d) \perp\!\!\!\perp D_i|Z_i, L_i$. Let $\pi_{2d}(L_i, y) = \text{pr}[D_i = d|Z_i = d, L_i, Y_i(d) = y]$ be the probability of treatment $D = d$ given $Z_i = d, L_i$, and $Y_i(d) = y$. Note that $\pi_{2d}(L_i, y)$ is not identifiable, because we do not observe $Y_i(d)$ for subjects for whom $D_i \neq d$. In a sensitivity analysis, we select plausible functions $\pi_{2d}(L_i, y)$ based on our prior beliefs. For a given $\pi_{2d}(L_i, y)$, $E[Y_i(d)]$ is given by the IPCW estimand with $\pi_{2d}(L_i, Y_i)$ substituted for $\pi_{2d}(L_i)$. Robins et al. (1995, p. 118) also considered estimation of $E[Y_i(d)]$ under the weaker assumption that $\pi_{2d}(L_i, y)$ followed a parametric model such as logit $\pi_{2d}(L_i, y) = \alpha_0 + \alpha_1 L_i + \alpha_2 Y_i$.

3. CONCLUSION

The ATE parameter can be of greater public health inter-

est than either the LATE or ITT parameter. We have proposed methods for setting bounds and for constructing estimators of the ATE parameter both in single active treatment trials and in bioequivalence trials. Both structural nested models and IPCW estimators can be applied to complex trials with randomized and nonrandomized time-dependent treatments, noncompliance, and dependent censoring with either failure time or repeated-measures outcomes (Robins 1989, 1993, 1994; Robins and Greenland 1994).

ADDITIONAL REFERENCES

- Balke, A., and Pearl, J. (1993), "Bounds on Treatment Effects of Studies With Imperfect Compliance," Technical Report R-199-J, UCLA Cognitive Systems Laboratory.
- Heyting, A., Tolboom, J. T. B. M., and Essers, J. G. A. (1992), "Statistical Handling of Drop-Outs in Longitudinal Clinical Trials," *Statistics in Medicine*, 11, 2043-2062.
- Manski, C. F. (1990), "Nonparametric Bounds on Treatment Effects," *American Economic Reviews, Papers and Proceedings*, 80, 319-323.
- (1994), "Learning About Social Programs From Experiments With Random Assignment of Treatment," Social Science Research Institute, 9505, University of Wisconsin.
- Robins, J. M. (1987), "Addendum to A New Approach to Causal Inference in Mortality Studies With Sustained Exposure Periods—Application to Control of the Healthy Worker Survivor Effect," *Computers and Mathematics with Applications*, 14, 923-945.
- (1989), "The Analysis of Randomized and Nonrandomized AIDS Treatment Trials Using a New Approach to Causal Inference in Longitudinal Studies," in *Health Service Research Methodology: A Focus on AIDS*, eds. L. Sechrest, H. Freeman, and A. Mulley, Washington, D.C.: U.S. Public Health Service, pp. 113-159.
- (1993), "Analytic Methods for Estimating HIV Treatment and Co-factor Effects," in *Methodological Issues of AIDS Mental Health Research*, eds. D. G. Ostrow and R. Kessler, New York: Plenum Publishing, pp. 213-290.
- (1994), "Correcting for Noncompliance in Randomized Trials Using Structural Nested Mean Models," *Communications in Statistics*, 23, 2379-2412.
- Robins, J. M., and Greenland, S. (1994), "Adjusting for Differential Rates of PCP Prophylaxis in High- Versus Low-Dose AZT Treatment Arms in an AIDS Randomized Trial," *Journal of the American Statistical Association*, 89, 737-749.
- Robins, J. M., and Rotnitzky, A. (1992), "Recovery of Information and Adjustment for Dependent Censoring Using Surrogate Markers," in *AIDS Epidemiology—Methodological Issues*, eds. N. Jewell, K. Dietz, and V. Farewell, Boston: Birkhäuser, pp. 297-331.
- Robins, J. M., Rotnitzky, A., Zhao, L.-P. (1995), "Analysis of Semiparametric Regression Models for Repeated Outcomes in the Presence of Missing Data," *Journal of the American Statistical Association*, 90, 106-121.