Marginal Mean Models for Dynamic Regimes

S. A. Murphy, **M. J. van der Laan**, **J. M. Robins**, and **CPPRG**[1]

Abstract

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, the use of 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.

Keywords

dynamic treatment regimes; nondynamic treatment regimes; causal inference; confounding

1. INTRODUCTION

Dynamic treatment regimes are individually tailored treatments that provide treatment to individuals only when and if they need the treatment and 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 prior to the beginning of treatment; these rules are based on time varying measurements of subject-specific need for treatment. 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 paper is to provide methodology for estimation of the marginal mean response to a dynamic treatment regime when the data is observational. We employ a structural model in that it is a model for the marginal mean of counterfactual or potential responses.

This work is 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 due to elevated behavior problems (Bierman et al., 1996; CPPRG, 1999ab; McMahon et al., 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 6 item "home visiting process measure," describing the quality of parenting and family functioning. Based on 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 greater corresponded to 4 home visits, a score from 9

to 16 corresponded to 8 home visits and a score of 8 or below corresponded to 16 home visits during the following semester. Thus the rule for assigning dose was

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

where S_t is an assessment of severity, the home visiting process measure score, at the end of the tth 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 above dynamic treatment regime would have been, had 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 necessitating cessation of treatment or when the optimal dosage at a given time should depend on the subjects 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 pre-determined goal. The treatment schedule was quite complex. Initially, all treated subjects were assigned the same dosage of anti-hypertensive medication; at eight weeks, subjects whose SBP exceeded the goal had their dosage doubled; at sixteen weeks, subjects whose SBP still exceeded the goal had a second anti-hypertensive added. At twenty-four weeks, a subject whose SBP continued to exceed the goal had their dose of the second anti-hypertensive doubled. The treatment schedule specified that subjects with evidence of severe clinical or laboratory toxicity would first have their dose of anti-hypertensive 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 due to side effects is a useful aid in reducing unexplained noncompliance. Indeed subjects whose medication is switched in, accordance with the rules are compliers.

In this paper we consider methodology for estimation of the mean response to both nonrandom and random dynamic treatment regimes when the available data is 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 the present need. In Section 3 we provide quantitative definitions of these two types of dynamic treatment regimes. The modeling of responses to dynamic treatment regimes based on observational data has received very little statistical attention outside of a series of papers by Robins (1986, [1989], [1993], 1997). In these papers Robins considers 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) discusses estimating parameters in the response to a nonrandom dynamic treatment regime by censoring the subject at the first time the subject's treatment differs from the treatment as specified by the dynamic regime.

In Section 2 we review the underlying causal theory based on counterfactual or potential outcomes. Next 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 (1935) 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 first specify our observations in a unified way regardless of the manner in which treatment is selected/assigned.

Suppose that the treatment lasts for K intervals; during this time 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 $\bar{a}_K = (a_1, a_2, \ldots, a_K)$; a_t is the level of treatment in the tth interval. In general we use a bar over a variable to denote that variable and all past values of the same variable, so $\bar{a}_t = (a_1, \ldots, a_t)$. Let $\mathscr{A}_K$ 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 $Y(\bar{a}_K)$. 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). That is, first we assume that the treatment is defined so that the treatment selection mechanism does not alter the subject's potential responses; for example, it does not matter why the subject was treated (e.g., whether the subject is randomly assigned the treatment or whether the subject's parent solicited the treatment), the entirety of the subject's intermediate potential outcomes and the potential responses are the same. In the time varying treatment setting this is the consistency assumption of Robins (1997). And second 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. These assumptions allow us to conceptualize a single potential outcome/ response corresponding to each possible $\bar{a}_K \in \mathscr{A}_K$ 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; Angrist, Imbens and Rubin, 1996).

SUTVA may well be violated in the Fast Track data. In addition to the home visiting, the Fast Track intervention included "friendship groups" in which small numbers of the intervention children were brought together in order to improve social skills. Thus in close friendships, the effect of home visiting on one child may affect the friend's response. Alternately whether or not 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 as an appropriate generalization of the proposed methodology would distract from the main points of this paper.

Capital letters are used to denote random variables. Treatment selection/assignment is allowed to be stochastic and are denoted by the random vector, $\bar{A}_K$. So $\bar{A}_K$ takes values in $\mathscr{A}_K$. Denote the time t ($t > 0$) variables 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 prior to and at time t. (S_t, V_t) is an intermediate outcome of 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 above. Corresponding to each treatment vector, $\bar{a}_K$, SUTVA implies that the intermediate outcome at time t can be written unambiguously as $(S_t(\bar{a}_t), V_t(\bar{a}_t))$ and the potential response as $Y(\bar{a}_K)$.

In both randomized and observational studies, the observed data is the pretreatment information plus the potential outcomes corresponding to the treatment pattern $\bar{A}_K$. Assuming SUTVA we may write the data in a unified way regardless of the manner in which treatment is selected/ assigned. A subject's potential outcomes, that is the "complete" severity and response data is $O_{sr} = \{Z, S_0, S_1(a_1), \ldots, S_{K-1}(\bar{a}_{K-1}), Y(\bar{a}_K); \bar{a}_K \in \mathscr{A}_K\}$ and the complete auxiliary data is $O_{aux} = \{V_0, V_1(a_1), \ldots, V_{K-1}(\bar{a}_{K-1}); \bar{a}_K \in \mathscr{A}_K\}$. In a given experiment or observational study we observe only a subset of the subject's outcomes plus the treatment pattern, $X = \{Z, S_0, V_0, A_1, S_1(A_1), V_1(A_1), A_2, \ldots, S_{K-1}(\bar{A}_{K-1}), V_{K-1}(\bar{A}_{K-1}), A_K, Y(\bar{A}_K)\}$. Additionally in randomized studies we also have data on the treatment regime to which a study subject was assigned in addition, to the observed data, X. In the following we use either Y or $Y(\bar{A}_K)$ to denote the observable response, either S_t or $S_t(\bar{A}_t)$ to denote the observable severity at time t and either V_t or $V_t(\bar{A}_t)$ to denote the observable auxiliary variables at time t.

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, $\bar{d}_K = (d_1, \ldots, d_K)$; 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 $\bar{p}_K = (p_1, \ldots, p_K)$; 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 prior to 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 post-treatment observations (i.e. the rule, d_t is not a function of post-treatment 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 which 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 et al. (1998) and Hernán et al. (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/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 section 11, of Robins, 1997). In this paper we assume:

Sequential Randomization

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

That is, the auxiliary data, the V's, are sufficiently rich so the observational distribution satisfies the above sequential randomization. This means that within levels of $L_0, A_1, L_1, A_2, …, L_{t-1}$, 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 have chosen 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 above 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. THE ESTIMAND

The sequential randomization assumption is used only to identify the parameter of interest (e.g., a treatment regime effect) in terms of the observational distribution. The identification is achieved by the following lemma. Suppose we are given a vector of treatment rules, $\bar{p}_K$. We wish to estimate parameters in $P_{\bar{p}_K}$ distribution using data from the observational data, the P_{obs} distribution. 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

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

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

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

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

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

Proof—Below 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

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

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,

$$\sum_{\bar{a}_K}E_{obs}\left[I\{(Y(\bar{a}_K), \bar{S}_{K-1}(\bar{a}_{K-1}), \bar{a}_K, Z) \in U\}B_1(O_{sr}, \bar{a}_K))\prod_{t=1}^{K}p_t(a_t|\bar{S}_{t-1}(\bar{a}_{t-1}))\right] \tag{4.3}$$

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,

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

Since the distribution of O_{sr} is the same under both $P_{\bar{p}_K}$ and P_{obs} distributions, (4.3) can also 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

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

Since B_1 can be expressed as repeated expectations of a probability (b_K) which is nessarily bounded above by one, (4.1) implies that $B_1(O_{sr},\overline{A}_K)=1$ $a.e.$ Thus we have absolute continuity and (4.2) is the Radon-Nikodym derivative of $P_{\bar{p}_K}$ with respect to P_{obs}.

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 occuring in the observational study." This assumption is emininently sensible since if a particular treatment pattern cannot occur in the observational study, then inference involving the response to this treatment pattern requires further knowledge/assumptions.

From a statistical perspective there are two important consequences of the fact that (4.2) is a Radon-Nikodym derivative. Recall that the goal of this paper is estimate, using observational data, the mean response that would have been observed had, contrary to fact, the entire study population followed a particular dynamic regime. Thus the estimand is

$$E_{\bar{p}_K}[Y|Z]=E_{obs}\left[E_{obs}[W_{\bar{p}_K}(\overline{A}_K,\overline{L}_{K-1})|Y,\overline{S}_{K-1},\overline{A}_K,Z]\,Y\big|Z\right]$$
$$=E_{obs}\left[W_{\bar{p}_K}(\overline{A}_K,\overline{L}_{K-1})Y\big|Z\right]$$

and this estimand can be expressed in terms of the observational distribution. Or writing the right hand side in terms of conditional densities, $E_{\bar{p}K}[Y|Z=z]=$

$$\sum_{\bar{a}_K\in\mathscr{A}}\int_{\bar{l}_{K-1}}\int_y y f_Y(y|\bar{a}_K,\bar{v}_{K-1},\bar{s}_{K-1})\prod_{t=1}^{K}p_t(a_t|s_{t-1})\prod_{t=1}^{K-1}f_t(l_t|\bar{a}_t,\bar{l}_{t-1})f_0(l_0|z)dy\,d\bar{l}_{K-1}$$

(4.4)

where $f_Y,f_K,\ldots,f_0$ denote conditional densities for the observational distribution. The conditional density of Y given $\overline{A}_K,\overline{L}_{K-1}$ is f_Y and the conditional density of L_t given $\overline{A}_{t-1},\overline{L}_{t-1}$ is f_t. Thus the conditional mean is a function of the observational distribution via the conditional densities f_Y and $f_{K-1},\ldots,f_0$. This is Robins (1986) G-computation formula.

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,

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

(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 in order to estimate parameters of marginal structural models and direct effect structural models for time varying treatments. Robins refers to such estimators as inverse probability of treatment weighted estimators.

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/non-adherence 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 which is the product of the modelled adherence probabilities. Note that in this paper, 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/non-adherence 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 non-adherence to the treatment regime and $\bar{A}_{t-1} = \bar{a}_{t-1}$, $\bar{L}_{t-1} = \bar{l}_{t-1}$.

We may view $\mu(\beta, Z)$ as a parameterization of (4.4). 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. If sequential randomization does not hold, then we only have that the estimation procedure will, under general conditions, result in a valid estimator of (4.4). The sequential randomization assumption allows us to intrepret this estimand as the mean response to the regime $\bar{p}_K$.

5. ESTIMATION

Our goal is to conduct inference concerning the mean response to a dynamic treatment regime, possibly conditional on pretreatment variables. Assuming SUTVA, sequential randomization and (4.1), (4.4) can be interpreted as the mean response to a dynamic treatment regime, conditional on Z. We separate this interpretation from the estimation. In this section we consider a parameterization of (4.4) denoted by $\mu(\beta, Z)$. In order to estimate β we no longer need SUTVA, sequential randomization or (4.1); indeed we no longer need to assume the existnec of potential outcomes. However we assume regularity conditions for the asymptotic theory and that

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

(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 $f_{k-1}, \ldots, f_0$. On-the other hand, the likelihood of the observational data can be written as a product of conditional densities:

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

(5.2)

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 $(\bar{\pi}_K)$. An optimal estimating function (one resulting in an efficient estimator of β) will be orthogonal to the score functions for the treatment selection probabilities (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). Recall the observed data for one person is denoted by X. Given an estimating function, say $U(X,\beta)$, the projection of U 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).

Setting $U(X, \beta)$ equal to $W_{\bar{p}_K}(\bar{A}_K, \bar{L}_K)\dot{\mu}(\beta, Z)(Y - \mu(\beta, Z))$, we get after algebraic simplification that

$$U(X,\beta) - \sum_{t=1}^{K} E_{obs}[U(X,\beta)|\bar{L}_{t-1}, \bar{A}_t] + E_{obs}[U(X,\beta)|\bar{L}_{t-1}, \bar{A}_{t-1}] =$$

$$W_{\bar{p}_K}(\bar{A}_K, \bar{L}_{K-1})\dot{\mu}(\beta, Z)(Y - \mu(\beta, Z)) - \sum_{t=1}^{K} W_{\bar{p}_t}(\bar{A}_t, \bar{L}_{t-1})\dot{\mu}(\beta, Z)(g_t(\bar{A}_t, \bar{L}_{t-1}) - \mu(\beta, Z))$$

$$+\sum_{t=1}^{K}\sum_{a_t} \pi_t(a_t|\bar{A}_{t-1}, \bar{L}_{t-1}) W_{\bar{p}_t}(a_t, \bar{A}_{t-1}, \bar{L}_{t-1})\dot{\mu}(\beta, Z)(g_t(a_t, \bar{A}_{t-1}, \bar{L}_{t-1}) - \mu(\beta, Z))$$

(5.3)

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

$$g_K(\bar{A}_K, \bar{L}_{K-1}) = E_{obs}(Y|\bar{A}_K, \bar{L}_{K-1})$$

$$g_t(\bar{A}_t, \bar{L}_{t-1}) = E_{obs}\left(\sum_{a_{t+1}} p_{t+1}(a_{t+1}|S_t)g_{t+1}(a_{t+1}, \bar{A}_t, \bar{L}_t)\Big|\bar{A}_t, \bar{L}_{t-1}\right)$$

(5.4)

for $t = 1,\ldots, K - 1$. If the parameterization of $\mu(\beta, Z)$ is saturated then the right hand 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 $\{g_t(\bar{a}_t, \bar{l}_{t-1}), t = 1,\ldots,K\}$. Given a parametrization of the treatment selection probabilities, say $\pi_t(a_t|\bar{a}_{t-1}, \bar{l}_{t-1}; \alpha_w)$, $t = 1,\ldots,K$ with parameter, α_w, we can use maximum likelihood based on (5.2) to estimate α_w. We maximize,

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

(5.5)

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

$$0 = \mathbb{P}_n \left(Y - g_K(\overline{A}_K, \overline{L}_{K-1}; \alpha_g) \right) \frac{\partial}{\partial \alpha_g} g_K(\overline{A}_K, \overline{L}_{K-1}; \alpha_g) +$$

$$\sum_{t=1}^{K-1} \mathbb{P}_n \left(\sum_{a_{t+1}} p_{t+1}(a_{t+1}|S_t) g_{t+1}(a_{t+1}, \overline{A}_t, \overline{L}_t; \alpha_g) - g_t(\overline{A}_t, \overline{L}_{t-1}; \alpha_g) \right) \frac{\partial}{\partial \alpha_g} g_t(\overline{A}_t, \overline{L}_{t-1}; \alpha_g) \tag{5.6}$$

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,

$$\mathbb{P}_n W_{\overline{p}_K}(\overline{A}_K, \overline{L}_{K-1}; \widehat{\alpha}_w) \dot{\mu}(\beta, Z)(Y - \mu(\beta, Z))$$

$$- \mathbb{P}_n \sum_{t=1}^{K} W_{\overline{p}_t}(\overline{A}_t, \overline{L}_{t-1}; \widehat{\alpha}_w) \dot{\mu}(\beta, Z)(g_t(\overline{A}_t, \overline{L}_{t-1}; \widehat{\alpha}_g) - \mu(\beta, Z))$$

$$+ \mathbb{P}_n \sum_{t=1}^{K} \sum_{a_t} \pi_t(a_t|\overline{A}_{t-1}, \overline{L}_{t-1}; \widehat{\alpha}_w) W_{\overline{p}_t}(a_t, \overline{A}_{t-1}, \overline{L}_{t-1}; \widehat{\alpha}_w) \dot{\mu}(\beta, Z)$$

$$(g_t(a_t, \overline{A}_{t-1}, \overline{L}_{t-1}; \widehat{\alpha}_g) - \mu(\beta, Z)), \tag{5.7}$$

to 0 and solve for β.

Under regularity conditions, the following two properties hold.

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

2. The asymptotic variance of $\hat{\beta}$ can be estimated by a sandwich estimator. Define $\hat{\Sigma} =$

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

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 $\mathbb{P}_n S_w(X; \alpha_w)$), $\hat{I}_w$ 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 $\mathbb{P}_n S_g$),

$\hat{I}_g$ is minus the derivative of (5.6) with respect to α_g, $\widehat{S}_{\beta,w} = \mathbb{P}_n \dfrac{\partial}{\partial \alpha_w} S_\beta(X; \widehat{\alpha}_w, \widehat{\alpha}_g, \widehat{\beta})$ and

$\widehat{S}_{\beta,g} = \mathbb{P}_n \dfrac{\partial}{\partial \alpha_g} S_\beta(X; \widehat{\alpha}_w, \widehat{\alpha}_g, \widehat{\beta})$. Also define $\hat{B} = \mathbb{P}_n \dot{\mu}(\beta, Z) \dot{\mu}(\beta, Z)^T$. Then a consistent estimator of the asymptotic variance of $\hat{\beta}$ is given by $\hat{B}^{-1} \hat{\Sigma} \hat{B}^{-1}$.

See the appendix for proofs of these results.

Consistent Parameterization of $g_1, \dots, g_K$

The relationships among the $g_1, \dots, g_K$ are highly constrained. Thus it is difficult to formulate parameterizations of $g_1, \dots, g_K$ which are consistent with one another. For example, it must be possible for the parameterization of g_K and g_{K-1} to satisfy

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

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 a parametrization of g_K; and g_{K-1} there must be at least some conditional density f_{K-1} for which the above display is true. See the data example for a consistent parametrization. Also see Robins and Wasserman (1997) and Scharfstein et al. (1999) for similar problems with consistent parametrization.

Double Robustness

Instead of using (5.7) to estimate β, we could use (4.5) with estimated weights. Consistency of the resulting β̂ 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

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

(5.8)

As noted previously, this is the G-computation formula of Robins (1986). Thus to directly estimate β using (4.4) we may estimate the g_t's in order to arrive at an estimate of g_1 and then solve

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

(5.9)

We can only expect the resulting β̂ to be a consistent estimator of β if we have correctly parameterized the g_t's.

The primary 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 π_t's are parameterized correctly (see (1) above). This is the "double robustness" property; see Scharfstein et al. (1999) and Robins et al. (2000) 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) so as 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, multi-site, multi-cohort trial of a preventive intervention versus a control. For the purposes of this illustration we use a subset of the children for whom data was 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 grade 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,

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

in assigning 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, i.e.,

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

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%) have treatment recommendation patterns that coincide with the rule for all 4 time periods. The staff did not/could not deviate from the rule for the control children as 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 will use past recommended intervention level, the home visiting process measure and site and cohort indicators to model the distribution of home visiting recommendations (these will be 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 $\{L_0, A_1 L_1, A_2, \ldots, L_{t-1}\}$ 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 recommended 16 home visits 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 recommended 4 home visits. It also means that (within a cross-stratification) children recommended 16 home visits do not have higher chance of high values of $Y(16, 16, 16, 16) - Y(4, 4, 4, 4)$ than children recommended 4 home visits.

To implement (5.7) we first estimate the weights, $W_{\bar{p}4}$. Since the staff were not to and did not assign any intervention to control children, the weights for the control children are equal to one. Note that the numerator of $W_{\bar{p}4}$ is one for an intervention child whose treatment recommendations coincide with the treatment rule and zero otherwise. This means that of the 202 intervention children, 98 (202–104) will have $W_{\bar{p}4} = 0$ and thus the first term in (5.7) will be zero 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 below. To estimate the denominator of the weight, we fit a proportional odds model to the staff home visiting recommendations. For each time t, we fit, via maximum likelihood,

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

where $\alpha_w = \{\alpha_{w,t01}, \alpha_{w,t02}, \alpha_{w,1}, t = 1, 2, 3, 4\}$ and $ZZ(\overline{L}_{t-1}, \overline{A}_{t-1})$ is a vector summary of past Information. See Table 2 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. Since there were only 4 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 one model with possibly time varying intercepts and regression coefficients to all 4 intervals simultaneously (see Hernán et al. for an example, 1998a).

To implement (5.7) we must also estimate α_g. We model g_1 through g_4 by linear models:

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

where $XX(\overline{L}_{t-1})$ is a summary vector of past information and estimate $\alpha_g = \{\alpha_{g0}, \alpha_{g1}\}$ by solving (5.6). Note that we assume that the g_t are functionally independent of the past assigned home visiting levels ($\overline{A}_t$). It is easy to check that this assumption plus the linearity of g_t imply that the parameterizations of $g1, \ldots, g_K$ will be consistent with one another. In this example the assumption that the g_t are functionally independent of the past assigned home visiting levels is plausible. Preliminary analyses (not shown) found that given past measures of L, both the most recent assigned home visiting levels and the average level of past home visiting were not predictive in the g_t's. This may not be surprising since the composition of children with a high past level of home visiting is strongly skewed toward children with higher severity whereas the composition of the children with a lower past level of home visiting is skewed toward children with low severity. Indeed if home visiting is effective, one could expect that within levels of past L, higher severity children will be similar to lower severity children after both groups receive their respective treatments. The estimates of α_g are given in Table 3.

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

$$\mu(\beta, Z, treatment) = \beta_0 + \beta_1 S BP1 + \beta_2 treatment$$

where *treatment* is an indicator with the value one for the intervention children and when grade3 teacher rating of school behavior problems is the response, grade 1 teacher rating of school behavior problems *SBP*1 is included in the regression. When the teacher rating of behavior change over the course of grade 3 is the response, *SBP*1 is not included. 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 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 4a, 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 Table 2) 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 and/or due to using the home visiting measure only as prescribed by the regime rule. This is discussed further below. In contrast, in Table 4b, we see that while there is a significant ITT treatment effect on 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. Similarly to before it is unclear whether the decreased treatment effect is due to forced uniformity across site and cohort in assigning treatment and/or due to ignoring past level of assigned treatment in making future recommendations and/or 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. In order to better utilize these responses and also 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 $\bar{p}_{4,h}$, $h \in \{0, .2, .4, \ldots, 1\}$. $p_{t,h}(a|s_{t-1})$ 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 of these random dynamic regimes do not allow site, cohort and past home visiting recommendations to influence 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 one 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:

$$\text{logit}\,(q_t(16|S_{t-1};\gamma)) = \gamma_{t01} + \gamma_{t1}I\{d(S_{t-1}) = 16\} + \gamma_{t2}I\{d(S_{t-1}) = 4\}$$
$$\text{logit}\,(1 - q_t(4|S_{t-1};\gamma)) = \gamma_{t02} + \gamma_{t1}I\{d(S_{t-1}) = 16\} + \gamma_{t2}I\{d(S_{t-1}) = 4\}$$

where $\gamma = \{\gamma_{t01}, \gamma_{t02}\, \gamma_{t1}, \gamma_{t2}, t = 1, 2, 3, 4\}$. The estimated values of γ are given in Table 5. 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 the q_t and π_t will not be consistent with one another. However we only view q_t as an approximation

to the observational distribution of treatment recommendation given severity. Also note that the $p_{t,h}$ $(\cdot|S_{t-1}; \hat{\gamma})$ are used to form the numerator of the weights and as before the $\pi_t(\cdot|\bar{L}_{t-1}, \bar{A}_{t-1}; \hat{\alpha}_w)$ 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; Table 4 contains the results. The formula for the estimated standard errors changes since the numerator of the weights is now estimated. See the appendix for the appropriate formula. Consider the results for $h = 0$ in Table 4. 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 even though 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 analyses. Certainly one should give careful thought to which responses are of primary interest; variations in the dynamic treatment regime may differentially impact different responses.

7. DISCUSSION

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 of the sensitivity analyses of Robins et al. (1999a) and Scharfstein et al. (1999) to this setting is needed. A second needed extension is to allow for missing severity measures. There are two ways to allow for missing severity measures. 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 includes individuals with missing severity measures then the methodology provided here must be adapted. One possibility is by weighting as described in Hernán et al. (1998). A. third extension would be to grouped data. Many intervention/prevention studies take place in a school based setting and thus grouped data dominates. Lastly more systematic work is needed on how to formulate/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.

Acknowledgments

APPENDIX

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

$$S_\beta(X;\alpha_w,\alpha_g,\beta) = W_{\bar{p}_K}(\bar{A}_K,\bar{L}_{K-1})\dot{\mu}(\beta,Z)(Y-\mu(\beta,Z))$$

$$-\sum_{t=1}^{K}\sum_{a_t}\left(I\{A_t=a_t\}-\pi_t(a_t|\bar{A}_{t-1},\bar{L}_{t-1})\right)W_{\bar{p}_t}(a_t,\bar{A}_{t-1},\bar{L}_{t-1})\dot{\mu}(\beta,Z)$$

$$(g_t(a_t,\bar{A}_{t-1},\bar{L}_{t-1})-\mu(\beta,Z)). \tag{8.1}$$

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

$$S_\beta(X;\alpha_w,\alpha_g,\beta) = W_{\bar{p}_K}(\bar{A}_K,\bar{L}_{K-1})\dot{\mu}(\beta,Z)(Y-g_K(\bar{A}_K,\bar{L}_{K-1}))$$

$$-\sum_{t=1}^{K-1}W_{\bar{p}_t}(\bar{A}_t,\bar{L}_{t-1})\dot{\mu}(\beta,Z)\left(g_t(\bar{A}_t,\bar{L}_{t-1})-\sum_{a_{t+1}}p_{t+1}(a_{t+1}|S_t)g_{t+1}(a_t,\bar{A}_t,\bar{L}_t)\right)$$

$$+\dot{\mu}(\beta,Z)\left(\sum_{a_1}g_1(a_1,L_0)p_1(a_1|S_0)-\mu(\beta,Z)\right). \tag{8.2}$$

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 and/or $\beta \neq \beta_0$.

Lemma 8.1

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 $t \in \{1,\ldots,K\}$ and for $U_t(\bar{A}_{t-1},\bar{L}_{t-1})$ for $t \in \{1,\ldots,K\}$ arbitrary bounded functions. Thus we need only show that (5.3) or equivalently (8.1) 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 (8.1). The covariance between $B_j(a,\bar{A}_j,L_{j-1})$ and (8.1) is thus

$$E_{obs}[W_{\bar{p}_K}(\bar{A}_K,\bar{L}_{K-1})\dot{\mu}(\beta,Z)(Y-\mu(\beta,Z))B_j(a,\bar{A}_j,\bar{L}_{j-1})]$$

$$-E_{obs}\left(\sum_{a_j}(I\{A_j=a_j\}-\pi_j(a_j|\bar{A}_{j-1},\bar{L}_{j-1}))\right. \tag{8.3}$$

$$\left.W_{\bar{p}_j}(a_j,\bar{A}_{j-1},\bar{L}_{j-1})\dot{\mu}(\beta,Z)(g_j(a_j,\bar{A}_{j-1},\bar{L}_{j-1})-\mu(\beta,Z))\right)B_j(a,\bar{A}_j,\bar{L}_{j-1}). \tag{8.4}$$

Since given $\bar{A}_{j-1},\bar{L}_{j-1}$, $B_j(a,\bar{A}_j,L_{j-1})$ has mean zero, the second term (8.4) reduces to

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

$$B_j(a,\overline{A}_j,\overline{L}_{j-1}).$$

But this is the negative of the first term (8.3) (repeatedly take conditional expectations and use (5.1)).

Lemma 8.2

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

Proof

Because of the factorization, (5.2) of the likelihood and the fact that the functional $\mu(\beta, Z)$ does not depend on the factor $\prod_{t=1}^{K}\pi_t(a_t|\overline{a}_{t-1},\overline{l}_{t-1})$, the efficient influence function is the same in the model for which $\overline{\pi}_K$ is completely unknown as in the model where $\overline{\pi}_K$ 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 $\overline{\pi}_K$ completely unknown is a nonparametric model for the observed data. Bickel et al. (1993) show 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 $\overline{\pi}_K$ completely unknown, this will be the efficient influence function for this nonparametric model and also be the efficient influence function for the model in which $\overline{\pi}_K$ is parameterized. Now when $\overline{\pi}_K$ is completely known, $U(X,\beta,\overline{p}_K) = W_{\overline{p}_K}(\overline{A}_K,\overline{L}_K)\dot\mu(\beta,Z)(Y-\mu(\beta,Z))$ 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 $\overline{\pi}_K$ (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 $\overline{\pi}_K$ is parameterized.

Assume:

1. There exists a finite vector α_w^* such that $P_{obs}(S_w(X;\alpha_w^*))=0$ where $\mathbb{P}_nS_w(X;\alpha_w)$ 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 $\mathbb{P}_nS_g(X;\alpha_g)$ 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 $\overline{\pi}_K$ is correctly parameterized then put $\alpha_w^*=\alpha_{w0}$, the true parameter in the observational distribution. Similarly if $\{g_1,\ldots,g_K\}$ 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

$$\left\{ \frac{\partial}{\partial\alpha_g}S_\beta(X;\alpha_w,\alpha_g,\beta),\ \frac{\partial}{\partial\alpha_w}S_\beta(X;\alpha_w,\alpha_g,\beta),\ \frac{\partial}{\partial\beta}S_\beta(X;\alpha_w,\alpha_g,\beta),\ \frac{\partial}{\partial\alpha_w}S_w(X;\alpha_w) \right.$$
$$\left. \frac{\partial}{\partial\alpha_g}S_g(X;\alpha_g), S_w(X;\alpha_w)^{\otimes2}, S_g(X;\alpha_g)^{\otimes2}, S_\beta(X;\gamma,\alpha_w,\alpha_g,\beta)^{\otimes2} \text{ for } (\alpha_w,\alpha_g,\beta)\in\Theta \right\}$$

    is a Glivenko-Cantelli class. (For vector V, $V^{\otimes2} = VV^T$.)

3.

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

Lemma 8.3

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

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

where $S_{\beta, \alpha_w} = E_{obs} \dfrac{\partial}{\partial \alpha_w} S_\beta(X; \alpha_w, \alpha_g^*, \beta_0)|_{\alpha_w = \alpha_w^*}$ and $S_{\beta, \alpha_g} = E_{obs} \dfrac{\partial}{\partial \alpha_g} S_\beta(X; \alpha_w^*, \alpha_g, \beta_0)|_{\alpha_g = \alpha_g^*}$. 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 (8.1) we have that

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

This is zero since $\mu(\beta_0, Z)$ is equal to (4.4). Alternately suppose that $\alpha_g^* = \alpha_{g0}$. Then as noted in the discussion concerning (8.2) we have that

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

As before this is zero since $\mu(\beta_0, Z)$ 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, $\widehat{\alpha}_w \overset{P}{\mapsto} \alpha_w^*$ to $\mathbb{P}_n(S_w(X; \alpha_w)) = 0$ and a consistent sequence of solutions, $\widehat{\alpha}_g \overset{P}{\mapsto} \alpha_g^*$ 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 consistence solution, $\widehat{\beta} \overset{P}{\mapsto} \beta_0$. 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

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

and that

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

Next a Taylor series yields,

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

Note from the previous discussion, the first term in the above expansion has mean zero. Next apply the Gliveriko-Cantelli property in 2), along with the invertibility in 3) to yield,

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

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

Given a parametrization 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

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

and

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

And define,

$$
\begin{aligned}
S_\beta(X; \alpha_w, \alpha_g, \gamma, \beta) &= W_{\bar{p}_K}(\overline{A}_K, \overline{L}_{K-1}; \widehat{\alpha}_w) \dot{\mu}(\beta, Z)(Y - \mu(\beta, Z)) \\
&\quad - \mathbb{P}_n \sum_{t=1}^{K} W_{\bar{p}_t}(\overline{A}_t, \overline{L}_{t-1} \widehat{\alpha}_w) \dot{\mu}(\beta, Z)(g_t(\overline{A}_t, \overline{L}_{t-1}; \widehat{\alpha}_g) - \mu(\beta, Z)) \\
&\quad + \mathbb{P}_n \sum_{t=1}^{K} \sum_{a_t} \pi_t(a_t | \overline{A}_{t-1}, \overline{L}_{t-1}; \widehat{\alpha}_w) W_{\bar{p}_t}(a_t, \overline{A}_{t-1}, \overline{L}_{t-1}; \widehat{\alpha}_w) \dot{\mu}(\beta, Z) \\
&\qquad\qquad (g_t(a_t, \overline{A}_{t-1}, \overline{L}_{t-1}; \widehat{\alpha}_g) - \mu(\beta, Z))
\end{aligned}
$$

Define

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

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

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

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

and $\widehat{S}_{\beta,w} = \mathbb{P}_n \frac{\partial}{\partial \alpha_w} S_\beta(X;\widehat{\gamma}, \alpha_w, \widehat{\alpha}_g, \widehat{\beta})|_{\alpha_w = \widehat{\alpha}_w}$, $\widehat{S}_{\beta,g} = \mathbb{P}_n \frac{\partial}{\partial \alpha_g} S_\beta(X;\widehat{\gamma}, \widehat{\alpha}_w, \alpha_g, \widehat{\beta})|_{\alpha_g = \widehat{\alpha}_g}$, and

$\widehat{S}_{g,num} = \mathbb{P}_n \frac{\partial}{\partial \gamma} S_g(X;\gamma, \widehat{\alpha}_g)|_{\gamma = \widehat{\gamma}}$.

(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 $\mu(\beta_0^*, Z) =$

$$\sum_{\overline{a}_K \in A} \int_{\overline{l}_{K-1}} \int_y y f_Y(y|\overline{a}_K, \overline{v}_{K-1}, \overline{s}_{K-1}) \prod_{t=1}^{K} p_t(a_t|s_{t-1};\gamma^*) \prod_{t=1}^{K-1} f_t(l_t|\overline{a}_t, \overline{l}_{t-1}) f_0(l_0|z) dy \, d\overline{l}_{K-1}.$$

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 $\{g_1, \ldots, g_K\}$ 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 $0 = \mathbb{P}_n S_g(X;\hat{\gamma}, \alpha_g)$. Lastly estimate β^* by solving $0 = \mathbb{P}_n S_\beta(X;\hat{\alpha}_w, \hat{\alpha}_g, \hat{\gamma}, \beta)$ for β.

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

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

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

(3′)

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

Lemma 8.4

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.

REFERENCES

1. Aitchison J, Silvey SD. Maximum-Likelihood Estimation of Parameters Subject to Restraints. Ann. Math. Statist 1958;29:813–828.

2. Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. Journal of the American Statistical Association 1996;91:434–471.

3. Bickel, P.; Klaassen, C.; Ritov, Y.; Wellner, J. Efficient and adaptive estimation for semiparametric models. Baltimore: Johns Hopkins University Press; 1993.

4. Bierman, KL.; Greenberg, MT. The Conduct Problems Prevention Research Group. Social skill training in the FAST Track program. In: Peters, R DeV; McMahon, RJ., editors. Preventing childhood disorders, substance abuse, and delinquency. Newbury Park, CA: Sage; 1996. p. 65-89.

5. Borhani NO, Applegate WB, Cutler JA, Davis BR, Furberg CD, Lakatos E, Page L, Perry HM, Smith WM, Probstfield JL. Part 1: Rationale and Design. Hypertension 1991;17:II-2–II-15. [PubMed: 1999373]

6. Conduct Problems Prevention Research Group. A developmental and clinical model for the prevention of conduct disorders: The FAST Track Program. Development and Psychopathology 1992;4:509–528.

7. Conduct Problems Prevention Research Group. Initial impact of the Fast Track prevention trial for conduct problems: I. The high-risk sample to appear in the Journal of Consulting and Clinical Psychology. 1999a

8. Conduct Problems Prevention Research Group. Initial impact of the Fast Track prevention trial for conduct problems: II. Classroom effects to appear in the Journal of Consulting and Clinical Psychology. 1999b

9. Cooperative Research Group. Rationale and design of a randomized clinical on prevention of stroke in isolated systolic hypertension. Journal of Clinical Epidemiology 1988;41:1197–1208. [PubMed: 2905387]

10. Cox, DR. The Design and Planning of Experiments. London: Chapman and Hall; 1958.

11. Hernán MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. 1998 *unpublished manuscript*.

12. Holland P. Statistics and Causal Inference. Journal of the American Statistical Association 1986;81:945–961.

13. Manski, CF. Identification Problems in the Social Sciences. Cambridge, MA: Harvard University Press; 1995.

14. McMahon, RJ.; Slough, N. Conduct Problems Prevention Research Group. Family-based intervention in the FAST Track Program. In: Peters, R DeV; McMahon, RJ., editors. Preventing childhood disorders, substance abuse, and delinquency. Newbury Park, CA: Sage; 1996. p. 65-89.

15. Neyman J, Iwaszkriewicz K, Kolodziejczyk S. Statistical Porblems in Agricultural Experimentation (with discussion). Supplement of Journal of the Royal Statistical Society bf 1935;2:107–180.

16. Robins JM. 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 1986;14:1393–1512.

17. Robins JM. 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 1987;14:923–945.

18. Robins, JM. The Analysis of Randomized and Nonrandomized AIDS Treatment Trials Using a New Approach to Causal Inference in Longitudinal Studies. In: Sechrest, L.; Freeman, H.; Mulley, A., editors. Health Service Research Methodology: A Focus on AIDS. U.S. Public Health Service: NCHSR; 1989. p. 113-159.

19. Robins, JM.; Rotnitzky, A. Recovery of information and adjustment for dependent censoring using surrogate markers. In: Jewell, N.; Dietz, K.; Farewell, V., editors. AIDS Epidemiology - Methodological Issues. Boston, MA: Birkhäuser; 1992. p. 297-331.

20. Robins JM. Information recovery and bias adjustment in proportional hazards regression analysis of randomized trials using surrogate markers. Proceedings of the Biopharmaceutical Section, American Statistical Association 1993:24–33.

21. Robins JM, Rotnitzky A. Semiparametric efficiency in multivariate regression models with missing data. Journal of the American Statistical Association 1995;90:122–129.

22. Robins, JM. Causal Inference from Complex Longitudinal Data. In: Berkane, M., editor. Latent Variable Modeling and Applications to Causality. Lecture Notes in Statistics. New York: Springer-Verlag, Inc; 1997. p. 69-117.

23. Robins, JM.; Wasserman, L. Estimation of Effects of Sequential Treatments by Reparameterizing Directed Acyclic Graphs. In: Geiger, D.; Shenoy, P., editors. Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence; Morgan Kaufmann; San Francisco. 1997. p. 409-442.

24. Robins, JM. Marginal structural models; 1997 Proceedings of the American Statistical Association. Section on Bayesian Statistical Science; 1998. p. 1-10.

25. Robins JM, Hernán MA, Brumback B. Marginal structural models and causal inference in epidemiology. 1998 *unpublished manuscript*.

26. Robins, JM.; Rotnitzky, A.; Scharfstein, DO. Sensitivity Analysis for Selection Bias and Unmeasured Confounding in Missing Data and Causal Inference Models. In: Halloran, E., editor. Statistical Methods in Epidemiology. New York: Springer-Verlag, Inc; 1999a. p. 1-92.

27. Robins JM, Greenland S, Hu FC. 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 1999b;94:687–700.

28. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. to appear, Proceedings of the American Statistical Association. 2000

29. Rubin DB. Bayesian Inference for Causal Effects: The Role of Randomization. The Annals of Statistics 1978;6:34–58.

30. Rubin DB. Comment: Which Ifs Have Causal Answers. Journal of the American Statistical Association 1986;81:961–962.

31. Scharfstein DO, Rotnitzky A, Robins JM. Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models. Journal of the American Statistical Association 1999;94:1096–1120.

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]	Beginning of Every Semester
Academic Performance[1]	End of Grades 1 and 2
Social Skills[1]	Average for Grades 1 and 2
Site (4 Sites)[1]	Kindergarten
Cohort (2 Cohorts)[1]	Kindergarten

[1] Auxiliary Covariates (V_0 or V_t).

[2] Response (Y).

[3] Severity (S_t).

[4] Treatment Level (A_t).

Table 2

Table 2a: Model for Home Visiting Assignment Probabilities		
Fall Grade 2		
Covariate	Estimate	Stderr
Intercept 1	−1.45	0.36
Intercept 2	1.40	0.35
Cohort 1	−0.58	0.29
Site 1	−0.58	0.42
Site 2	−0.24	0.43
Site 3	0.42	0.36
HVPM16[1]	2.22	0.75
HVPM4[1]	−2.58	0.31

Table 2b: Model for Home Visiting Assignment Probabilities		
Spring Grade 2		
Covariate	Estimate	Stderr
Intercept 1	−2.61	0.47
Intercept 2	1.10	0.40
Cohort 1	−0.51	0.33
Site 1	1.64	0.49
Site 2	0.01	0.56
Site 3	1.76	0.45
HVPM16[1]	3.04	0.96
HVPM4[1]	−2.36	0.38
Past High Home Visiting Assignment[2]	2.18	.60
Past Low Home Visiting Assignment[2]	−1.55	0.39

Table 2c: Model for Home Visiting Assignment Probabilities		
Fall Grade 3		
Covariate	Estimate	Stderr
Intercept 1	−3.19	0.69
Intercept 2	3.00	0.60
Cohort 1	−0.58	0.38
Site 1	−1.26	0.55
Site 2	−0.75	0.62

Table 2c: Model for Home Visiting Assignment Probabilities		
Fall Grade 3		
Covariate	**Estimate**	**Stderr**
Site 3	1.23	0.48
HVPM16[1]	0.19	1.08
HVPM4[1]	−3.65	0.68
Past High Home Visiting Assignment[2]	2.98	.68
Past Low Home Visiting Assignment[2]	−1.31	0.41

Table 2d: Model for Home Visiting Assignment Probabilities		
Spring Grade 3		
Covariate	**Estimate**	**Stderr**
Intercept 1	−3.95	1.09
Intercept 2	3.35	0.61
Cohort 1	−0.89	0.43
Site 1	−0.82	0.55
Site 2	−0.56	0.63
Site 3	1.18	0.54
HVPM16[1]	4.67	1.61
HVPM4[1]	−2.75	0.47
Past High Home Visiting Assignment[2]	5.94	1.61
Past Low Home Visiting Assignment[2]	−2.53	0.43

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

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

[2] Past High Home Visiting Assignment if 1 if home visiting assignment in fall of 2nd Grade is 16; otherwise it is 0. Past Low Home Visiting Assignment is 1 if home visiting assignment in fall of 2nd grade is 4; otherwise it is 0.

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

[2] Past High Home Visiting Assignment if 1 if home visiting assignment in spring of 2nd Grade is 16; otherwise it is 0. Past Low Home Visiting Assignment is 1 if home visiting assignment in spring of 2nd grade is 4; otherwise it is 0.

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

[2] Past High Home Visiting Assignment if 1 if home visiting assignment in fall of 3rd Grade is 16; otherwise it is 0. Past Low Home Visiting Assignment is 1 if home visiting assignment in fall of 3rd grade is 4; otherwise it is 0.

Table 3

Table 3a: Model for g_1 through g_4		
School Behavior Problems		
Covariate	**Estimate[1]**	**Stderr**
Intercept	0.30	0.21
Cohort 1	−0.28	0.10
Site 1	0.43	0.14
Site 2	0.48	0.15
Site 3	0.45	0.12
Suspected Abuse	0.20	0.06
School Behavior Problems-Grade 1	0.26	0.06
Most Recent Academic Performance Score	0.09	0.03
Most Recent Social Skills Score	0.16	0.05

Table 3b: Model for g_1 through g_4		
School Behavior Change		
Covariate	**Estimate[1]**	**Stderr**
Intercept	−0.93	0.34
Site 1	−0.62	0.19
Site 2	−0.44	0.18
Site 3	−0.15	0.19
Suspected Abuse	−0.06	0.10
School Behavior Problems-Grade 1	0.18	0.10
Most Recent Academic Performance Score	0.05	0.03
Most Recent Social Skills Score	−0.15	0.07

[1] All variables are coded so that high values indicate poorer functioning.

[1] All variables are coded so that high values indicate poorer functioning.

Table 4

Table 4a: Estimated Treatment Effect on School Behavior Problems[1]		
h	**Estimated Treatment Effect[2]**	**Stderr**
ITT[3]	−0.23	0.07
0	−0.21	0.08
.2	−0.24	0.08
.4	−0.28	0.08
.6	−0.31	0.08
.8	−0.35	0.08
1	−0.40	0.09

Table 4b: Estimated Treatment Effect on School Behavior Change[1]		
h	**Estimated Treatment Effect[2]**	**Stderr**
ITT[3]	−0.33	0.09
0	−0.39	0.10
.2	−0.38	0.10
.4	−0.36	0.10
.6	−0.32	0.10
.8	−0.27	0.11
1	−0.19	0.13

[1] The estimated intercepts and the 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 social health profile; all weights are equal to 1.

[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 5

Table 5a: Model for Numerator of the Weights				
	Fall Grade 2		**Spring Grade 2**	
Covariate	Estimate	Stderr*	Estimate	Stderr[1]
Intercept 1	−1.69	0.23	−1.55	0.25
Intercept 2	1.06	0.20	1.27	0.24
HVPM16[2]	2.43	1.03	2.67	0.82
HVPM4[2]	−2.61	0.34	−2.69	0.34

Table 5b: Model for Numerator of the Weights				
	Fall Grade 3		**Spring Grade 3**	
Covariate	Estimate	Stderr*	Estimate	Stderr*
Intercept 1	−2.57	0.42	−2.87	0.45
Intercept 2	1.71	0.32	1.28	0.26
HVPM16**	1.95	1.22	3.79	0.93
HVPM4**	−3.00	0.36	−2.43	0.34

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

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

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

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