



Marginal Mean Models for Dynamic Regimes

Author(s): S. A. Murphy, M. J. van der Laan, J. M. Robins, Conduct Problems Prevention Research Group

Source: *Journal of the American Statistical Association*, Vol. 96, No. 456 (Dec., 2001), pp. 1410-1423

Published by: American Statistical Association

Stable URL: <http://www.jstor.org/stable/3085909>

Accessed: 24/11/2008 15:03

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=astata>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



American Statistical Association is collaborating with JSTOR to digitize, preserve and extend access to *Journal of the American Statistical Association*.

<http://www.jstor.org>

Marginal Mean Models for Dynamic Regimes

S. A. MURPHY, M. J. VAN DER LAAN, J. M. ROBINS,
and CONDUCT PROBLEMS PREVENTION RESEARCH GROUP

A dynamic treatment regime is a list of rules for how the level of treatment will be tailored through time to an individual's changing severity. In general, individuals who receive the highest level of treatment are the individuals with the greatest severity and need for treatment. Thus, there is planned selection of the treatment dose. In addition to the planned selection mandated by the treatment rules, staff judgment results in unplanned selection of the treatment level. Given observational longitudinal data or data in which there is unplanned selection of the treatment level, the methodology proposed here allows the estimation of a mean response to a dynamic treatment regime under the assumption of sequential randomization.

KEY WORDS: Causal inference; Confounding; Dynamic treatment regime; Nondynamic treatment regime.

1. INTRODUCTION

Dynamic treatment regimes are individually tailored treatments that provide treatment to individuals only when and if treatment is needed and that adjust the level of treatment to the individual's need. In a dynamic regime, rules for how the treatment level and type should vary with time are specified before treatment is started; these rules are based on time-varying measurements of subject-specific need for treatment. Thus, the assignment to treatment depends on past individual information. Dynamic treatment regimes are potentially attractive to public policy makers because they treat only subjects who show need for treatment, freeing public and private funds for more intensive treatment of the needy.

The goal of this article is to provide an estimation method for the mean response that would have been observed if, contrary to fact, the entire study population followed a particular dynamic regime. The available data are observational; that is, in the available data, the treatment level may depend on past individual information that is not used in, or that is not used according to, the particular dynamic regime rules. We employ a structural model; thus the model is a model for the marginal mean of counterfactual or potential responses.

This work was motivated by the Fast Track prevention program. This is an ongoing randomized trial of a complex preventive intervention versus control. The intervention was designed to prevent the emergence of and reduce the level of conduct disorders and drug use in children at risk because of elevated behavior problems (Bierman, Greenberg, and Conduct Problems Prevention Research Group 1996; to appear a, b; McMahon, Slough, and Conduct Problems Prevention Research Group 1996). Part of the intervention involved implementation of a dynamic treatment regime designed to improve family functioning. At the end of each semester,

beginning with the spring semester of first grade, the family counselor filled out a six-item "home visiting process measure," which described the quality of parenting and family functioning. On the basis of the home visiting process measure total score, home visiting assignments for the next semester were made. A simplified version is as follows. A score of 17 or higher corresponds to an assignment of 4 home visits, a score from 9 to 16 corresponds to an assignment of 8 home visits, and a score of 8 or below corresponds to an assignment of 16 home visits during the following semester. Thus, the rule for assigning dose was

$$d_t(S_{t-1}) = 16I\{S_{t-1} \leq 8\} + 8I\{9 \leq S_{t-1} \leq 16\} \\ + 4I\{17 \leq S_{t-1}\}, \quad t = 1, 2, 3, 4,$$

where S_t is an assessment of severity, the home visiting process measure score, at the end of the t th semester, with low values indicating greater severity. Staff were told that in exceptional cases they might need to deviate from the rule for assigning home visiting level; in practice, staff assignments deviated from the rule for approximately 50% of the intervention children. In Section 6, our goal is to estimate what the marginal effect of the dynamic treatment regime would have been if staff followed the rules exactly. We demonstrate that by identifying the intervention group (with home visiting assignments as implemented) as an observational study, one can use the methods developed here to achieve this goal.

Dynamic treatment regimes are routinely assigned to subjects participating in randomized medical trials when there is danger of serious side effects that necessitate cessation of treatment or when the optimal dosage at a given time should depend on the subject's clinical status or laboratory values. For example, in the active treatment arm of the Systolic Hypertension in the Elderly Program randomized trial (Cooperative Research Group 1988; Borhani et al. 1991), subjects were assigned to a dynamic treatment regime that was intended to minimize the amount of medication required to maintain a subject's systolic blood pressure (SBP) at or below a predetermined goal. The treatment schedule was complex. Initially, all treated subjects were assigned the same dosage of antihypertensive medication; at 8 weeks, subjects whose SBP exceeded

S. A. Murphy is Professor, Department of Statistics, University of Michigan, Ann Arbor, MI 48109-1285 (E-mail: samurphy@umich.edu). M. J. van der Laan is Professor, Division of Biostatistics, University of California, Berkeley, CA, 94720 (E-mail: laan@stat.berkeley.edu). J. M. Robins is Professor, Department of Epidemiology and Biostatistics, Harvard University, Boston, MA 02115 (E-mail: robins@epinet.harvard.edu). Conduct Problems Prevention Research Group, members include K. L. Bierman (Pennsylvania State University), J. D. Coie (Duke University), M. T. Greenberg (Pennsylvania State University), J. E. Lochman (University of Alabama), R. J. McMahon (University of Washington), and E. Pinderhughes (Vanderbilt University). The research of the first author was partially supported by NIDA grant A P50 DA 10075, DMS 9802885, and SBR 9811983. The research of the second author was partially supported by NIAID grant 1R01AI46182. The research of the third author was partially supported by NIAID grant RO1AI32475.

the goal had their dosage doubled; at 16 weeks, subjects whose SBP still exceeded the goal had a second antihypertensive added. At 24 weeks, a subject whose SBP continued to exceed the goal had their dose of the second antihypertensive doubled. The treatment schedule specified that subjects with evidence of severe clinical or laboratory toxicity would first have their dose of antihypertensive medications reduced; if the toxicity persisted, they would be switched to an alternate medication. As in the Fast Track program, this dynamic treatment is used to tailor the dosage to the subject. Furthermore, the use of rules for reducing dosage or switching medications because of side effects is a useful aid in reducing unexplained non-compliance. Indeed, subjects whose medication is switched in accordance with the rules are compliers.

In this article, we consider methodology for estimation of the mean response to both nonrandom and random dynamic treatment regimes when the available data are observational. In a nonrandom dynamic treatment regime, the rules specifying the treatment level, as outputs of present need, are nonrandom. Both the Fast Track program and the Systolic Hypertension in the Elderly Program use nonrandom dynamic treatment regimes. In a random dynamic treatment regime, the time-specific treatment level is drawn from a conditional probability distribution depending only on measures of present needs. In Section 3, we provide quantitative definitions of these two types of dynamic treatment regime. The modeling of responses to dynamic treatment regimes based on observational data has received very little statistical attention outside of a series of articles by Robins (1986, 1989, 1993, 1997). In these papers, Robins considered the use of nested structural models to estimate a variety of conditional treatment effect parameters in the distribution of the potential response. Also, Robins (1993) discussed estimating parameters in the response to a nonrandom dynamic treatment regime by censoring the subject the first time the subject's treatment differs from the treatment specified by the dynamic regime.

In Section 2, we review the underlying causal theory based on counterfactual or potential outcomes. In section 3, we discuss the assumption of sequential randomization along with the definitions of nonrandom and random dynamic treatment regimes. Section 4 uses the work in the previous sections to precisely specify the estimand in terms of the observational data distribution. We provide an estimating function in Section 5, and then in the last section we return to the Fast Track example.

2. POTENTIAL OUTCOMES

We use counterfactual or potential outcome models to quantify the desired treatment effect and to state assumptions. Neyman (1990) introduced counterfactual outcomes to analyze the causal effect of time-independent treatments in randomized studies. Rubin (1978) explicated Neyman's ideas and extended Neyman's work to the analysis of causal effects of time-independent treatments from observational data. Robins (1986, 1987) proposed a formal theory of causal inference that extended both Neyman's and Rubin's work to assess the direct and indirect-effects of time-varying treatments from experimental and observational longitudinal studies. We use these

works to specify our observations in a unified way regardless of the manner in which treatment is selected or assigned.

Suppose the treatment lasts for K intervals; during this period, intermediate outcomes may be measured, and at the end of the K intervals, a response is measured. We denote the treatment regime/vector across the K intervals by $\bar{a}_K = (a_1, a_2, \dots, a_K)$; a_t is the level of treatment in the t th interval. In general, we use a bar over a variable to denote that variable and all past values of the same variable, so $\bar{a}_t = (a_1, \dots, a_t)$. Let \mathcal{A}_K be the collection of all possible treatment vectors. Corresponding to each fixed value of the treatment vector, \bar{a}_K , we conceptualize a potential (counterfactual) response denoted by $Y(\bar{a}_K)$, where $Y(\bar{a}_K)$ is the response at the end of the K th interval that a subject would have had if he or she followed the treatment regime, \bar{a}_K . The set of all potential responses is $\{Y(\bar{a}_K) : \bar{a}_K \text{ varying in } \mathcal{A}_K\}$. This statement implicitly assumes that there is no need to index Y by others' treatment or by the mechanism by which the treatment was selected; this is the stable unit treatment assumption [SUTVA; see Rubin (1986)]. In the time-varying treatment setting, this is the consistency assumption of Robins (1997). In particular, we assume that treatment of a subject does not influence any of the outcomes of any other subject (Cox 1958); indeed, we will make the stronger assumption that the observations on the sample of subjects are independent draws from one distribution. This assumption allows us to conceptualize a single potential outcome or response corresponding to each possible $\bar{a}_K \in \mathcal{A}_K$ and thus permit a well-defined, simple notation for an intermediate outcome and final response [see Robins (1986) and, for time-independent treatment, Rubin (1986) and Angrist, Imbens, and Rubin (1996)].

SUTVA may well be violated in the Fast Track data. In addition to home visiting, the Fast Track intervention included friendship groups, in which small numbers of the intervention children were brought together to improve social skills. Thus, in close friendships, the effect of home visiting on one child may affect the friend's response. Alternately, whether a dominant child is receiving home visits may alter the dynamics of the friendship groups and thus alter other children's response to treatment. We do not deal with these complications here, because an appropriate generalization of the proposed methodology would distract from the main points of this article.

Capital letters are used to denote random variables. Treatment selection and assignment are allowed to be stochastic and are denoted by the random vector \bar{A}_K . So \bar{A}_K takes values in \mathcal{A}_K . Denote the variables available at the beginning of time interval t ($t > 0$) by $L_t = (S_t, V_t)$, where S_t is the severity vector and V_t is a vector of auxiliary information. In general, both V_t and S_t may be summaries of information that are available before and at the beginning of time interval t . (S_t, V_t) is an intermediate outcome of (past) treatment. Denote the initial, time 0, variables by $L_0 = (Z, S_0, V_0)$, where Z consists of variables defining interesting subpopulations and S_0, V_0 have the same interpretation as earlier. Additionally, in randomized studies Z may contain the indicator of the treatment regime to which a study subject was assigned. Corresponding to each treatment vector, \bar{a}_K , SUTVA implies that the intermediate outcome at time t can be written unambiguously as $(S_t(\bar{a}_t), V_t(\bar{a}_t))$ (Robins 1997) and the potential response as $Y(\bar{a}_K)$.

A subject's potential outcomes, that is, the complete severity and response data, are $O_{sr} = \{Z, S_0, S_1(a_1), \dots, S_{K-1}(\bar{a}_{K-1}), Y(\bar{a}_K); \bar{a}_K \in \mathcal{A}_K\}$, and the complete auxiliary data are $O_{aux} = \{V_0, V_1(a_1), \dots, V_{K-1}(\bar{a}_{K-1}); \bar{a}_K \in \mathcal{A}_K\}$. Assuming SUTVA, we may write the observed data in a unified way regardless of the manner in which treatment is selected or assigned. In both randomized and observational studies, the observed data are the pretreatment information plus the potential outcomes corresponding to the treatment pattern \bar{A}_K . That is, we observe $X = \{Z, S_0, V_0, A_1, S_1(A_1), V_1(A_1), A_2, \dots, S_{K-1}(\bar{A}_{K-1}), V_{K-1}(\bar{A}_{K-1}), A_K, Y(\bar{A}_K)\}$. In the following, we use either Y or $Y(\bar{A}_K)$ to denote the observed response at the end of time interval K , either S_t or $S_t(\bar{A}_t)$ to denote the observed severity at the beginning of time interval t , and either V_t or $V_t(\bar{A}_t)$ to denote the observed auxiliary variables available at the beginning of time interval t . Thus, we may denote an individual's data in order of observation by

$$X = \{(Z, S_0, V_0), A_1, (S_1, V_1), A_2, \dots, (S_{K-1}, V_{K-1}), A_K, Y\}.$$

3. DYNAMIC REGIMES AND SEQUENTIAL RANDOMIZATION

In a dynamic treatment regime, the rule for treatment assignment in each of the K intervals may be random or nonrandom; in either case, the rule depends only on prior severity. If the rules are nonrandom, denote the regime by the K vector, $\bar{d}_K = (d_1, \dots, d_K)$; at time t , treatment assignment is $A_t = d_t(S_{t-1})$. (The Fast Track example uses nonrandom rules.) If the rules are random (stochastic), the treatment regime is characterized by a K vector of stochastic rules or, equivalently, treatment assignment probabilities, say, $\bar{p}_K = (p_1, \dots, p_K)$; at time t , treatment assignment is drawn from the conditional distribution, $p_t(\cdot | S_{t-1})$. We shall denote by $P_{\bar{p}_K}$ the distribution of $(O_{sr}, O_{aux}, \bar{A}_K)$ had the entire study population been subjected to the dynamic regime \bar{p}_K . The associated expectation is denoted by $E_{\bar{p}_K}$. Note that a nonrandom dynamic treatment regime, \bar{d}_K , is just a special case of a random dynamic treatment regime with $p_t(a | S_{t-1})$ degenerate at the point, $a = d_t(S_{t-1})$. Henceforth, we do not, except where necessary, distinguish in our notation between random and nonrandom dynamic treatment regimes.

Dynamic treatment regimes should not be confused with nondynamic treatment regimes. In a nondynamic treatment regime, the variation in treatment level across time and subjects is specified before treatment onset. In other words, a nondynamic treatment regime is a special case of a dynamic treatment regime in which the treatment assignments do not vary by posttreatment observations (i.e., the rule d_t is not a function of posttreatment information in S_{t-1} and is at most a function of pretreatment variables, L_0). Concrete examples of nondynamic treatment regimes are "assign 16 home visits in every semester following kindergarten," or "assign 16 home visits in the first semester following kindergarten and 8 home visits in each semester thereafter." Regression analyses that model the mean of $Y(\bar{a}_K)$ as \bar{a}_K varies are regression analyses comparing different nondynamic treatment regimes, each corresponding to a different treatment pattern, \bar{a}_K . For an example of this type of analysis, see Robins, Hernán, and Brumback (2000) and Hernán, Brumback, and Robins (1998).

Denote the distribution of $(O_{sr}, O_{aux}, \bar{A}_K)$ in our observational study by P_{obs} and the expectation with respect to this distribution by E_{obs} . Note that the marginal distribution of the complete data (O_{sr}, O_{aux}) is the same for both distributions, $P_{\bar{p}_K}$ and P_{obs} . Only the marginal or conditional distribution of the treatment (\bar{A}_K) varies between our observational distribution and the distribution $P_{\bar{p}_K}$. Thus, we denote the marginal distribution of (O_{sr}, O_{aux}) by P without a subscript and expectation with respect to this marginal distribution by E without a subscript.

If the study population follows the dynamic treatment regime, \bar{p}_K , then A_t is conditionally independent of the potential outcomes, (O_{sr}, O_{aux}) , conditional on S_{t-1} . However in our observational study, A_t may not be conditionally independent of the potential outcomes, (O_{sr}, O_{aux}) , conditional on S_{t-1} , because there may be unmeasured confounders that determine treatment and are associated with the potential outcomes. In general, assumptions, in addition to SUTVA, about this distributional relationship (based on substantive knowledge) must be used to identify causal effects and thus permit causal inference. [For discussion, see Robins (1997, sec. 11)]. In this article, we use additional confounders in place of the unmeasured confounders and we assume

Sequential Randomization. For each $t = 1, \dots, K$, A_t is independent of O_{sr} given $\{L_0, A_1, L_1, A_2, \dots, L_{t-1}\}$.

That is, the auxiliary data, the V 's, are sufficiently rich so the observational distribution satisfies the sequential randomization. This means that within levels of $L_0, A_1, L_1, A_2, \dots, L_{t-1}$, we assume that the subjects with different levels of treatment do not vary systematically by the potential severity measures or responses. See Robins (1997) for further discussion of sequential randomization.

We chose to assume sequential randomization of the treatment level \bar{A}_K with respect to O_{sr} only and not with respect to both (O_{sr}, O_{aux}) . There are two reasons for this. First, there are practical, real-life examples in which sequential randomization of \bar{A}_K with respect to O_{sr} may be plausible, but sequential randomization of \bar{A}_K with respect to O_{aux} is clearly untrue [see example 1 of Robins (1987)]. The second reason is more philosophical. If a time-varying correlate of both treatment level selection and response becomes available, and the sequential randomization assumption is more plausible when the V 's include this time-varying covariate, then we would like to include this covariate in the V 's without having to make the additional assumption that the treatment levels are sequentially randomized with respect to this covariate as well.

4. IDENTIFICATION OF THE ESTIMAND IN TERMS OF THE OBSERVATIONAL DATA

Recall that the goal of this article is to estimate, by using observational data, the mean response that would have been observed had, contrary to fact, the entire study population followed a particular dynamic regime with rules \bar{p}_K . Thus, we wish to estimate $E_{\bar{p}_K}[Y|Z]$ for Z denoting an interesting subpopulation. However, we do not have data from a random sample of individuals following the dynamic regime, \bar{p}_K . Rather, the available data are observational (from the P_{obs} distribution). The following lemma allows us to write $E_{\bar{p}_K}[Y|Z]$

in terms of the observational distribution, thus demonstrating that $E_{\bar{p}_K}[Y|Z]$ is identifiable from the available data. The lemma achieves this by providing a Radon–Nikodym derivative connecting the distribution $P_{\bar{p}_K}$ resulting from implementing the dynamic regime to the observational distribution P_{obs} . A secondary consequence of the lemma is to provide a way to construct estimating equations for $E_{\bar{p}_K}[Y|Z]$.

Let $\pi_t(\cdot|\bar{a}_{t-1}, \bar{l}_{t-1})$ represent the conditional probability mass function of A_t given $\bar{A}_{t-1} = \bar{a}_{t-1}$, $\bar{L}_{t-1} = \bar{l}_{t-1}$ in the observational distribution. Define

$$B(o_{sr}, \bar{a}_K) = P_{obs} \left[\prod_{t=1}^K \pi_t(a_t|\bar{a}_{t-1}, \bar{L}_{t-1}(\bar{a}_{t-1})) > 0 \mid O_{sr} = o_{sr} \right]$$

and

$$W_{\bar{p}_j}(\bar{a}_j, \bar{l}_{j-1}) = \prod_{t=1}^j \frac{p_t(a_t|s_{t-1})}{\pi_t(a_t|\bar{a}_{t-1}, \bar{l}_{t-1})}$$

for any $j \leq K$.

Lemma 4.1. Assume that the observational distribution P_{obs} satisfies the sequential randomization assumption. If

$$P_{\bar{p}_K}[B(O_{sr}, \bar{A}_K) = 1] = 1, \quad (4.1)$$

then the distribution of $(Y, \bar{S}_{K-1}, \bar{A}_K, Z)$ under $P_{\bar{p}_K}$ is absolutely continuous with respect to the distribution of $(Y, \bar{S}_{K-1}, \bar{A}_K, Z)$ under P_{obs} , and a version of the Radon–Nikodym derivative is

$$E_{obs}[W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1})|Y=y, \bar{S}_{K-1}=\bar{s}_{K-1}, \bar{A}_K=\bar{a}_K, Z=z]. \quad (4.2)$$

See the Appendix for a proof of this lemma.

An intuitive interpretation of the lemma's assumption (4.1) is that, "Any treatment pattern that can result in the implementation of the dynamic treatment regime, \bar{p}_k , must also have a positive probability of occurring in the observational study." This assumption is eminently sensible because if a particular treatment pattern cannot occur in the observational study, then inference involving the response to this treatment pattern requires further knowledge and assumptions.

As stated, there are two important consequences of the fact that (4.2) is a Radon–Nikodym derivative. By using the lemma, we have that the estimand can be written in terms of the observational distribution as

$$\begin{aligned} E_{\bar{p}_K}[Y|Z] &= E_{obs} \left[E_{obs} [W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1})|Y, \bar{S}_{K-1}, \bar{A}_K, Z] Y|Z \right] \\ &= E_{obs} [W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}) Y|Z]. \end{aligned} \quad (4.3)$$

More explicitly, the right side is, in terms of conditional densities,

$$\begin{aligned} E_{\bar{p}_K}[Y|Z=z] &= \sum_{\bar{a}_K \in \mathcal{A}} \int_{\bar{l}_{K-1}} \int_y y f_Y(y|\bar{a}_K, \bar{v}_{K-1}, \bar{s}_{K-1}) \\ &\quad \times \prod_{t=1}^K p_t(a_t|s_{t-1}) \prod_{t=1}^{K-1} f_t(l_t|\bar{a}_t, \bar{l}_{t-1}) \\ &\quad \times f_0(l_0|z) dy d\bar{l}_{K-1}, \end{aligned} \quad (4.4)$$

where f_Y, f_K, \dots, f_0 denote conditional densities for the observational distribution. The conditional density of Y given \bar{A}_K, \bar{L}_{K-1} is f_Y , and the conditional density of L_t given $\bar{A}_{t-1}, \bar{L}_{t-1}$ is f_t . Thus, the conditional mean is a function of the observational distribution via the conditional densities f_Y

and f_{K-1}, \dots, f_0 . This is Robins's (1986) G-computation formula for a dynamic regime.

The second consequence of the fact that (4.2) is a Radon–Nikodym derivative is that given any $P_{\bar{p}_K}$ -unbiased estimating equation for parameters in the distribution $P_{\bar{p}_K}$, we can weight this estimating equation by the Radon–Nikodym derivative and thereby produce an unbiased estimating equation distribution for the same parameters but now using data from the observational study, the P_{obs} distribution. Suppose for the moment that $\bar{\pi}_K$ is known. Let $\mu(\beta, Z)$ be a parameterization of $E_{\bar{p}_K}[Y|Z]$, where β is an unknown vector parameter and μ is a known function. Given data from $P_{\bar{p}_K}$, a simple estimating function for β might be based on $\dot{\mu}(\beta, Z)(Y - \mu(\beta, Z))$, where $\dot{\mu}(\beta, Z)$ is the derivative of μ with respect to β . Given data from the P_{obs} distribution, we weight to produce the unbiased estimating function,

$$\mathbb{P}_n [W_{\bar{p}_K}(\bar{A}_K, \bar{L}_K) \dot{\mu}(\beta, Z)(Y - \mu(\beta, Z))], \quad (4.5)$$

where for a function g , $\mathbb{P}_n g(X)$ is defined as the average of g over the observational sample ($\mathbb{P}_n g(X) = 1/n \sum_{i=1}^n g(X_i)$).

Note that the denominator of $W_{\bar{p}_K}$ is informally the probability of having the treatment pattern one did indeed have in the observational study. Robins (1998) introduced the idea of weighting estimating functions by the inverse of this probability to estimate parameters of marginal structural models and direct-effect structural nested models for time-varying treatments. Robins refers to such estimators as inverse probability of treatment-weighted estimators (Robins 1998, 1999).

In the special case in which \bar{p}_K represents a nonrandom dynamic treatment regime, an alternate strategy is to use the inverse probability of censoring weighted estimators as considered by Robins and Rotnitzky (1992) and Rotnitzky and Robins (1995). Instead of modeling the treatment selection given past information, they model the adherence or nonadherence to the treatment regime given past information. Subjects are censored when their treatment first deviates from the dynamic treatment regime; these subjects receive a weight of zero. Subjects whose treatment patterns match the dynamic treatment regime receive a weight that is the product of the modeled adherence probabilities. Note that in this article, for each t , we could separate the selection probability $\pi_t(a_t|\bar{a}_{t-1}, \bar{l}_{t-1})$ into the product of the probability of adherence or nonadherence to the treatment regime given $\bar{A}_{t-1} = \bar{a}_{t-1}$, $\bar{L}_{t-1} = \bar{l}_{t-1}$ times the probability of $A_t = a_t$ given adherence or nonadherence to the treatment regime and $\bar{A}_{t-1} = \bar{a}_{t-1}$, $\bar{L}_{t-1} = \bar{l}_{t-1}$.

If the observational distribution satisfies sequential randomization, then $\mu(\beta, Z) = E_{\bar{p}_K}[Y|Z]$ and β can be interpreted as a parameter from the $P_{\bar{p}_K}$ distribution. That is, the sequential randomization assumption allows us to interpret $\mu(\beta, Z)$ as the mean response to the regime \bar{p}_K . Alternately, we may view $\mu(\beta, Z)$ simply as a parameterization of (4.4); this viewpoint does not require sequential randomization and does not involve consideration of potential outcomes. We use this latter viewpoint to form the estimating equation and when considering efficiency improvements.

5. AN ESTIMATING EQUATION

Recall that we wish to estimate $E_{\bar{p}_K}[Y|Z]$ for Z denoting an interesting subpopulation and the available data are observational (from the P_{obs} distribution). Assuming SUTVA, (4.1) and sequential randomization, $E_{\bar{p}_K}[Y|Z]$ is equal to (4.4) or, equivalently, (4.3). We separate this interpretation from the estimation. In this section, we simply consider a parameterization of (4.4) denoted by $\mu(\beta, Z)$. To estimate β , we no longer need SUTVA, sequential randomization, or (4.1); indeed, we no longer need to assume the existence of potential outcomes. However, we assume regularity conditions for the asymptotic theory and that

$$E_{obs}W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}) = 1. \tag{5.1}$$

This is easily shown to be a consequence of (4.1); however, (5.1) makes no reference to potential outcomes.

Note that $\mu(\beta, Z)$ is a function of the conditional densities f_Y and f_{K-1}, \dots, f_0 . On the other hand, the likelihood of the observational data can be written as a product of conditional densities:

$$\begin{aligned} & \left(f_Y(y|\bar{a}_K, \bar{l}_{K-1}) \prod_{t=1}^{K-1} f_t(l_t|\bar{a}_t, \bar{l}_{t-1})f_0(l_0|z) \right) \\ & \times \prod_{t=1}^K \pi_t(a_t|\bar{a}_{t-1}, \bar{l}_{t-1}), \tag{5.2} \end{aligned}$$

where $\pi_t(\cdot|\bar{a}_{t-1}, \bar{l}_{t-1})$ represents the conditional probability mass function of A_t given $\bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_{t-1} = \bar{l}_{t-1}$ in the observational distribution. Thus, the estimand, (4.4), depends only on the first term in the likelihood. The second term in the likelihood concerns only the treatment selection probabilities ($\bar{\pi}_K$), which are nuisance parameters. An optimal estimating function (one resulting in an efficient estimator of β) will be orthogonal to the score functions for the nuisance parameters (Bickel et al. 1993). This means we have the potential to improve estimators based on (4.5) by replacing (4.5) by itself minus its projection on the score functions for the treatment selection probabilities (Robins 1999). We do this in the following. Recall that the observed data for one person is denoted by X . Set $U(X, \beta)$ equal to $W_{\bar{p}_K}(\bar{A}_K, \bar{L}_K)\dot{\mu}(\beta, Z)(Y - \mu(\beta, Z))$. The estimating function (4.5) is just $\mathbb{P}_n U(X, \beta)$. The projection of $U(X, \beta)$ on the score functions of the treatment selection probabilities is given by $\sum_{t=1}^K E_{obs}[U(X, \beta)|\bar{L}_{t-1}, \bar{A}_t] - E_{obs}[U(X, \beta)|\bar{L}_{t-1}, \bar{A}_{t-1}]$ (Robins 1999).

After algebraic simplification, we have that

$$\begin{aligned} & U(X, \beta) - \sum_{t=1}^K E_{obs}[U(X, \beta)|\bar{L}_{t-1}, \bar{A}_t] \\ & \quad + E_{obs}[U(X, \beta)|\bar{L}_{t-1}, \bar{A}_{t-1}] \\ & = W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1})\dot{\mu}(\beta, Z)(Y - \mu(\beta, Z)) \\ & \quad - \sum_{t=1}^K W_{\bar{p}_t}(\bar{A}_t, \bar{L}_{t-1})\dot{\mu}(\beta, Z)(g_t(\bar{A}_t, \bar{L}_{t-1}) - \mu(\beta, Z)) \\ & \quad + \sum_{t=1}^K \sum_{a_t} \pi_t(a_t|\bar{A}_{t-1}, \bar{L}_{t-1})W_{\bar{p}_t}(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}) \\ & \quad \times \dot{\mu}(\beta, Z)(g_t(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}) - \mu(\beta, Z)), \tag{5.3} \end{aligned}$$

where $W_{\bar{p}_t}(\bar{a}_t, \bar{l}_{t-1})$ represents $\prod_{u=1}^t p_u(a_u|s_{u-1})/(\pi_u(a_u|\bar{a}_{u-1}, \bar{l}_{u-1}))$ and

$$\begin{aligned} & g_K(\bar{A}_K, \bar{L}_{K-1}) = E_{obs}(Y|\bar{A}_K, \bar{L}_{K-1}), \\ & g_t(\bar{A}_t, \bar{L}_{t-1}) = E_{obs}\left(\sum_{a_{t+1}} p_{t+1}(a_{t+1}|S_t) \right. \\ & \quad \left. \times g_{t+1}(a_{t+1}, \bar{A}_t, \bar{L}_t)|\bar{A}_t, \bar{L}_{t-1}\right) \tag{5.4} \end{aligned}$$

for $t = 1, \dots, K - 1$. If the parameterization of $\mu(\beta, Z)$ is saturated, then the right side of (5.3) is the efficient influence function for β . See the appendix for proofs.

Note that whereas the treatment selection probabilities for the $P_{\bar{p}_K}$ are known—they are \bar{p}_K —in our observational study, the treatment selection probabilities, $\bar{\pi}_K$, are unknown and thus must be estimated from the observational data. Additionally, we will need to estimate $\{g_t(\bar{a}_t, \bar{l}_{t-1}), t = 1, \dots, K\}$. Given a model for the treatment selection probabilities, say, $\pi_t(a_t|\bar{a}_{t-1}, \bar{l}_{t-1}; \alpha_w), t = 1, \dots, K$, with parameter α_w , we can use maximum likelihood based on (5.2) to estimate α_w . We maximize

$$\mathbb{P}_n \sum_{t=1}^K \log \pi_t(A_t|\bar{A}_{t-1}, \bar{L}_{t-1}; \alpha_w) \tag{5.5}$$

to get $\hat{\alpha}_w$ and then by substitution the vector $\{\hat{\pi}_t, t = 1, \dots, K\}$. Next, we parameterize $\{g_t(\bar{a}_t, \bar{l}_{t-1}), t = 1, \dots, K\}$ with the vector parameter α_g and denote the parameterization by $\{g_t(\bar{a}_t, \bar{l}_{t-1}; \alpha_g), t = 1, \dots, K\}$. A variety of methods may be used to estimate α_g ; we use the following estimating equation, that is, solve

$$\begin{aligned} & 0 = \mathbb{P}_n(Y - g_K(\bar{A}_K, \bar{L}_{K-1}; \alpha_g))\frac{\partial}{\partial \alpha_g} g_K(\bar{A}_K, \bar{L}_{K-1}; \alpha_g) \\ & \quad + \sum_{t=1}^{K-1} \mathbb{P}_n\left(\sum_{a_{t+1}} p_{t+1}(a_{t+1}|S_t)g_{t+1}(a_{t+1}, \bar{A}_t, \bar{L}_t; \alpha_g) \right. \\ & \quad \left. - g_t(\bar{A}_t, \bar{L}_{t-1}; \alpha_g)\right)\frac{\partial}{\partial \alpha_g} g_t(\bar{A}_t, \bar{L}_{t-1}; \alpha_g) \tag{5.6} \end{aligned}$$

to get $\hat{\alpha}_g$.

To use (5.3) for estimation of β , we substitute estimators of the parameters α_w and α_g in the weights and in the g_t 's and set

$$\begin{aligned} & \mathbb{P}_n W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}; \hat{\alpha}_w)\dot{\mu}(\beta, Z)(Y - \mu(\beta, Z)) \\ & \quad - \mathbb{P}_n \sum_{t=1}^K W_{\bar{p}_t}(\bar{A}_t, \bar{L}_{t-1}; \hat{\alpha}_w)\dot{\mu}(\beta, Z)(g_t(\bar{A}_t, \bar{L}_{t-1}; \hat{\alpha}_g) \\ & \quad - \mu(\beta, Z)) + \mathbb{P}_n \sum_{t=1}^K \sum_{a_t} \pi_t(a_t|\bar{A}_{t-1}, \bar{L}_{t-1}; \hat{\alpha}_w) \\ & \quad \times W_{\bar{p}_t}(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}; \hat{\alpha}_w)\dot{\mu}(\beta, Z) \\ & \quad \times (g_t(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}; \hat{\alpha}_g) - \mu(\beta, Z)), \tag{5.7} \end{aligned}$$

to 0 and solve for β .

Under regularity conditions, the following two properties hold.

First, assuming the model for μ is correct, and if the model for $\{g_t(\bar{a}_t, \bar{l}_{t-1}), t = 1, \dots, K\}$, $R \bar{\pi}_K$ is correct, the resulting estimator of β is consistent and asymptotically normal.

Second, the asymptotic variance of $\hat{\beta}$ can be estimated by using a sandwich estimator. Define

$$\hat{\Sigma} = \mathbb{P}_n(S_\beta(X; \hat{\alpha}_w, \hat{\alpha}_g, \hat{\beta}) + \widehat{S}_{\beta,w} \widehat{I}_w^{-1} S_w(X; \hat{\alpha}_w) + \widehat{S}_{\beta,g} \widehat{I}_g^{-1} S_g(X; \hat{\alpha}_g))^{\otimes 2},$$

where for vector V , $V^{\otimes 2} = VV^T$, $S_\beta(X; \hat{\alpha}_w, \hat{\alpha}_g, \beta)$ is the estimating function for β [formula (5.7) is $\mathbb{P}_n S_\beta$], $S_w(X; \alpha_w)$ is the score function for α_w [the derivative of (5.5) with respect to α_w is $\mathbb{P}_n S_w(X; \alpha_w)$], \widehat{I}_w is the observed information matrix for α_w [minus the Hessian of (5.5)], $S_g(X; \alpha_g)$ is the estimating function for α_g [formula (5.6) is $\mathbb{P}_n S_g$], \widehat{I}_g is minus the derivative of (5.6) with respect to α_g , $\widehat{S}_{\beta,w} = \mathbb{P}_n \frac{\partial}{\partial \alpha_w} S_\beta(X; \hat{\alpha}_w, \hat{\alpha}_g, \hat{\beta})$, and $\widehat{S}_{\beta,g} = \mathbb{P}_n \frac{\partial}{\partial \alpha_g} S_\beta(X; \hat{\alpha}_w, \hat{\alpha}_g, \hat{\beta})$. Also define $\widehat{B} = \mathbb{P}_n \dot{\mu}(\beta, Z) \dot{\mu}(\beta, Z)^T$. Then, a consistent estimator of the asymptotic variance of $\hat{\beta}$ is given by $\widehat{B}^{-1} \widehat{\Sigma} \widehat{B}^{-1}$.

See the appendix for proofs of these results.

5.1 Consistent Models for g_1, \dots, g_K

The relationships among the g_1, \dots, g_K are highly constrained. Thus, it is difficult to formulate models for g_1, \dots, g_K that are consistent with one another. For example, models for g_K and g_{K-1} are constrained by the following relationships:

$$g_{K-1}(\bar{a}_{K-1}, \bar{l}_{K-2}) = \int_{l_{K-1}} \sum_{a_K} p_K(a_K | s_{K-1}) g_K(\bar{a}_K, \bar{l}_{K-1}) \times f_{K-1}(l_{K-1} | \bar{a}_{K-1}, \bar{l}_{K-2}) dl_{K-1},$$

where p_K is specified by the treatment regime and f_{K-1} is the unspecified conditional density of L_{K-1} given $(\bar{A}_{K-1}, \bar{L}_{K-2})$. Or, equivalently, given models for g_K and g_{K-1} , there must exist a conditional density f_{K-1} for which the preceding display is true. One way to formulate g_1, \dots, g_K to satisfy the interrelationships is to parameterize each g_t as linear in \bar{l}_{t-1} and independent of \bar{a}_t . This particular parameterization [combine the display and (5.8)] is consistent with no treatment effect. If there is a treatment effect, modeling the g_t 's in this way is used to improve efficiency. We use this parameterization in our example. See Robins and Wasserman (1997) and Scharfstein, Rotnitzky, and Robins (1999) for discussions about consistent parameterization.

5.2 Double Robustness

Instead of using (5.7) to estimate β , we could use (4.5) with estimated weights. Consistency of the resulting $\hat{\beta}$ would require that $\bar{\pi}_K$ is correctly parameterized. Alternately, we could base estimation of β on the formula in (4.4). Note that (4.4) can be written in terms of the g_t 's as

$$\mu(\beta, Z) = \sum_{a_1} E[g_1(a_1, L_0) p_1(a_1 | S_0) | Z]. \quad (5.8)$$

As noted previously, this is the G-computation formula of Robins (1986). Thus, to directly estimate β by using (4.4), we

may estimate the g_t 's to arrive at an estimate of g_1 and then solve

$$0 = \mathbb{P}_n \dot{\mu}(\beta, Z) \left(\sum_{a_1} g_1(a_1, L_0; \hat{\alpha}_g) p_1(a_1 | S_0) - \mu(\beta, Z) \right). \quad (5.9)$$

We can expect the resulting $\hat{\beta}$ to be a consistent estimator of β only if we correctly parameterized the g_t 's.

An advantage of the use of (4.5) over (5.9) is that it is easier to parameterize the $\bar{\pi}_K$ than it is to parameterize the \bar{g}_K , because of the constrained relationships among the \bar{g}_K . An advantage of (5.7) over (5.9) and over the use of (4.5) is that the use of (5.7) to estimate μ leads to a consistent estimator if either the g_t 's or the π_t 's are parameterized correctly [see the first property following (5.7)]. This is the double robustness property; see Scharfstein et al. (1999) and Robins (to appear) for further discussion. Furthermore, note that the first term in (5.7) is (4.5) with estimated weights, and the $t = 1$ summand in the last term is (5.9). Thus, (5.7) combines both the direct estimation method and the weighted estimation method (4.5) to achieve consistency even when only the g_t 's or only the π_t 's are parameterized correctly.

6. EXAMPLE

The Fast Track trial is a longitudinal, multisite, multicohort trial of a preventive intervention versus a control. For the purposes of this illustration, we use a subset of the children for whom data were available at the time of the analysis; this is 579 high-risk children, of which 202 are in the intervention group. This is not a representative subset; furthermore, the variables used in the analysis were selected for illustrative purposes only. A more complete analysis is forthcoming. We consider two endpoints, specifically chosen to illustrate the diversity of results that can follow from the proposed analyses. Both endpoints are from end of grade 3 teacher evaluations. The first endpoint is a teacher rating of school behavior problems. The second endpoint is a teacher rating of behavior change over the course of grade 3. This second rating assesses improvement across the year in social and emotional adjustment.

Beginning in the fall of grade 2, a nonrandom dynamic treatment regime for home visiting was planned as part of the intervention. The home visits were designed to improve parental functioning. The time intervals are fall and spring of grades 2 and 3 ($K = 4$). The severity for semester t is S_t , the total score on the home visiting process measure; S_t is a measure, taking values in the interval $[1, 24]$, evaluating the quality of family functioning. High values indicate better functioning. The treatment rule is, "If assigned to control, there is no assignment of home visiting by staff; if assigned to treatment, the staff should follow the treatment rule

$$d_t(S_{t-1}) = 16I\{S_{t-1} \leq 8\} + 8I\{9 \leq S_{t-1} \leq 16\} + 4I\{17 \leq S_{t-1}\}, \quad t = 1, 2, 3, 4,$$

in assigning the home visiting level." Thus, if a child is in the intervention group and the total on the home visiting process measure from the end of grade 1, S_0 , is less than 8, staff were

to assign 16 home visits for the fall of grade 2, and so on. This means that $A_t \in \{4, 8, 16\}$.

Our primary goal is an evaluation of the average effect of the nonrandom dynamic treatment regime on the two endpoints. This regime corresponds to degenerate treatment selection probabilities, that is,

$$p_t(a|s_{t-1}) = 1 \quad \text{if } a = d_t(s_{t-1}) \quad \text{and 0 otherwise}$$

for $t = 1, 2, 3, 4$.

Staff were told that in exceptional cases they might need to deviate from the rule for recommending home visiting level; in practice, staff recommendations deviated from the rule on 22% of the 202×4 occasions. The staff recommended more home visits than the rule on 15% of the occasions, and on 7% of the occasions the staff recommended fewer home visits. In all, 104 intervention children (52%) had treatment recommendation patterns that coincide with the rule for all four time periods. The staff did not or could not deviate from the rule for the control children because no home visiting assignments were made for the control children. Because the staff deviated from the rule (for the intervention children), we treat the group of intervention children as an observational sample in which the home visiting recommendations may have been influenced by severity (S_{t-1}) in ways varying from the rule, and, indeed, the recommendations may have been influenced by measures other than severity. The Fast Track research team collected a vast amount of information relating to why staff recommended intervention levels. For the purposes of this example, we use past recommended intervention level, the home visiting process measure, and site and cohort indicators to model the distribution of home visiting recommendations; these are included in the L 's. Also included in the L 's are time-varying predictors of the two endpoints. (Recall that these variables are chosen for illustrative purposes; they do not encompass the full set of variables that might be used in a substantive treatment.) See Table 1 for a description of L .

We assume that sequential randomization holds; that is, we assume that within any cross stratification of $\{L_0, A_1, L_1, A_2, \dots, L_{t-1}\}$, the staff treatment recommendation, A_t , is not predictive of the array of potential outcomes, O_{sr} . For example, this means that (within a cross stratification) children for

whom 16 home visits are recommended do not have a higher chance of high values of $\{Y(a_1, a_2, a_3, a_4), a_t = 4, 8, 16, t = 1, 2, 3, 4\}$ than children for whom 4 visits are recommended. It also means that (within a cross stratification) children recommended for 16 home visits do not have a higher chance of high values of $Y(16, 16, 16, 16) - Y(4, 4, 4, 4)$ than children recommended for 4 home visits.

To implement (5.7), we first estimate the weights $W_{\hat{p}_4}$. Because the staff were not to assign and did not assign any intervention to control children, the weights for the control children are equal to one. Furthermore, we do not need to estimate the treatment selection probabilities, $\bar{\pi}_K$, for the control children because we know that each probability has mass one at 0 home visiting. Consider the intervention group children. The numerator of $W_{\hat{p}_4}$ is 1 for an intervention group child whose treatment recommendations coincide with the treatment rule and 0 otherwise. This means that of the 202 intervention children, 98 ($202 - 104$) will have $W_{\hat{p}_4} = 0$, and thus the first term in (5.7) will be 0 for these children. This makes sense because without further (smoothness) assumptions, the responses from the 98 children whose treatment recommendations are inconsistent with the regime are not informative of responses to the treatment regime. This issue is discussed in further detail later. To estimate the denominator of the weight, we fit a proportional odds model to the staff home visiting recommendations for the intervention group children. For each time t , we fit, via maximum likelihood,

$$\begin{aligned} \text{logit}(\pi_t(16|\bar{L}_{t-1}, \bar{A}_{t-1}; \alpha_w)) &= \alpha_{w,t01} + \alpha'_{w,1} ZZ(\bar{L}_{t-1}, \bar{A}_{t-1}), \\ \text{logit}(1 - \pi_t(4|\bar{L}_{t-1}, \bar{A}_{t-1}; \alpha_w)) &= \alpha_{w,t02} + \alpha'_{w,1} ZZ(\bar{L}_{t-1}, \bar{A}_{t-1}), \end{aligned}$$

where $\alpha_w = \{\alpha_{w,t01}, \alpha_{w,t02}, \alpha_{w,1}, t = 1, 2, 3, 4\}$ and $ZZ(\bar{L}_{t-1}, \bar{A}_{t-1})$ is a vector summary of past information. See Tables 2, 3, 4, and 5 for the estimates of α_w . Note that in all four time periods, a low home visiting process measure score is predictive of higher staff home visiting recommendation; also, in the last three time periods, a past high staff home visiting recommendation is predictive of a higher staff home visiting recommendations. Additionally, there are some minor differences in home visiting recommendations by cohort and site. Because there were only four intervals and the sample size was not low, we choose to fit separate models for the home visiting recommendations in each interval; alternately, we could have fit

Table 1. Variables Used in Example

Variable	Collected during
Suspected abuse ¹	Kindergarten
Race ¹	Kindergarten
School behavior problems ¹	End of grade 1
School behavior problems ²	End of grade 3
School behavior change ²	End of grade 3
Home visiting process measure ³	End of every semester
Home visiting assignment ⁴	Beg. of every semester
Academic performance ¹	End of grades 1 and 2
Social skills ¹	Avg. over grades 1 and 2
Site (4 sites) ¹	Kindergarten
Cohort (2 cohorts) ¹	Kindergarten

¹Auxiliary covariates (V_0 or V_T).

²Response (Y).

³Severity (S_t).

⁴Treatment level (A_t).

Table 2. Model for Home Visiting Assignment Probabilities, Fall Grade 2

Covariate	Estimate	Stderr
Intercept 1	-1.45	.36
Intercept 2	1.40	.35
Cohort 1	-.58	.29
Site 1	-.58	.42
Site 2	-.24	.43
Site 3	.42	.36
HVPM16*	2.22	.75
HVPM4*	-2.58	.31

* HVPM16 is 1 only if the home visiting process measure is ≤ 8 ; otherwise it is 0. HVPM4 is 1 only if the home visiting process measure is ≥ 17 ; otherwise it is 0.

one model with possibly time-varying intercepts and regression coefficients to all four intervals simultaneously (Hernán et al. 1998a).

To implement (5.7), we must also estimate α_g . We estimate α_g separately for control and intervention groups. However, because the last two terms in (5.7) cancel for control group children (because the weights are equal to one and the π_t have mass one for 0 home visits), we do not use the estimator of α_g for the control group. Thus, we only need to estimate α_g for the intervention group children. We model g_1 through g_4 by linear models,

$$g_t(\bar{A}_t, \bar{L}_{t-1}; \alpha_g) = \alpha_{g0} + \alpha'_{g1} XX(\bar{L}_{t-1}),$$

where $XX(\bar{L}_{t-1})$ is a summary vector of past information, and we estimate $\alpha_g = \{\alpha_{g0}, \alpha_{g1}\}$ by solving (5.6). Note that the g_t 's are assumed to be functionally independent of both the present time t home visiting level and the past assigned home visiting levels (\bar{A}_t). It is easy to check that this assumption plus the linearity of g_t imply that the parameterizations of g_1, \dots, g_K will be consistent with one another. However, as mentioned earlier, this parameterization can be correct only under the null hypothesis of no treatment effect. Thus, the use of this model for \bar{g}_K implies that we assume our model for the $\bar{\pi}_K$ is correct. The estimates of α_g are given in Tables 6 and 7.

For each response, we model the mean response to the treatment regime \bar{p}_4 by

$$\mu(\beta, Z, \text{treatment}) = \beta_0 + \beta_1 SBP1 + \beta_2 \text{treatment},$$

where treatment is an indicator with the value one for the intervention children, and when grade 3 teacher rating of school behavior problems is the response, grade 1 teacher rating of school behavior problems $SBP1$ is included in the regression. When the teacher rating of behavior change over the course of grade 3 is the response, $SBP1$ is not included. We estimate β by setting (5.7) to zero and solving for $\hat{\beta}$. See Table 4, $h = 1$, for estimates of β_2 . As a comparison, we include the results of an intent to treat analysis (ITT); this is a simple linear regression of the response on treatment (and possibly pretreatment

Table 4. Model for Home Visiting Assignment Probabilities, Fall Grade 3

Covariate	Estimate	Stderr
Intercept 1	-3.19	.69
Intercept 2	3.00	.60
Cohort 1	-.58	.38
Site 1	-1.26	.55
Site 2	-.75	.62
Site 3	1.23	.48
HVPM16 ¹	.19	1.08
HVPM4 ¹	-3.65	.68
Past high home visiting assignment ²	2.98	.68
Past low home visiting assignment ²	-1.31	.41

¹ HVPM16 is 1 only if the home visiting process measure is ≤ 8 ; otherwise it is 0. HVPM4 is 1 only if the home visiting process measure is ≥ 17 ; otherwise it is 0.

² Past high home visiting assignment is 1 if home visiting assignment in the spring of grade 2 is 16; otherwise it is 0. Past low home visiting assignment is 1 if home visiting assignment in spring grade 2 is 4; otherwise it is 0.

covariates). The ITT analysis does not adjust for the level of staff deviance from the rules of the treatment regime. In contrast, the analysis of the dynamic treatment regime yields an estimate of the mean treatment effect in the setting in which staff follow the treatment regime rules in assigning home visiting. In Table 8, we see that for the response school behavior problems, both the ITT treatment effect and the treatment effect for the treatment regime ($h = 1$) are highly significant. But the estimated treatment effect in the scenario in which the staff follow the rules is nearly twice as large as the ITT treatment effect. Recall that (see Tables 2–5) in addition to the home visiting process measure, past treatment recommendation, site, and cohort were predictive of the treatment recommendation probabilities ($\bar{\pi}_4$) in the observational data. The dynamic regime \bar{d}_4 forces uniformity across past treatment recommendation, site, and cohort in making present treatment recommendations. It is unclear whether the increased treatment effect is due to this forced uniformity or is due to using the home visiting measure only as prescribed by the regime rule. This is discussed further later. In contrast, in Table 9, we see that although there is a significant ITT treatment effect on

Table 3. Model for Home Visiting Assignment Probabilities, Spring Grade 2

Covariate	Estimate	Stderr
Intercept 1	-2.61	.47
Intercept 2	1.10	.40
Cohort 1	-.51	.33
Site 1	1.64	.49
Site 2	.01	.56
Site 3	1.76	.45
HVPM16 ¹	3.04	.96
HVPM4 ¹	-2.36	.38
Past high home visiting assignment ²	2.18	.60
Past low home visiting assignment ²	-1.55	.39

¹ HVPM16 is 1 only if the home visiting process measure is ≤ 8 ; otherwise it is 0. HVPM4 is 1 only if the home visiting process measure is ≥ 17 ; otherwise it is 0.

² Past high home visiting assignment is 1 if home visiting assignment in the fall of grade 2 is 16; otherwise it is 0. Past low home visiting assignment is 1 if home visiting assignment in fall of grade 2 is 4; otherwise it is 0.

Table 5. Model for Visiting Assignment Probabilities, Spring Grade 3

Covariate	Estimate	Stderr
Intercept 1	-3.95	1.09
Intercept 2	3.35	.61
Cohort 1	-.89	.43
Site 1	-.82	.55
Site 2	-.56	.63
Site 3	1.18	.54
HVPM16 ¹	4.67	1.61
HVPM4 ¹	-2.75	.47
Past high home visiting assignment ²	5.94	1.61
Past low home visiting assignment ²	-2.53	.43

¹ HVPM16 is 1 only if the home visiting process measure is ≤ 8 ; otherwise it is 0. HVPM4 is 1 only if the home visiting process measure is ≥ 17 ; otherwise it is 0.

² Past high home visiting assignment is 1 if home visiting assignment in fall of grade 3 is 16; otherwise it is 0. Past low home visiting assignment is 1 if home visiting assignment in fall of grade 3 is 4; otherwise it is 0.

Table 6. Model for g_1 Through g_4 , School Behavior Problems

Covariate	Estimate*	Stderr
Intercept	.30	.21
Cohort 1	-.28	.10
Site 1	.43	.14
Site 2	.48	.15
Site 3	.45	.12
Suspected abuse	.20	.06
School behavior problems—grade 1	.26	.06
Most recent academic performance score	.09	.03
Most recent social skills score	.16	.05

* All variables are coded so that high values indicate poorer functioning.

behavior change, there is no treatment effect for the treatment regime ($h = 1$). In this case, it appears that staff judgment in deviating from the rules has enhanced the treatment effect. It is unclear whether the decreased treatment effect is due to forced uniformity across site and cohort in assigning treatment, or is due to ignoring the past level of assigned treatment in making future recommendations, or is due to using the home visiting measure only as prescribed by the regime rule.

Recall that 98 of the 202 intervention children have $W_{\bar{p}_4} = 0$. These children's responses enter (5.7) only through the estimated g_t 's. To better utilize these responses and to obtain a better understanding of the reasons for the difference in treatment effects between the regime and the ITT analyses, we examine a small variety of random dynamic treatment regimes, denoted by $\bar{p}_{4,h}$, $h \in \{0, .2, .4, \dots, 1\}$. $p_{t,h}(a|s_{t-1})$ will be the probability of recommending home visiting level a at time t when the home visiting measure is equal to s_{t-1} . For each $t \in \{1, 2, 3, 4\}$, set $p_{t,h}$ equal to a mixture of the indicator, $I\{a = d_t(s_{t-1})\}$ and $q_t(a|s_{t-1})$,

$$p_{t,h}(a|s_{t-1}; \gamma) = h I\{a = d_t(s_{t-1})\} + (1 - h) q_t(a|s_{t-1}; \gamma),$$

where q_t is a conditional probability mass function for treatment selection given severity, s_{t-1} . We choose q_t to approximate the observational distribution of treatment recommendation at time t given only the home visiting process measure. All these random dynamic regimes do not allow site, cohort, and past home visiting recommendations to influence

Table 7. Model for g_1 Through g_4 , School Behavior Change

Covariate	Estimate*	Stderr
Intercept	-.93	.34
Site 1	-.62	.19
Site 2	-.44	.18
Site 3	-.15	.19
Suspected abuse	-.06	.10
School behavior problems—grade 1	.18	.10
Most recent academic performance score	.05	.03
Most recent social skills score	-.15	.07

* All variables are coded so that high values indicate poorer functioning.

Table 8. Estimated Treatment Effect on School Behavior Problems¹

h	Estimated treatment effect ²	Stderr
ITT ³	-.23	.07
0	-.21	.08
.2	-.24	.08
.4	-.28	.08
.6	-.31	.08
.8	-.35	.08
1	-.40	.09

¹ Estimated intercepts and regression coefficients of grade 1 school behavior problems score are omitted.

² High values on the school behavior problems response indicate behavior problems at school.

³ Results of a simple linear regression of response on treatment and grade 1 school behavior problems; all weights are equal to 1.

present home visiting recommendation. Note that $h = 1$ corresponds to originally planned treatment regime. The random dynamic treatment regime, $\{p_{1,h}, p_{2,h}, p_{3,h}, p_{4,h}\}$, corresponds to drawing a Bernoulli random variable with success probability equal to h at each time t ; if the Bernoulli outcome is 1, then set A_t equal to $d_t(S_{t-1})$; otherwise, draw A_t from the probability mass function, $q_t(a|S_{t-1}; \gamma)$.

Of course, the observational distribution of treatment recommendations given only the home visiting process measure is unknown and must be approximated. We estimate this approximation. To construct q_t , $t = 1, 2, 3, 4$, at each time we fit, via maximum likelihood, the following proportional odds model to the data:

$$\begin{aligned} \text{logit}(q_t(16|S_{t-1}; \gamma)) &= \gamma_{t01} + \gamma_{t1} I\{d(S_{t-1}) = 16\} \\ &+ \gamma_{t2} I\{d(S_{t-1}) = 4\}, \end{aligned}$$

$$\begin{aligned} \text{logit}(1 - q_t(4|S_{t-1}; \gamma)) &= \gamma_{t02} + \gamma_{t1} I\{d(S_{t-1}) = 16\} \\ &+ \gamma_{t2} I\{d(S_{t-1}) = 4\}, \end{aligned}$$

where $\gamma = \{\gamma_{t01}, \gamma_{t02}, \gamma_{t1}, \gamma_{t2}, t = 1, 2, 3, 4\}$. The estimated values of γ are given in Tables 10 and 11. We use $q_t(\cdot|S_{t-1}; \hat{\gamma})$ to form the $p_{t,h}(\cdot|S_{t-1}; \hat{\gamma})$ when $h < 1$.

Note that because the proportional odds model q_t is not a saturated model, the models for q_t and π_t will not be consistent with one another. However, we view q_t as only an approximation to the observational distribution of treatment recommendation given severity. Also note that $p_{t,h}(\cdot|S_{t-1}; \hat{\gamma})$

Table 9. Estimated Treatment Effect on School Behavior Change¹

h	Estimated treatment effect ²	Stderr
ITT ³	-.33	.09
0	-.39	.10
.2	-.38	.10
.4	-.36	.10
.6	-.32	.10
.8	-.27	.11
1	-.19	.13

¹ The estimated intercepts are omitted.

² Low values on the school behavior change indicate improvement across the year in social and emotional adjustment.

³ Results of a simple linear regression of response on treatment; all weights are equal to 1.

Table 10. Model for Numerator of the Weights, Grade 2

Covariate	Fall		Spring	
	Estimate	Stderr ¹	Estimate	Stderr ¹
Intercept 1	-1.69	.23	-1.55	.25
Intercept 2	1.06	.20	1.27	.24
HVPM16 ²	2.43	1.03	2.67	.82
HVPM4 ²	-2.61	.34	-2.69	.34

¹This model is used to approximate the distribution of treatment assignment given the home visiting process measure; thus, robust standard errors are provided.

²HVPM16 is 1 only if the home visiting process measure is ≤ 8 ; otherwise it is 0. HVPM4 is 1 only if the home visiting process measure is ≥ 17 ; otherwise it is 0.

are used to form the numerator of the weights, and, as before, $\pi_i(\cdot|\bar{L}_{t-1}, \bar{A}_{t-1}; \hat{\alpha}_w)$ are used to form the denominator of the weights. For each value of h and each of the two responses, we use (5.6) to construct estimators of α_g , but for $h < 1$, p_t is replaced by $p_{t,h}(\cdot|S_{t-1}; \hat{\gamma})$. The estimates of γ are omitted. The model for the mean treatment response is as before; Tables 8 and 9 contain the results. The formula for the estimated standard errors changes because the numerator of the weights is now estimated. See the appendix for the appropriate formula. Consider the results for $h = 0$ in Tables 8 and 9. These results are essentially equivalent to the ITT results. Recall that for $h = 0$, the dynamic treatment regime corresponds to drawing A_t from the probability mass function, $q_t(a|S_{t-1}; \gamma)$. Thus, although there was systematic variation from the treatment assignment rule, systematic variation that can be accounted for by past home visiting recommendation and site and cohort differences, this systematic variation does not appear to alter the mean response. Rather, we see that as the random dynamic treatment regime approaches (h approaches 1) the original dynamic treatment regime, the mean effect changes. This indicates that the change in mean effect from the ITT analysis to the original dynamic treatment regime analysis is due to only using the home visiting measure as prescribed by the regime rule.

Clearly a more complete analysis of the Fast Track data is needed to investigate these substantive issues. As stated previously, the two responses were chosen specifically to illustrate the diversity of thought-provoking results that can be obtained by these types of analysis. Certainly, one should give careful thought to which responses are of primary interest; variations in the dynamic treatment regime may differentially impact different responses.

Table 11. Model for Numerator of the Weights, Grade 3

Covariate	Fall		Spring	
	Estimate	Stderr ¹	Estimate	Stderr ¹
Intercept 1	-2.57	.42	-2.87	.45
Intercept 2	1.71	.32	1.28	.26
HVPM16 ²	1.95	1.22	3.79	.93
HVPM4 ²	-3.00	.36	-2.43	.34

¹This model is used to approximate the distribution of treatment assignment given the home visiting process measure; thus, robust standard errors are provided.

²HVPM16 is 1 only if the home visiting process measure is ≤ 8 ; otherwise it is 0. HVPM4 is 1 only if the home visiting process measure is ≥ 17 ; otherwise it is 0.

7. DISCUSSION OF RESULTS

This methodology provides a way to evaluate the effects of dynamic treatment regimes using observational data. It should be extended in a variety of ways. First, the interpretation of the results depends heavily on the assumption of sequential randomization. It is difficult to believe that this assumption really holds in the observational setting. Thus, an extension to this setting of the sensitivity analyses of Robins, Rotnitzky, and Scharfstein (1999b) and Scharfstein et al. (1999) is needed. A second needed extension is to allow for missing severity measures. There are two ways to do this. First, if one can expect severity measures to be occasionally missing in practice, then it makes sense to extend the dynamic regime rules to include rules for assignment when severity is missing. On the other hand, if one wishes to make inference for a rule that does not allow for missing severity measures, yet one's observational data include individuals with missing severity measures, then the methodology provided here must be adapted. One possibility is a modification of the weighting method as described by Hernán et al. (1998). A third extension would be to grouped data. Many intervention and prevention studies take place in a school-based setting, and thus grouped data dominates. Last, more systematic work is needed on how to formulate and evaluate dynamic treatment rules and in particular the evaluation and detection of which posttreatment conditions (such as the development of side effects or the occurrence of family crises) may negate or enhance the effect of treatment.

APPENDIX

Proof of Lemma 4.1. We repeatedly use the notation, $Y == Y(\bar{A}_K)$, $L_t == (S_t, V_t)$, and $L_t == L_t(\bar{A}_t)$. Let U be a Borel set. We have that

$$\begin{aligned}
& E_{obs}[I\{(Y, \bar{S}_{K-1}, \bar{A}_K, Z) \in U\}] \\
& \quad \times E_{obs}[W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1})|Y, \bar{S}_{K-1}, \bar{A}_K, Z]] \\
& = E_{obs}[I\{(Y, \bar{S}_{K-1}, \bar{A}_K, Z) \in U\}W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1})] \\
& = \sum_{\bar{a}_K} E_{obs}[I\{(Y, \bar{S}_{K-1}, \bar{a}_K, Z) \in U\}] \\
& \quad \times W_{\bar{p}_K}(\bar{a}_K, \bar{L}_{K-1})I\{\bar{A}_K = \bar{a}_K\}] \\
& = \sum_{\bar{a}_K} E_{obs}[I\{(Y(\bar{a}_K), \bar{S}_{K-1}, \bar{a}_K, Z) \in U\}] \\
& \quad \times W_{\bar{p}_{K-1}}(\bar{a}_K, \bar{L}_{K-2})p_K(a_K|S_{K-1}) \\
& \quad \times P_{obs}[\pi_K(a_K|\bar{a}_{K-1}, \bar{L}_{K-1}) > 0|O_{sr}, \bar{L}_{K-2}, \\
& \quad \quad \bar{A}_{K-1} = \bar{a}_{K-1}]I\{\bar{A}_{K-1} = \bar{a}_{K-1}\}].
\end{aligned}$$

Continuing in this fashion, eliminating the indicator variables concerning A_j and accumulating the conditional expectations concerning the $\pi_j(a_j|\bar{L}_{j-1}(\bar{a}_{j-1}), \bar{a}_{j-1})$, we arrive at

$$\begin{aligned}
& \sum_{\bar{a}_K} E_{obs} \left[I\{(Y(\bar{a}_K), \bar{S}_{K-1}(\bar{a}_{K-1}), \bar{a}_K, Z) \in U\} B_1(O_{sr}, \bar{a}_K) \right. \\
& \quad \left. \times \prod_{t=1}^K p_t(a_t|S_{t-1}(\bar{a}_{t-1})) \right], \quad (A.1)
\end{aligned}$$

where B_1 is formed from repeated conditional expectations and can be defined iteratively, beginning with

$$B_K(O_{sr}, \bar{L}_{K-2}, \bar{a}_K) = P_{obs} \left[\prod_{t=1}^K \pi_t(a_t | \bar{a}_{t-1}, \bar{L}_{t-1}(\bar{a}_{t-1})) > 0 | O_{sr}, \bar{L}_{K-2}, \bar{A}_{K-1} = \bar{a}_{K-1} \right],$$

and then for each $j \leq K - 1$,

$$B_j(O_{sr}, \bar{L}_{j-2}, \bar{a}_K) = E_{obs} [B_{j+1}(O_{sr}, \bar{L}_{j-1}, \bar{a}_K) | O_{sr}, \bar{L}_{j-2}, \bar{A}_{j-1} = \bar{a}_{j-1}]$$

and for $j = 1$ as

$$B_1(O_{sr}, \bar{a}_K) = E_{obs} [B_2(O_{sr}, L_0, \bar{a}_K) | O_{sr}].$$

Because the distribution of O_{sr} is the same under both $P_{\bar{p}_K}$ and P_{obs} distributions, (A.1) also can be written with $E_{\bar{p}_K}$ in place of E_{obs} . But then the sum over \bar{a}_K can be taken inside the expectation to yield

$$E_{obs} [I\{(Y, \bar{S}_{K-1}, \bar{a}_K, Z) \in U\} E_{\bar{p}_K} [W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}) | Y, \bar{S}_{K-1}, \bar{A}_K]] = E_{\bar{p}_K} [B_1(O_{sr}, \bar{A}_K) I\{(Y, \bar{S}_{K-1}, \bar{A}_K, Z) \in U\}].$$

Because B_1 can be expressed as repeated expectations of a probability (B_K) that is necessarily bounded above by one, (4.1) implies that $B_1(O_{sr}, \bar{A}_K) = 1$ almost everywhere. Thus, we have absolute continuity, and (4.2) is the Radon–Nikodym derivative of $P_{\bar{p}_K}$ with respect to P_{obs} .

In the following, we assume (5.1). Denote the true value of β by β_0 . It is useful to note two alternate ways to rewrite (5.3). The first emphasizes the correct specification of $\bar{\pi}_K$:

$$S_\beta(X; \alpha_w, \alpha_g, \beta) = W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}) \dot{\mu}(\beta, Z) (Y - \mu(\beta, Z)) - \sum_{t=1}^K \sum_{a_t} (I\{A_t = a_t\} - \pi_t(a_t | \bar{A}_{t-1}, \bar{L}_{t-1})) \times W_{\bar{p}_t}(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}) \dot{\mu}(\beta, Z) \times (g_t(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}) - \mu(\beta, Z)). \quad (A.2)$$

Note that if $\bar{\pi}_K$ is correctly specified, then the second term of (A.2) has mean zero even if the g_t 's are incorrectly specified or $\beta \neq \beta_0$ or both. The second way to rewrite (5.3) emphasizes the correct estimation of the g_t 's:

$$S_\beta(X; \alpha_w, \alpha_g, \beta) = W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}) \dot{\mu}(\beta, Z) (Y - g_K(\bar{A}_K, \bar{L}_{K-1})) - \sum_{t=1}^{K-1} W_{\bar{p}_t}(\bar{A}_t, \bar{L}_{t-1}) \dot{\mu}(\beta, Z) \left(g_t(\bar{A}_t, \bar{L}_{t-1}) - \sum_{a_{t+1}} p_{t+1}(a_{t+1} | S_t) g_{t+1}(a_{t+1}, \bar{A}_t, \bar{L}_t) \right) + \dot{\mu}(\beta, Z) \left(\sum_{a_1} g_1(a_1, L_0) p_1(a_1 | S_0) - \mu(\beta, Z) \right). \quad (A.3)$$

Note that if the g_t 's are correctly specified, then the first two terms have expectation zero even if $\bar{\pi}_K$ is incorrectly specified or $\beta \neq \beta_0$ or both.

Lemma A.1. Formula (5.3) is orthogonal to, or equivalently, uncorrelated with the score functions for the treatment assignment probabilities, $\bar{\pi}_K$; thus, (5.3) is $W_{\bar{p}_K}(\bar{A}_K, \bar{L}_K) \dot{\mu}(\beta, Z) (Y - \mu(\beta, Z))$ minus its projection on the score functions for the treatment assignment probabilities.

Proof. First note that all score functions for the $\bar{\pi}_K$ belong to the linear span of

$$B_t(a, \bar{A}_t, \bar{L}_{t-1}) = U_t(\bar{A}_{t-1}, \bar{L}_{t-1}) (I\{A_t = a\} - \pi_t(a | \bar{A}_{t-1}, \bar{L}_{t-1}))$$

for each possible treatment level, a and $t \in \{1, \dots, K\}$ and for $U_t(\bar{A}_{t-1}, \bar{L}_{t-1})$ for $t \in \{1, \dots, K\}$ arbitrary bounded functions. Thus, we need only show that (5.3) or equivalently (A.2) is orthogonal to each of these functions. Note that $B_j(a, \bar{A}_j, \bar{L}_{j-1})$ is orthogonal to the terms with $t \neq j$ in the sum in (A.2). The covariance between $B_j(a, \bar{A}_j, \bar{L}_{j-1})$ and (A.2) is thus

$$E_{obs} [W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}) \dot{\mu}(\beta, Z) (Y - \mu(\beta, Z)) B_j(a, \bar{A}_j, \bar{L}_{j-1})] - E_{obs} \left(\sum_{a_j} (I\{A_j = a_j\} - \pi_j(a_j | \bar{A}_{j-1}, \bar{L}_{j-1})) \times W_{\bar{p}_j}(a_j, \bar{A}_{j-1}, \bar{L}_{j-1}) \dot{\mu}(\beta, Z) (g_j(a_j, \bar{A}_{j-1}, \bar{L}_{j-1}) - \mu(\beta, Z)) \right) B_j(a, \bar{A}_j, \bar{L}_{j-1}). \quad (A.4)$$

Because, given $\bar{A}_{j-1}, \bar{L}_{j-1}$, $B_j(a, \bar{A}_j, \bar{L}_{j-1})$ has mean zero, the second term (A.5) reduces to

$$- E_{obs} \sum_{a_j} I\{A_j = a_j\} W_{\bar{p}_j}(a_j, \bar{A}_{j-1}, \bar{L}_{j-1}) \dot{\mu}(\beta, Z) \times (g_j(a_j, \bar{A}_{j-1}, \bar{L}_{j-1}) - \mu(\beta, Z)) B_j(a, \bar{A}_j, \bar{L}_{j-1}).$$

But this is the negative of the first term (A.4) [repeatedly take conditional expectations and use (5.1)].

Lemma A.2. If the parameterization of $\mu(\beta, Z)$ is saturated and g_1, \dots, g_K are nonparametric, then (5.3) is the efficient influence function for β .

Proof. Because of the factorization (A.4) of the likelihood and the fact that the functional $\mu(\beta, Z)$ does not depend on the factor $\prod_{t=1}^K \pi_t(a_t | \bar{a}_{t-1}, \bar{l}_{t-1})$, the efficient influence function is the same in the model for which $\bar{\pi}_K$ is completely unknown as in the model where $\bar{\pi}_K$ is a known function up to a parameter α_w (Robins et al. 1994; Robins and Ritov 1997). However, when $\mu(\beta, Z)$ is saturated, the model with $\bar{\pi}_K$ completely unknown is a nonparametric model for the observed data. Bickel et al. (1993) showed that there is only one influence function for any nonparametric model, which therefore must necessarily be the efficient influence function. Thus, if we find an influence function of the model in which $\bar{\pi}_K$ is completely unknown, this will be the efficient influence function for this nonparametric model and also will be the efficient influence function for the model in which $\bar{\pi}_K$ is parameterized. Now, when $\bar{\pi}_K$ is completely known, $U(X, \beta, \bar{p}_K) = W_{\bar{p}_K}(\bar{A}_K, \bar{L}_K) \dot{\mu}(\beta, Z) (Y - \mu(\beta, Z))$ is an unbiased estimating function for β regardless of the distribution of $Y - \mu$. It follows that $U(X, \beta)$ minus its projection on all scores for $\bar{\pi}_K$ [as given in (5.3)] is orthogonal to the nuisance tangent space for the nonparametric models. Furthermore, (5.3) is an influence function (the expectation of its derivative with respect to β is -1). Thus, (5.3) is the efficient influence function when $\bar{\pi}_K$ is parameterized.

Assume the following:

1. There exists a finite vector α_w^* such that $P_{obs}(S_w(X; \alpha_w^*)) = 0$, where $P_n S_w(X; \alpha_w)$ is the derivative with respect to α_w of (5.5) and there exists a finite vector, α_g^* , such that $P_{obs}(S_g(X; \alpha_g^*)) = 0$, where $P_n S_g(X; \alpha_g)$ is given in (5.6). $\mu(\beta_0, Z)$ is equal to (4.4); that is, the parameterization of (4.4) in terms of β is correct. If $\bar{\pi}_K$ is correctly

parameterized, then put $\alpha_w^* = \alpha_{w0}$, the true parameter in the observational distribution. Similarly, if $\{g_1, \dots, g_K\}$ are correctly parameterized, then $\alpha_g^* = \alpha_{g0}$, the true parameter in the observational distribution.

2. Let Θ be neighborhood of $(\alpha_w^*, \alpha_g^*, \beta_0)$. The class of functions

$$\left\{ \begin{aligned} &\frac{\partial}{\partial \alpha_g} S_\beta(X; \alpha_w, \alpha_g, \beta), \frac{\partial}{\partial \alpha_w} S_\beta(X; \alpha_w, \alpha_g, \beta), \\ &\frac{\partial}{\partial \beta} S_\beta(X; \alpha_w, \alpha_g, \beta), \frac{\partial}{\partial \alpha_w} S_w(X; \alpha_w) \frac{\partial}{\partial \alpha_g} S_g(X; \alpha_g), \\ &S_w(X; \alpha_w)^{\otimes 2}, S_g(X; \alpha_g)^{\otimes 2}, S_\beta(X; \gamma, \alpha_w, \alpha_g, \beta)^{\otimes 2} \\ &\text{for } (\alpha_w, \alpha_g, \beta) \in \Theta \end{aligned} \right\}$$

is a Glivenko–Cantelli class. (For vector V , $V^{\otimes 2} = VV^T$.)

3. Assume that $I_w = E_{obs} \frac{\partial}{\partial \alpha_w} S_w(X; \alpha_w)$ is invertible at $\alpha_w = \alpha_w^*$, $I_g = E_{obs} \frac{\partial}{\partial \alpha_g} S_g(X; \alpha_g)$ is invertible at $\alpha_g = \alpha_g^*$, and $B = E_{obs} \dot{\mu}(\beta_0, Z) \dot{\mu}(\beta_0, Z)^T$ is invertible.

Lemma A.3. Suppose that either $\bar{\pi}_K$ or the g_t 's are correctly parameterized. Assuming points 1–3, there exists a sequence of $\hat{\beta}$ solutions to (5.7) for which $\sqrt{n}(\hat{\beta} - \beta_0)$ converges in distribution to a multivariate normal with mean zero and var-covariance matrix, $B^{-1} \Sigma B^{-1}$, where

$$\begin{aligned} \Sigma = &E_{obs} (S_\beta(X; \alpha_w^*, \alpha_g^*, \beta_0) + S_{\beta, \alpha_w} I_w^{-1} S_w(X; \alpha_w^*) \\ &+ S_{\beta, \alpha_g} I_g^{-1} S_g(X; \alpha_g^*))^{\otimes 2}, \end{aligned}$$

where $S_{\beta, \alpha_w} = E_{obs} \frac{\partial}{\partial \alpha_w} S_\beta(X; \alpha_w, \alpha_g^*, \beta_0)|_{\alpha_w = \alpha_w^*}$ and $S_{\beta, \alpha_g} = E_{obs} \times \frac{\partial}{\partial \alpha_g} S_\beta(X; \alpha_w^*, \alpha_g, \beta_0)|_{\alpha_g = \alpha_g^*}$. Furthermore, $\hat{B}^{-1} \hat{\Sigma} \hat{B}^{-1}$ is a consistent estimator of $B^{-1} \Sigma B^{-1}$.

Proof. Suppose that $\alpha_w^* = \alpha_{w0}$. Then, as noted in the discussion concerning (A.2), we have that

$$\begin{aligned} E_{obs} S_\beta(X, \alpha_{w0}, \alpha_g^*, \beta_0) \\ = E_{obs} W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}; \alpha_{w0}) \dot{\mu}(\beta_0, Z) (Y - \mu(\beta_0, Z)). \end{aligned}$$

This is zero because $\mu(\beta_0, Z)$ is equal to (4.4). Alternately, suppose that $\alpha_g^* = \alpha_{g0}$. Then, as noted in the discussion concerning (A.3), we have that

$$\begin{aligned} E_{obs} S_\beta(X, \alpha_w^*, \alpha_{g0}, \beta_0) \\ = E_{obs} \dot{\mu}(\beta_0, Z) \left(\sum_{a_1} g_1(a_1, L_0; \alpha_{g0}) p_1(a_1 | S_0) - \mu(\beta_0, Z) \right). \end{aligned}$$

As before, this is zero because $\mu(\beta_0, Z)$ is equal to (4.4) [recall that (4.4) can be written as in (5.8)].

Proofs that there exists a consistent sequence of solutions, $\hat{\alpha}_w \xrightarrow{P} \alpha_w^*$ to $\mathbb{P}_n(S_w(X; \alpha_w)) = 0$, and a consistent sequence of solutions, $\hat{\alpha}_g \xrightarrow{P} \alpha_g^*$ to $\mathbb{P}_n(S_g(X; \alpha_g)) = 0$, are very similar and use assumptions 1, 2, and 3 along with lemma 2 of Aitchison and Silvey (1958). Next, using existence of consistent $\hat{\alpha}_w$ and $\hat{\alpha}_g$ and again using Aitchison and Silvey's lemma 2 along with assumptions 1, 2, 3, and 5, one can show there exist a consistency solution, $\hat{\beta} \xrightarrow{P} \beta_0$. These proofs are omitted.

To prove asymptotic normality, use Taylor series arguments, along with the Glivenko–Cantelli property in 2 and the invertibility in 3, to deduce that

$$\sqrt{n}(\hat{\alpha}_w - \alpha_w^*) = I_w^{-1} \sqrt{n} \mathbb{P}_n(S_w(X; \alpha_w^*)) + o_P(1)$$

and that

$$\sqrt{n}(\hat{\alpha}_g - \alpha_g^*) = I_g^{-1} \sqrt{n} \mathbb{P}_n(S_g(X; \alpha_g^*)) + o_P(1).$$

Next, a Taylor series yields

$$\begin{aligned} 0 = &\mathbb{P}_n S_\beta(X, \hat{\alpha}_w, \hat{\alpha}_g, \hat{\beta}) = \mathbb{P}_n S_\beta(X, \alpha_w^*, \alpha_g^*, \beta_0) \\ &+ \mathbb{P}_n \frac{\partial}{\partial \alpha_w} S_\beta(X; \alpha_w, \hat{\alpha}_g, \hat{\beta})|_{\alpha_w = \alpha_w^*} (\hat{\alpha}_w - \alpha_w^*), \\ &\text{(where } \alpha_w^{**} \text{ is a vector of values intermediate} \\ &\text{between } \hat{\alpha}_w \text{ and } \alpha_w^* \text{),} \\ &+ \mathbb{P}_n \frac{\partial}{\partial \alpha_g} S_\beta(X; \alpha_w, \alpha_g, \hat{\beta})|_{\alpha_g = \alpha_g^*} (\hat{\alpha}_g - \alpha_g^*), \\ &\text{(where } \alpha_g^{**} \text{ is a vector of values intermediate} \\ &\text{between } \hat{\alpha}_g \text{ and } \alpha_g^* \text{),} \\ &+ \mathbb{P}_n \frac{\partial}{\partial \beta} S_\beta(X; \alpha_w^*, \alpha_g^*, \beta^{**}) (\hat{\beta} - \beta_0), \\ &\text{(where } \beta^{**} \text{ is a vector of values intermediate} \\ &\text{between } \hat{\beta} \text{ and } \beta_0 \text{),} \end{aligned}$$

Note from the previous discussion that the first term in the expansion has mean zero. Next, apply the Glivenko–Cantelli property in 2, along with the invertibility in 3, to yield

$$\begin{aligned} \sqrt{n}(\hat{\beta} - \beta_0) = &B^{-1} \mathbb{P}_n \left(S_\beta(X; \alpha_w^*, \alpha_g^*, \beta_0) + S_{\beta, \alpha_w} I_w^{-1} S_w(X; \alpha_w^*) \right. \\ &\left. + S_{\beta, \alpha_g} I_g^{-1} S_g(X; \alpha_g^*) \right) + o_P(1). \end{aligned}$$

To show that $\hat{B}^{-1} \hat{\Sigma} \hat{B}^{-1}$ converges in probability to $B^{-1} \Sigma B^{-1}$, we may use the asymptotic normality result along with 2 and 3.

Given a parameterization of \bar{p}_K in terms of the parameter, γ , define

$$S_{num}(X; \gamma) = \frac{\partial}{\partial \gamma} \log p_t(A_t | S_{t-1}; \gamma).$$

Define

$$W_{\bar{p}_t}(\bar{A}_K, \bar{L}_{K-1}; \gamma, \alpha_w) = \prod_{j=1}^t \frac{p_j(A_j | S_{j-1}; \gamma)}{\pi_j(A_j | \bar{L}_{j-1}, \bar{A}_{j-1}; \alpha_w)}$$

and

$$\begin{aligned} S_g(X; \gamma, \alpha_g) = &(Y - g_K(\bar{A}_K, \bar{L}_{K-1}; \alpha_g)) \frac{\partial}{\partial \alpha_g} g_K(\bar{A}_K, \bar{L}_{K-1}; \alpha_g) \\ &+ \sum_{t=1}^{K-1} \mathbb{P}_n \left(\sum_{a_{t+1}} p_{t+1}(a_{t+1} | S_t; \gamma) \right. \\ &\left. \times g_{t+1}(a_{t+1}, \bar{A}_t, \bar{L}_t; \alpha_g) - g_t(\bar{A}_t, \bar{L}_{t-1}; \alpha_g) \right) \\ &\times \frac{\partial}{\partial \alpha_g} g_t(\bar{A}_t, \bar{L}_{t-1}; \alpha_g). \end{aligned}$$

Define

$$\begin{aligned} S_\beta(X; \alpha_w, \alpha_g, \gamma, \beta) \\ = &W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1}; \hat{\alpha}_w) \dot{\mu}(\beta, Z) \\ &\times (Y - \mu(\beta, Z)) - \mathbb{P}_n \sum_{t=1}^K W_{\bar{p}_t}(\bar{A}_t, \bar{L}_{t-1}; \hat{\alpha}_w) \end{aligned}$$

$$\begin{aligned} & \times \dot{\mu}(\beta, Z)(g_t(\bar{A}_t, \bar{L}_{t-1}; \hat{\alpha}_g) - \mu(\beta, Z)) \\ & + \mathbb{P}_n \sum_{t=1}^K \sum_{a_t} \pi_t(a_t | \bar{A}_{t-1}, \bar{L}_{t-1}; \hat{\alpha}_w) \\ & \times W_{\hat{p}_t}(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}; \hat{\alpha}_w) \dot{\mu}(\beta, Z) \\ & \times (g_t(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}; \hat{\alpha}_g) - \mu(\beta, Z)). \end{aligned}$$

Define

$$\begin{aligned} U(X) &= S_{\beta}(X; \gamma^*, \alpha_w^*, \alpha_g^*, \beta_0^*) + S_{\beta, w} I_w^{-1} S_w(X; \alpha_w^*) \\ & + S_{\beta, g} I_g^{-1} (S_g(X; \gamma^*, \alpha_g^*) + S_{g, num} I_{num}^{-1} S_{num}(X; \gamma^*)) \\ & + S_{\beta, num} I_{num}^{-1} S_{num}(X; \gamma^*), \end{aligned}$$

where $S_{\beta, w} = E_{obs} \frac{\partial}{\partial \alpha_w} S_{\beta}(X; \gamma, \alpha_w, \alpha_g^*, \beta_0^*) |_{\alpha_w = \alpha_w^*}$, $S_{\beta, g} = E_{obs} \frac{\partial}{\partial \alpha_g} S_{\beta}(X; \gamma, \alpha_w^*, \alpha_g, \beta_0^*) |_{\alpha_g = \alpha_g^*}$, and $S_{g, num} = E_{obs} \times \frac{\partial}{\partial \gamma} S_g(X; \gamma, \alpha_g^*) |_{\gamma = \gamma^*}$. Similarly, let $\hat{\Sigma}^* = \mathbb{P}_n \hat{U}(X)^{\otimes 2}$, where

$$\begin{aligned} \hat{U}(X) &= S_{\beta}(X; \hat{\gamma}, \hat{\alpha}_w, \hat{\alpha}_g, \hat{\beta}) + \hat{S}_{\beta, w} \hat{I}_w^{-1} S_w(X; \hat{\alpha}_w) \\ & + \hat{S}_{\beta, g} \hat{I}_g^{-1} (S_g(X; \hat{\gamma}, \hat{\alpha}_g) + \hat{S}_{g, num} \hat{I}_{num}^{-1} S_{num}(X; \hat{\gamma}^*)) \\ & + \hat{S}_{\beta, num} \hat{I}_{num}^{-1} S_{num}(X; \hat{\gamma}^*) \end{aligned}$$

and $\hat{S}_{\beta, w} = \mathbb{P}_n \frac{\partial}{\partial \alpha_w} S_{\beta}(X; \hat{\gamma}, \alpha_w, \hat{\alpha}_g, \hat{\beta}) |_{\alpha_w = \hat{\alpha}_w}$, $\hat{S}_{\beta, g} = \mathbb{P}_n \times \frac{\partial}{\partial \alpha_g} S_{\beta}(X; \hat{\gamma}, \hat{\alpha}_w, \alpha_g, \hat{\beta}) |_{\alpha_g = \hat{\alpha}_g}$, and $\hat{S}_{g, num} = \mathbb{P}_n \frac{\partial}{\partial \gamma} S_g(X; \gamma, \hat{\alpha}_g) |_{\gamma = \hat{\gamma}^*}$.

1' There exists a finite vector α_w^* such that $P_{obs}(S_w(X; \alpha_w^*)) = 0$, and there exists a finite vector α_g^* such that $P_{obs}(S_g(X; \alpha_g^*)) = 0$. There exists a finite vector γ^* such that $P_{obs}(S_{num}(X; \gamma^*)) = 0$ and also a finite vector β_0^* for which

$$\begin{aligned} & \mu(\beta_0^*, Z) \\ & = \sum_{\bar{a}_K \in \mathcal{A}} \int_{\bar{l}_{K-1}} \int_y y f_Y(y | \bar{a}_K, \bar{v}_{K-1}, \bar{s}_{K-1}) \prod_{t=1}^K p_t(a_t | s_{t-1}; \gamma^*) \\ & \quad \times \prod_{t=1}^{K-1} f_t(l_t | \bar{a}_t, \bar{l}_{t-1}) f_0(l_0 | z) dy d\bar{l}_{K-1}. \end{aligned}$$

Suppose there exists a finite vector α_g^* such that $P_{obs}(S_g(X; \gamma^*, \alpha_g^*)) = 0$. If $\bar{\pi}_K$ is correctly parameterized, then put $\alpha_w^* = \alpha_{w0}$ the true parameter in the observational distribution. Similarly, if $\{g_1, \dots, g_K\}$ are correctly parameterized, then $\alpha_g^* = \alpha_{g0}$, the true parameter in the observational distribution. The parameter α_w is estimated as before. Calculate $\hat{\gamma}$ by solving $0 = \mathbb{P}_n S_{num}(X; \gamma)$. Calculate $\hat{\alpha}_g$ by solving $0 = \mathbb{P}_n S_g(X; \hat{\gamma}, \alpha_g)$. Last, estimate $\hat{\beta}^*$ by solving $0 = \mathbb{P}_n S_{\beta}(X; \hat{\alpha}_w, \hat{\alpha}_g, \hat{\gamma}, \beta)$ for $\hat{\beta}$.

2' Let Θ be neighborhood of $(\gamma^*, \alpha_w^*, \alpha_g^*, \beta_0^*)$. The class of functions

$$\begin{aligned} & \left\{ \frac{\partial}{\partial \alpha_g} S_{\beta}(X; \gamma, \alpha_w, \alpha_g, \beta), \frac{\partial}{\partial \alpha_w} S_{\beta}(X; \gamma, \alpha_w, \alpha_g, \beta), \right. \\ & \frac{\partial}{\partial \beta} S_{\beta}(X; \gamma, \alpha_w, \alpha_g, \beta), \frac{\partial}{\partial \gamma} S_{\beta}(X; \gamma, \alpha_w, \alpha_g, \beta), \\ & \left. \frac{\partial}{\partial \alpha_g} S_g(X; \gamma, \alpha_g), \frac{\partial}{\partial \gamma} S_g(X; \gamma, \alpha_g), \frac{\partial}{\partial \gamma} S_{num}(X; \gamma), \right. \end{aligned}$$

$$\begin{aligned} & \left. \frac{\partial}{\partial \alpha_w} S_w(X; \alpha_w), S_w(X; \alpha_w)^{\otimes 2}, S_g(X; \gamma, \alpha_g)^{\otimes 2}, \right. \\ & \left. S_{\beta}(X; \gamma, \alpha_w, \alpha_g, \beta_0)^{\otimes 2}, S_{num}(X; \gamma)^{\otimes 2} \right\} \\ & \text{for } (\gamma, \alpha_w, \alpha_g, \beta) \in \Theta \end{aligned}$$

is a Glivenko–Cantelli class. (For vector V , $V^{\otimes 2} = VV^T$.)

3' Assume that $I_w = E_{obs} \frac{\partial}{\partial \alpha_w} S_w(X; \alpha_w)$ is invertible at $\alpha_w = \alpha_w^*$, $I_g = E_{obs} \frac{\partial}{\partial \alpha_g} S_g(X; \alpha_g)$ is invertible at $\alpha_g = \alpha_g^*$, $I_{num} = E_{obs} \frac{\partial}{\partial \gamma} S_{num}(X; \gamma)$ is invertible at $\gamma = \gamma^*$, and $B = E_{obs} \dot{\mu}(\beta_0^*, Z) \dot{\mu}(\beta_0^*, Z)^T$ is invertible.

Lemma A.4. Suppose that either $\bar{\pi}_K$ or the g_t 's are correctly parameterized. Assuming 1' through 3', there exists a sequence of $\hat{\beta}$ solutions to (5.7) for which $\sqrt{n}(\hat{\beta} - \beta_0^*)$ converges in distribution to a multivariate normal with mean zero and var-covariance matrix $B^{-1} \Sigma^* B^{-1}$, where $\Sigma^* = E_{obs}(U(X))^{\otimes 2}$.

Furthermore, $\hat{B}^{-1} \hat{\Sigma}^* \hat{B}^{-1}$ is a consistent estimator of $B^{-1} \Sigma^* B^{-1}$.

Proof. The proof is similar to that of the previous lemma and is omitted.

[Received August 2000. Revised February 2001.]

REFERENCES

Aitchison, J., and Silvey, S. D. (1958), "Maximum-Likelihood Estimation of Parameters Subject to Restraints," *Annals of Mathematical Statistics*, 29, 813–828.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996), "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91, 434–471.

Bickel, P., Klaassen, C., Ritov, Y., and Wellner, J. (1993), *Efficient and Adaptive Estimation for Semiparametric Models*, Baltimore: Johns Hopkins University Press.

Bierman, K. L., Greenberg, M. T., and Conduct Problems Prevention Research Group (1996), "Social Skill Training in the Fast Track Program," in *Preventing Childhood Disorders, Substance Abuse, and Delinquency*, eds. R. DeV. Peters and R. J. McMahon, Newbury Park, CA: Sage.

Borhani, N. O., Applegate, W. B., Cutler, J. A., Davis, B. R., Furberg, C. D., Lakatos, E., Page, L., Perry, H. M., Smith, W. M., and Probstfield, J. L. (1991), "Part 1: Rationale and Design," *Hypertension*, 17, II-2–II-15.

Conduct Problems Prevention Research Group (1992), A Developmental and Clinical Model for the Prevention of Conduct Disorders: The Fast Track Program," *Development and Psychopathology*, 4, 509–528.

Conduct Problems Prevention Research Group (to appear a), "Initial Impact of the Fast Track Prevention Trial for Conduct Problems: I. The High-Risk Sample," *Journal of Consulting and Clinical Psychology*.

Conduct Problems Prevention Research Group (to appear b), "Initial Impact of the Fast Track Prevention Trial for Conduct Problems: II. Classroom Effects," *Journal of Consulting and Clinical Psychology*.

Cooperative Research Group (1988), "Rationale and Design of a Randomized Clinical on Prevention of Stroke in Isolated Systolic Hypertension," *Journal of Clinical Epidemiology*, 41, 1197–1208.

Cox, D. R. (1958), *The Design and Planning of Experiments*, London: Chapman & Hall.

Hernán, M. A., Brumback, B., and Robins, J. M. (2000), "Marginal Structural Models to Estimate the Causal Effect of Zidovudine on the Survival of HIV-Positive Men," *Epidemiology*, 11, 561–570.

Holland, P. (1986), "Statistics and Causal Inference," *Journal of the American Statistical Association*, 81, 945–961.

Manski, C. F. (1995), *Identification Problems in the Social Sciences*, Cambridge, MA: Harvard University Press.

McMahon, R. J., Slough, N., and Conduct Problems Prevention Research Group (1996), "Family-Based Intervention in the Fast Track Program," in *Preventing Childhood Disorders, Substance Abuse, and Delinquency*, eds. R. DeV. Peters and R. J. McMahon, Newbury Park, CA: Sage.

- Neyman, J. (1990), "On the Application of Probability Theory to Agricultural Experiments," *Statistical Science*, 5, 465–480.
- Robins, J. M. (1986), "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, 1393–1512.
- 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.
- Robins, J. M. (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, U.S. Public Health Service, pp. 113–159.
- Robins, J. M. (1993), "Information Recovery and Bias Adjustment in Proportional Hazards Regression Analysis of Randomized Trials Using Surrogate Markers," in *Proceedings of the Biopharmaceutical Section, American Statistical Association*, pp. 24–33.
- Robins, J. M. (1997), "Causal Inference From Complex Longitudinal Data," in *Latent Variable Modeling and Applications to Causality*, ed. M. Berkane, *Lecture Notes in Statistics*, 120, New York: Springer-Verlag, pp. 69–117.
- Robins, J. M. (1998), "Marginal Structural Models," *1997 Proceedings of the American Statistical Association, Section on Bayesian Statistical Science*, pp. 1–10.
- Robins, J. M. (1999), "Testing and Estimation of Direct Effects by Reparameterizing Directed Acyclic Graphs With Structural Nested Models," in *Computation, Causation, and Discovery*, eds. C. Glymour and G. Cooper, Menlo Park, CA & Cambridge, MA, AAAI Press/The MIT Press, pp. 349–405.
- Robins, J. M. (to appear), "Robust Estimation in Sequentially Ignorable Missing Data and Causal Inference Models," *Proceedings of the American Statistical Association*.
- Robins, J. M., Greenland, S., and Hu, F. C., (1999a), "Estimation of the Causal Effect of a Time-Varying Exposure on the Marginal Mean of a Repeated Binary Outcome," *Journal of the American Statistical Association*, 94, 687–700.
- Robins, J. M., Hernán, M. A., and Brumback, B. (2000), "Marginal Structural Models and Causal Inference in Epidemiology," *Epidemiology*, 11, 550–560.
- Robins, J. M., and Ritov, Y. (1997), "Toward a Curse of Dimensionality Appropriate (CODA) Asymptotic Theory for Semi-Parametric Models," *Statistics in Medicine*, 16, 285–319.
- 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, MA: Birkhäuser, pp. 297–331.
- Robins, J. M., and Rotnitzky, A. (1995), "Semiparametric Efficiency in Multivariate Regression Models With Missing Data," *Journal of the American Statistical Association*, 90, 122–129.
- Robins, J. M., Rotnitzky, A., and Scharfstein, D. O. (1999b), "Sensitivity Analysis for Selection Bias and Unmeasured Confounding in Missing Data and Causal Inference Models," in *Statistical Methods in Epidemiology*, ed. E. Halloran, New York: Springer-Verlag, pp. 1–92.
- Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994), "Estimation of Regression Coefficients When Some Regressors are Not Always Observed," *J. Am. Stat. Assoc.* 89, 846–866.
- Robins, J. M., and Wasserman, L. (1997), "Estimation of Effects of Sequential Treatments by Reparameterizing Directed Acyclic Graphs," in *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, eds. D. Geiger and P. Shenoy, San Francisco: Morgan Kaufmann.
- Rotnitzky, A., and Robins, J. M. (1995), "Semiparametric Regression Estimation in the Presence of Dependent Censoring," *Biometrika*, 82, 805–820.
- Rubin, D. B. (1978), "Bayesian Inference for Causal Effects: The Role of Randomization," *The Annals of Statistics*, 6, 34–58.
- Rubin, D. B., (1986), Comment: "Which Ifs Have Causal Answers," *Journal of the American Statistical Association*, 81, 961–962.
- Scharfstein, D. O., Rotnitzky, A., and Robins, J. M. (1999), "Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models," *Journal of the American Statistical Association*, 94, 1096–1120.