

Commentary on
‘Using inverse weighting and predictive inference to estimate
the effects of time-varying treatments on the discrete-time
hazard’

James M. Robins^{*,†}

*Departments of Epidemiology and Biostatistics, Harvard School of Public Health, 677 Huntington Avenue,
Boston, MA 02115-6096, U.S.A.*

1. INTRODUCTION

I would like to begin by thanking the editor Ralph D’Agostino for his recognition of the importance of the topic considered in Dawson and Lavori’s paper and for inviting this commentary. Dawson and Lavori (DL) propose several methods for estimating the effect of a monotone treatment pattern on the hazard function of a discrete failure time variable and compare their approaches to others previously proposed. In my opinion, DL’s methods in their Sections 1–3 are interesting and correct but are not the best available. However, I believe the methods DL propose in Sections 4 and 5 are biased without the addition of further, rather implausible, assumptions. Additionally, I feel that certain of DL comments indicate they are not fully aware of the motivation for and the range of the different methods proposed by myself and colleagues. My commentary is organized as follows. I first review DL’s methodological proposals in the order in which they appear in their paper and compare other proposals with theirs. In my final section, I review DL’s critique of methodologies proposed by myself and co-authors.

2. IGNORABLE TREATMENT ASSIGNMENT

DL’s assumption of sequential strongly ignorable treatment assignment is stronger than is needed to obtain the results in their Sections 1–3. Rather, given that treatment Z_t is monotone (that is, subjects off treatment never restart) and Y_t is a survival indicator, it is sufficient to impose the weaker assumption that for all $t, t' > t$ and M in the set $\{c, d_s, s = t, t + 1, \dots\}$ of treatment regimes

$$\text{pr}(Z_t = 1 \mid Z_{t-1} = 1, \bar{X}_t, Y_t = 0, Y_{M'}) = \text{pr}(Z_t = 1 \mid Z_{t-1} = 1, \bar{X}_t, Y_t = 0) > 0$$

where $\bar{X}_t = (X_0, X_1, \dots, X_t)$. This assumption states that among subjects surviving to t who have remained on treatment through $t - 1$ ($Z_{t-1} = 1$), the conditional probability of staying on

*Correspondence to: James M. Robins, Department of Epidemiology, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115-6096, U.S.A.

†E-mail: robins@hsph.harvard.edu

treatment at t given past covariate history \bar{X}_t does not further depend on the future potential (that is, counterfactual) survival outcomes Y_{M^t} for any treatment regime M that specifies continuous treatment at least through time $t-1$. I have referred to this assumption as the assumption of sequential randomization with respect to counterfactual survival or, equivalently, as the assumption of no unmeasured confounders for potential survival outcomes [1, 2]. This assumption is weaker than DL's sequential ignorability assumption in that it does not require the different regimes M to be considered jointly. More importantly, in contrast to DL's assumption, it does not require that there be no unmeasured confounders for the counterfactual time dependent covariate outcomes \bar{X}_t . In references [1] and [2], I consider substantive settings in which the assumption of no unmeasured confounders holds only for potential survival outcomes (so DL's ignorability assumption fails). I show that in such settings the effect of treatment on potential survival outcomes is still given by the G -computation algorithm formula estimand. This implies that the results obtained by DL in Sections 1–3 hold under this weaker assumption, as their results follow mathematically from the G -computation algorithm formula. For example, with c denoting the continuous treatment regime, their estimator (2) is converging to $\lambda_c(t) = \text{pr}(Y_{c(t+1)} = 1 \mid Y_{ct} = 0)$ which is the ratio of the G -computation estimand for $\text{pr}(Y_{c(t+1)} = 1, Y_{ct} = 0)$ to that for $\text{pr}(Y_{ct} = 0)$. Note the fact that DL's estimand of interest is based on the G -computation algorithm formula does not imply that the most robust way to estimate this estimand is with the plug-in G -computation algorithm estimator. Indeed, beginning with my original derivation of the G -computation algorithm estimand in 1986 [3], I have stressed that the G -computation algorithm estimator is rarely a robust estimation method. DL reiterate this point in their Section 3.3 and Appendix B. All the methods co-workers and I have developed since 1986 were motivated by the need to find more robust estimators.

3. DOUBLY ROBUST ESTIMATION OF THE HAZARD UNDER CONTINUOUS TREATMENT

In this section, we argue that, except when the sample size is very small, one should use one of the doubly robust locally efficient estimators initially proposed by colleagues and myself [4–7] instead of the estimator proposed by DL in their Section 3. A thorough discussion of double robustness can be found in reference [35].

3.1. The special case $t = 1$

To focus on the central issue, it will be useful to first consider the simplest case of DL's set-up in which $t = 1$. Specifically, let X_0 denote the column vector of baseline covariates, Z_0 denote the indicator of treatment at time 0, and Y_0 and Y_1 denote the survival indicators at times 0 and 1 ($Y_j = 1$ if a failure and $Y_j = 0$ if a survivor). Let Y_{c0} and Y_{c1} be a subject's survival indicator at times 0 and 1 had, possibly contrary to fact, a subject been treated at time 0. Our goal is to estimate the discrete hazard $\lambda_c = \text{pr}[Y_{c1} = 1 \mid Y_{c0} = 0]$ at time 1 under treatment. Because all subjects are survivors at time 0, $Y_{c0} = Y_0 = 0$ with probability 1, so that our estimand is $\lambda_c = \text{pr}[Y_{c1} = 1]$. Under DL's ignorability assumption that $Y_{c1} \perp\!\!\!\perp Z_0 \mid X_0$, we have that $\text{pr}[Y_{c1} = 1 \mid X_0] = \text{pr}[Y_1 = 1 \mid X_0, Z_0 = 1]$. Therefore, $\lambda_c = \text{pr}[Y_{c1} = 1]$ is given by the G -computation algorithm formula estimand $\lambda_c = \int \text{pr}[Y_1 = 1 \mid X_0, Z_0 = 1] dF(X_0)$ or equivalently by the inverse probability of treatment weighting (IPTW) formula $\lambda_c = E[Z_0 Y_1 / e(X_0)]$ where $e(X_0) = \text{pr}[Z_0 = 1 \mid X_0]$ is the propensity score. In Section 3.1, DL suggest fitting the

model $\text{logit pr}[Y_1 = 1 | X_0, Z_0 = 1] = \alpha X_0$ and then estimating λ_c with the regression estimator $\hat{\lambda}_{\text{reg}} = n^{-1} \sum_i \text{expit}\{\hat{\alpha} X_{0i}\}$ where i will always indexes the n study subjects, $\text{expit}(x) = e^x / (1 + e^x)$, $\hat{\alpha}$ is the maximum likelihood estimator (MLE) of α among subjects with $Z_0 = 1$, and 1 is a component of X_0 so that αX_0 includes an intercept term. Note that in the special case in which $t = 1$, DL's regression estimator does not require an estimate of propensity score because the conditional law of X_0 given $Z_0 = 1$ is equal to its marginal law. In contrast, Rosenbaum [8], Robins and Rotnitzky [9], and Heyting *et al.* [10] specified a propensity score model $\text{pr}[Z_0 = 1 | X_0] = e(X_0; \gamma)$ with $e(X_0; \gamma) = \text{expit}(\gamma X_0)$ and then estimated λ_c with the IPTW estimator $\hat{\lambda}_{\text{IPTW}} = \{\sum_i Z_{0i} Y_{1i} / e(X_{0i}; \hat{\gamma})\} / \{\sum_i Z_{0i} / e(X_{0i}; \hat{\gamma})\}$ where $\hat{\gamma}$ is the MLE of γ . In their Section 3, DL also mention the IPTW estimator. Now, in general, $\hat{\lambda}_{\text{IPTW}}$ will be consistent for λ_c only if the propensity model $e(X_0; \gamma)$ is correct, while $\hat{\lambda}_{\text{reg}}$ will be consistent only if the regression model $\text{expit}\{\alpha X_0\}$ is correct. When both models are correct, Robins and Wang [11] point out that $\hat{\lambda}_{\text{reg}}$ has smaller asymptotic variance than $\hat{\lambda}_{\text{IPTW}}$ and will have better sampling properties in very small samples. DL stress this latter point. However, since with observational data one can never be sure that either the propensity model or the regression model is correct, the best that can be hoped for is to find an estimator that is consistent asymptotically normal (CAN) for λ_c if either (but not necessarily both) of the models is correct. Robins *et al.* [6] refer to an estimator $\hat{\lambda}_{\text{dr}}$ with this property as doubly robust or doubly protected. In the following subsection and Appendix D, we review arguments in references [4–7] to show that the following estimators are doubly robust: (i) the augmented regression estimator $\hat{\lambda}_{\text{dr, aug}} = \hat{\lambda}_{\text{reg}} + \sum_i Z_{0i} (Y_i - \text{expit}(\hat{\alpha} X_{0i})) / e(X_{0i}; \hat{\gamma}) / \{\sum_i Z_{0i} / e(X_{0i}; \hat{\gamma})\}$; (ii) the doubly robust regression estimator $\hat{\lambda}_{\text{reg, dr}} = n^{-1} \sum_i \text{expit}(\tilde{\alpha} X_{0i} + \tilde{\kappa} / e(X_{0i}; \hat{\gamma}))$ where $(\tilde{\alpha}, \tilde{\kappa})$ is the MLE of (α, κ) in the extended regression model $\text{logit pr}[Y_1 = 1 | X_0, Z_0 = 1] = \alpha X_0 + \kappa e(X_0; \hat{\gamma})^{-1}$ that adds the 'robustifying covariate' $e(X_0; \hat{\gamma})^{-1}$ to the regression model αX_0 ; and (iii) the iterated IPTW estimator $\hat{\lambda}_{\text{IPTW, dr}} = \{\sum_i Z_{0i} Y_{1i} / e(X_{0i}; \tilde{\gamma}, \tilde{\theta})\} / \{\sum_i Z_{0i} / e(X_{0i}; \tilde{\gamma}, \tilde{\theta})\}$ where $(\tilde{\gamma}, \tilde{\theta})$ are the values at convergence (that is, as $k \rightarrow \infty$) of the MLEs $(\tilde{\gamma}_k, \tilde{\theta}_k)$ of (γ_k, θ_k) in the extended model $\text{pr}[Z_0 = 1 | X_0] = e(X_0; \gamma_k, \theta_k)$ with $e(X_0; \gamma_k, \theta_k) = \text{expit}(\gamma_k X_0 + \theta_k \{\text{expit}(\hat{\alpha} X_{0i}) - \hat{\lambda}_{\text{reg}}\} / e(X_0; \tilde{\gamma}_{k-1}, \tilde{\theta}_{k-1}))$ that adds at iteration k the 'robustifying covariate' $\{\text{expit}(\hat{\alpha} X_{0i}) - \hat{\lambda}_{\text{reg}}\} / e(X_0; \tilde{\gamma}_{k-1}, \tilde{\theta}_{k-1})$ to the propensity model $e(X_0; \gamma) = \text{expit}(\gamma X_0)$. The iteration is initialized by taking $e(X_0; \tilde{\gamma}_1, \tilde{\theta}_1)$ equal to $e(X_0; \hat{\gamma})$.

All doubly robust estimators are also locally semi-parametric efficient when both the regression model $\text{expit}(\alpha X_0)$ and the propensity model $e(X_0; \gamma)$ are correctly specified. That is, when both models are correct, the estimators are equally efficient and no doubly robust can have a smaller asymptotic variance. Now, one might view the double robustness property as being of only academic interest because, in reality, all models are misspecified. However, even when both the propensity and the regression models are misspecified, the bias of the doubly robust estimators (i) and (ii) will often be less than that of $\hat{\lambda}_{\text{IPTW}}$ or of DL's estimator $\hat{\lambda}_{\text{reg}}$. To see why, let γ^* and α^* be the probability limits of the MLEs $\hat{\gamma}$ and $\hat{\alpha}$ from the unextended propensity and regression models. Define the theoretical quantities $h_e(X_0)$ and $h_{\text{reg}}(X_0)$ by $h_e(X_0) = e(X_0) / e(X_0; \gamma^*) - 1$ and $h_{\text{reg}}(X_0) = \text{pr}[Y_1 = 1 | X_0, Z_0 = 1] - \text{expit}(\alpha^* X_0)$, and let $\sigma_e = \sup |h_e(x_0)|$ and $\sigma_{\text{reg}} = \sup |h_{\text{reg}}(x_0)|$ with the supremum taken over the support of X_0 . Letting λ_c^* be the probability limit of $\hat{\lambda}$, the large sample absolute bias $|\lambda_{\text{dr}}^* - \lambda_c|$ of any doubly robust estimator is at most order $\sigma_e \sigma_{\text{reg}}$. In contrast, the large sample biases $|\lambda_{\text{IPTW}}^* - \lambda_c|$ of

$\hat{\lambda}_{IPTW}$ and $|\lambda_{reg}^* - \lambda_c|$ of $\hat{\lambda}_{reg}$ are generally of order σ_e and σ_{reg} , respectively. Therefore, when both the propensity and regression models are approximately correct in the sense that that σ_e and σ_{reg} are small, the large sample bias of the doubly robust estimators will often be smaller than that of the IPTW or regression estimator. When one model is grossly wrong and the other approximately correct (but it is not known which is the nearly correct model), only the use of the doubly robust estimator can guarantee small bias. Finally, the doubly-robust estimators (i) and (ii) above have biases that are bounded even when both the propensity and outcome regression models are grossly incorrect.

3.2. Arbitrary t

In the case of arbitrary t , results of Robins [12] and Scharfstein *et al.* [4] can be used to derive doubly robust estimators of the hazard $\lambda_c(t)$ under continuous treatment. In the context of estimation of the cumulative (integrated hazard) of a continuous failure time variable, Robins [12] derived the analogue of $\hat{\lambda}_{dr, aug}$, and Robins *et al.* [5] derived the analogue of $\hat{\lambda}_{IPTW, dr}$. Here we provide the analogue of $\hat{\lambda}_{reg, dr}$. In Appendix D, we prove the double robustness of all three estimators. (Before beginning, we remark that for purposes of causal inference, one should generally make regime-specific survivor functions or cumulative hazards the ultimate targets of inference, rather than hazards or hazard ratios at particular times. This reflects the well known fact that the potential survival time of every subject could be greater under a regime **a** than under a regime **b** and yet, at certain times t , the hazard of failure under regime **a** can exceed that under regime **b**.)

Let $G_{(t+1)0} = Y_{t+1}$ and $G_{(t+1)1} = I(Y_t = 0)$. For $m = t, t - 1, \dots, 0$ and $j = 0, 1$, let $G_{mj} = E[G_{(m+1)j} | \bar{X}_m, Y_m = 0, Z_m = 1]I(Y_m = 0)$. Thus $G_{t0} = \lambda(t | \bar{X}_t, Z_t = 1)I(Y_t = 0)$ and $G_{t1} = I(Y_t = 0)$. Under DL's sequential ignorability assumption, $G_{m0} = \lambda_c(t | \bar{X}_m)pr[Y_{ct} = 0 | \bar{X}_m, Y_m = 0]I(Y_m = 0)$ and $G_{m1} = pr[Y_{ct} = 0 | \bar{X}_m, Y_m = 0]I(Y_m = 0)$. In particular note that $E[G_{00}]/E[G_{01}] = E\{pr[Y_{c(t+1)} = 1, Y_{ct} = 0 | X_0]\}/E\{pr[Y_{ct} = 0 | X_0]\} = pr[Y_{c(t+1)} = 1, Y_{ct} = 0]/pr[Y_{ct} = 0] = \lambda_c(t)$. Our goal will be to estimate $E[G_{00}]/E[G_{01}]$, as $E[G_{00}]/E[G_{01}] = \lambda_c(t)$ even under the weaker assumption [1] of no unmeasured confounders for potential survival outcomes.

Given an estimate $\hat{G}_{(m+1)j}$ of $G_{(m+1)j}$ for $j = 0, 1$, we estimate G_{mj} by specifying a model $\text{expit}(\alpha_{mj}W_{mj} + \kappa_{mj}\hat{e}_m^{-1})$ for the conditional expectation of $\hat{G}_{(m+1)j}$ given \bar{X}_m among subjects with $Y_m = 0$ and $Z_m = 1$ where W_{mj} is a user-supplied vector function of the covariates \bar{X}_m through time m and $\hat{e}_m = \hat{e}_{m|m-1} \times \hat{e}_{m-1|m-2} \times \dots \times \hat{e}_{1|0} \times \hat{e}_{0|-1}$ is a model-based estimate of the propensity score product through time m where $e_{m|m-1} = pr[Z_m = 1 | Z_{m-1} = 1, \bar{X}_m, Y_m = 0]$, $Z_{-1} \equiv 1$. We obtain estimates $(\hat{\alpha}_{mj}, \hat{\kappa}_{mj})$ as the solution to the following linear logistic score equation for the model $\text{logit } E[\hat{G}_{(m+1)j} | \bar{X}_m, Z_m = 1, Y_m = 0] = \text{expit}(\alpha_{mj}W_{mj} + \kappa_{mj}\hat{e}_m^{-1})$

$$0 = \sum_i I(Y_{m,i} = 0)I(Z_{m,i} = 1)(\hat{G}_{(m+1)j,i} - \text{expit}(\alpha_{mj}W_{mj,i} + \kappa_{mj}\hat{e}_{m,i}^{-1}))(W_{mj,i}, \hat{e}_{m,i}^{-1})^T$$

and set $\hat{G}_{mj} \equiv \hat{G}_{mj}(\hat{\alpha}_{mj}, \hat{\kappa}_{mj}) = \text{expit}(\hat{\alpha}_{mj}W_{mj} + \hat{\kappa}_{mj}/\hat{e}_m)I(Y_m = 0)$. Even though $\hat{G}_{(m+1)j,i}$ can take on values intermediate between zero and one, a standard binomial regression software program can be tricked into solving this score equation by multiplying the $\hat{G}_{(m+1)j,i}$ by 1000, rounding to the nearest integer, treating the inflated counts as the number of binomial successes in 1000 trials, and specifying the canonical logistic link. The above recursion is initialized by setting $\hat{G}_{(t+1)0} = G_{(t+1)0} = Y_{t+1}$ and $\hat{G}_{(t+1)1} = G_{(t+1)1} = I(Y_t = 0)$. (Note that this implies that

$\hat{G}_{t1} = G_{t1} = I(Y_t = 0)$ as well.) Finally, $\hat{\lambda}_c(t) = \sum_i \hat{G}_{00,i} / \sum_i \hat{G}_{01,i}$. Robins *et al.* [13] and Robins [14] refer to recursive regression estimators like $\hat{\lambda}_c(t)$ as (ratios of) iterated conditional expectation (ICE) estimators. Now $\hat{\lambda}_c(t)$ is CAN when the models $\text{expit}(\alpha_{mj}W_{mj} + \kappa_{mj}\hat{e}_m^{-1})$ for the conditional expectation of the $G_{(m+1)j}$ are correct because then $\sum_i \hat{G}_{00,i} / \sum_i \hat{G}_{01,i}$ is converging to $E[G_{00}] / E[G_{01}] = \lambda_c(t)$. In Appendix D we prove that $\hat{\lambda}_c(t)$ is doubly robust by showing it will be CAN if the model for the propensity score used to calculate \hat{e}_m is correct even if the models $\text{expit}(\alpha_{mj}W_{mj} + \kappa_{mj}\hat{e}_m^{-1})$ are incorrect. We also provide doubly robust estimators of augmented regression estimator form and of the IPTW form as well.

In contrast to the doubly robust ICE estimator $\hat{\lambda}_c(t)$, DL's estimator requires not only a correct model for the propensity scores but also a correct model $\text{expit}(\alpha_{t0}W_{t0})$ for $E[G_{(t+1)0} | \bar{X}_t, Y_t = 0, Z_t = 1] = \lambda[t | \bar{X}_t, Z_t = 1]$.

I wish to point out that Scharfstein *et al.* [4] have shown that, in contrast to propensity score models, it is not in general possible to specify mutually compatible logistic models $\text{expit}(\alpha_{mj}W_{mj})$ or $\text{expit}(\alpha_{mj}W_{mj} + \kappa_{mj}\hat{e}_m^{-1})$ for the conditional expectations $G_{(m+1)j}$. Specifically, they show there is no joint distribution for which all such models are correct with the α_{mj} non-zero. However, although incorrect, the models, if sufficiently high dimensional, can still approximate the conditional expectations $\hat{G}_{(m+1)j}$ closely, in which case, as argued earlier, we would expect our doubly robust estimators to generally have smaller large sample bias than their competitors, provided flexible models for the propensity score are used as well. Now with a good deal more effort it is possible to specify and fit mutually compatible models for the conditional expectations $\hat{G}_{(m+1)j}$; however, we doubt the extra effort will significantly improve on the performance obtained with mutually incompatible logistic models, provided these latter models are richly parameterized [35].

3.3. Improved inverse weighting

We next consider DL's claim in their Section 3.2 that the small sample performance of their estimator can be improved by replacing the covariate X_t in their discrete hazard model by an estimate of the propensity score $e_{t|t-1} = \text{pr}[Z_t = 1 | \bar{X}_t, Z(t-1) = 1, Y_t = 0]$. Returning to the special case that $t=1$, DL's claim is that the performance of $\hat{\mu}_{\text{reg}}$ is inferior to that of the estimator $\hat{\mu}_{\text{reg},e} = \sum_i \text{expit}(\hat{\alpha}_0 + \hat{\alpha}_1 e(X_{0i}; \hat{\gamma}))$ where $\hat{\alpha}_0$ and $\hat{\alpha}_1$ are the MLEs of α_0 and α_1 in the regression model $\text{logit pr}[Y_1 = 1 | e(X_0; \hat{\gamma}), Z_0 = 1] = \alpha_0 + \alpha_1 e(X_0; \hat{\gamma})$. We disagree with this claim for the following reasons. First, we believe that the arguments that DL offer in their Section 3.2 and in the first third of their Appendix A fail to provide any support for, much less prove, their claim. Second, and more importantly, their estimator $\hat{\mu}_{\text{reg},e}$ will only be consistent for λ_c if both the propensity score model $e(X_0; \gamma)$ and the linear logistic regression model of Y_1 on $(1, e(X_0; \hat{\gamma}))$ are both correct, even in the special case $t=1$. Thus, in our opinion, the estimator $\hat{\mu}_{\text{reg},e}$ should never be used, as its less robust than even their original estimator $\hat{\mu}_{\text{reg}}$, much less a doubly robust estimator $\hat{\mu}_{\text{dr}}$.

4. ESTIMATING THE EFFECTS OF DURATION OF TREATMENT

In the first part of Section 4, DL propose a method for estimating the effect of treatment duration. We shall argue that it is preferable to estimate duration effects by either IPTW

estimation of a marginal structural logistic discrete hazard model [15–17] or G -estimation of a structural nested failure time model [18, 19].

4.1. Marginal structural models

The problem with DL's approach is that it does not contain regression parameters that serve to quantify biologically relevant causal effects. As a result, it is difficult to test important causal hypotheses. For example, suppose we wanted to test the causal null hypothesis of no effect of treatment on survival. That is, the hypothesis that the hazard $\lambda_c(t) = \text{pr}(Y_{c(t+1)} = 1 \mid Y_{ct} = 0)$ had all subjects remained on treatment past t is equal to the hazard $\lambda_{d_s}(t)$ had all subjects stopped treatment at time s for $s = 0, 1, \dots, t, t = 0, 1, \dots, T$ where $T+1$ is the maximum follow-up time. Then to estimate $\lambda_{d_s}(t)$, DL in their equation (4) propose fitting separately for $s = 0, 1, \dots, t$, hazard models for the observed hazard at t among those who stopped treatment at time s of the form $\text{logit } \lambda(t \mid \bar{X}_s, Z_s = 0, Z_{s-1} = 1) = f(t)^\top \alpha_s + X_s \gamma_s + X_0 \delta_s$ where we have indexed the parameters by s to indicate that they will generally vary with s , and, for expositional convenience, following DL, we have assumed the hazard depends on the entire covariate history only through X_s and X_0 . To estimate $\lambda_c(t)$ they proposed fitting $\text{logit } \lambda(t \mid \bar{X}_t, Z_t = 1) = f(t)^\top \alpha_t + X_t \gamma_t + X_0 \delta_t$ to subjects who remained on treatment through t . Note the null hypothesis of no treatment effect on the hazard at t does not imply that the γ_k , the α_k or the δ_k are equal for $k = 0, \dots, t^3$. Rather after inverse probability weighting by the appropriate function of the propensity scores, DL obtain estimates $\hat{\lambda}_{c,t}(t), \hat{\lambda}_{d_s,t}(t)$ $s = 0, 1, \dots, t, t = 0, 1, \dots, T$. To test the null hypothesis of no treatment effect we must test the hypothesis that the $t+2$ estimates at each t are estimating the same number. To do so we must extend DL's approach by introducing models for the marginal counterfactual hazards. Then a test of the causal null hypothesis is a $\sum_{t=0}^T (t+1) = (T^2 + 3T + 1)/2$ degree of freedom test of the hypothesis $\theta_{1s} = 0, s = 0, 1, \dots, t, t = 0, 1, 2, \dots, T$ in the model

$$\text{logit } \lambda_M(t) = \theta_{0t} + \theta_{1s} I(s \leq t)$$

for M in the set $\{d_s; s = 0, 1, \dots, T, T+1\}$ where d_{T+1} is the continuous treatment regime c . Under this model, for s greater than t , $\text{logit } \lambda_{d_s}(t) = \text{logit } \lambda_c(t) = \theta_{0t}$ since stopping treatment at time s in the future does not affect the hazard at the earlier time t . Clearly such a test will have little power. To increase power, we would specify a submodel with fewer parameters – an extreme case being

$$\text{logit } \lambda_M(t) = \theta_{0t} + \theta_1(1 + t - s)I(s \leq t)$$

which says that the odds ratio e^{θ_1} at time t comparing $\lambda_{d_s}(t)$ and $\lambda_c(t)$ is simply a linear logistic function of the time $(1 + t - s)$ since stopping treatment. Then a one degree of freedom test of $\theta_1 = 0$ is a test of the causal null hypothesis. To carry out this 1 d.o.f. test one would generally treat DL's estimates $\text{logit } \hat{\lambda}_c(t)$ and $\text{logit } \hat{\lambda}_{d_s}(t)$ as multivariate normal with mean $\text{logit } \lambda_c(t)$ and $\text{logit } \lambda_{d_s}(t)$ and covariance matrix that would have to be estimated by the bootstrap. However, if, as would certainly be the case for large T , there are very few subjects who followed regime d_s for many s , the corresponding estimates $\text{logit } \hat{\lambda}_{d_s}(t)$ would fail to have an approximate normal sampling distribution and, as a consequence, the test would have the wrong level. Further, there might be a concern that the effect of treatment might be modified by a continuous pretreatment variable such as depression score X_0^* , a component of

X_0 , so that the optimal treatment might differ for those with different values of the depression score. Thus one might wish to fit the model

$$\text{logit } \lambda_M(t | X_0^*) = \theta_{0t} + \theta_1(1 + t - s)I(s \leq t) + \theta_2 X_0^* + \theta_3 X_0^*(1 + t - s)I(s \leq t)$$

with θ_3 quantifying the interaction. However, this is not possible using DL's methodology. All of these difficulties are solved by using IPTW estimation of marginal structural models (MSMs).

Indeed, all three of the above models are examples of discrete hazard MSMs [15–17] as they model the marginal hazard of the counterfactual hazards $\lambda_{d_s}(t)$ and $\lambda_c(t)$ (possibly within levels of baseline covariates) rather than the hazard of the observed failure time variable given past treatment history and baseline covariates. The first model is a saturated (non-parametric) MSM as it places no *a priori* restrictions on the values of $\text{logit } \lambda_{d_s}(t)$ or $\text{logit } \lambda_c(t)$. Under the assumption of no unmeasured confounders, the parameters of any discrete hazard MSM, can be fit by IPTW. For example, to fit the second of the MSMs above, one fits the discrete hazard model

$$\text{logit } \lambda(t | \bar{Z}_t = s) = \theta_{0t} + \theta_1(1 + t - s)I(s \leq t)$$

to the observed data by pooled weighted logistic regression where each unit of person-time contributes a separate observation; if $s \leq t$, $\bar{Z}_t = s$ means a subject first failed to take treatment at time s , \bar{Z}_t is recorded as $t + 1$ if a subject continues on treatment through t , and the weight given to a person at time t is $\text{SW}_t = \{\prod_{j=0}^{s-1} \hat{e}_{j|j-1}^* (1 - \hat{e}_{s|s-1}^*)^{I(s \leq t)}\} / \{\prod_{j=0}^{s-1} \hat{e}_{j|j-1} (1 - \hat{e}_{s|s-1})^{I(s \leq t)}\}$ where we set $s = t + 1$ if a subject stayed on treatment through time t , and $\hat{e}_{j|j-1}^*$ is the empirical probability of remaining on treatment at time j among surviving subjects who were on treatment at time $j - 1$ (without correcting for past covariate history \bar{X}_j). Informally, the denominator of SW_t is the probability that subject had the treatment history through time t that they did indeed have, given past treatment and covariate history. The resulting IPTW estimators of the parameters θ are CAN if the propensity model used to calculate the $\hat{e}_{j|j-1}$ is correct. Doubly robust IPTW estimators of regression form (that is, ICE), augmented regression form, and IPTW form that are analogous to those described earlier are available [7].

4.2. Dynamic treatment regimes and G-estimation of structural nested failure-time models

In this section, we continue to follow DL and restrict consideration to monotone treatments. However, in contrast to DL, we shall consider dynamic treatment regimes. A dynamic regime is one in which a subject's covariate history \bar{X}_s through time s determines whether treatment should be stopped at time s . To select an optimal treatment strategy for a patient, it is usually necessary to consider dynamic treatment regimes. For example, the optimal regime will be dynamic when the treatment can cause serious toxicity, because the optimal strategy must stop the treatment when toxicity develops. Even in the absence of toxicity, the optimal regime may be dynamic. Indeed, the optimal regime will be dynamic whenever the effect of treatment at s is qualitatively modified by time-varying covariates \bar{X}_s . For example, in DL's depression scenario, if antidepressant treatment at s was of benefit to those who have experienced symptoms of depression in the past three weeks, but was harmful (due to the possibility of side-effects) to those without symptoms for three weeks, then the optimal regime might be to continue an antidepressant until the patient has been symptom-free for three weeks. Robins [21] showed that under DL's sequential ignorability assumption IPTW estimators could

be used to estimate the effect of a dynamic regime: one simply defines a patient to be censored at the time m at which the treatment received by the patient first differs from that prescribed under the particular regime. The weight ascribed to a patient at time t who had followed the regime up to t and who stopped treatment at time $m < t$ (because they had been symptom-free for three weeks at m) would be the the inverse of the product of their propensity scores $e_{k|k-1}$ for $k=0, \dots, m-1$ times the probability $1-e_{m|m-1}$ they did indeed stop treatment at time m . As there are many more dynamic treatment regimes than the $T+2$ non-dynamic treatment regimes, $c, d_s, s=0, \dots, T$, it will often be the case that very few patients have followed a particular dynamic treatment regime of interest. Thus, the problem of inadequate sample size will be more severe for dynamic than non-dynamic regimes. As a consequence, some dimension reducing regression approach (analogous to the role of MSMs for non-dynamic regimes) is necessary. Structural nested failure-time models (SNFTMs) [18, 19, 22] were developed in part for this purpose. SNFTMs allow one to directly model the interaction between treatment at each time s and past time-dependent covariate history. As a consequence, SNFTMs can be used to estimate the effect of dynamic treatment regimes. In contrast, the modelling of the interaction between treatment and time-dependent covariates and the estimation of the effect of dynamic regimes with MSMs is computationally complex and beset with practical difficulties. In Reference [37], I describe how one might approach directly estimating the optimal treatment regime using an extension of SNFTMs called optimal-regime SNFTMs. Development of these optimal-regime models was inspired by Murphy's ground-breaking article [38] on estimation of optimal treatment regimes.

We will define SNFTMs in continuous time [19]. Let $U(s)$ be the potential failure time if a subject had followed his observed treatment up to s and took no treatment thereafter, where s can now be any non-negative real number. Thus if, in reality, a subject stopped treatment at $t > s$, $U(s)$ is the potential survival time U_{d_s} associated with the regime d_s . If, in reality, a subject stopped treatment at time s' for $s' \leq s$, $U(s)$ is the potential survival time $U_{d_{s'}}$ associated with the regime $d_{s'}$. Let T be a subject's observed failure time. Note that $U(T) = T$ since the time $U(T)$ a subject would fail if he followed his observed treatment history up to his actual failure time T is his actual failure time.

4.2.1. Rank preserving SNFTMs. We first define locally rank preserving (RP) SNFTMs. The assumption of local rank preservation states that if $U(s) > s$, the derivative $\dot{U}(s) = dU(s)/ds$ exists and is finite and equals $d(U(s), \bar{X}_s)Z_s$ where $d(u, \bar{x}_s)$ is an unknown function to be estimated. That is, if $U(s) > s$ then, for infinitesimal positive Δs

$$U(s + \Delta s) - U(s) = Z_s d(U(s), \bar{X}_s) \Delta s$$

The LHS of this equation is the additive increment in survival time attributable to a subject's actual treatment in the interval $[s, s + \Delta s)$. Thus the equation states that for a subject on treatment (so $Z_s = 1$), the additional increment $d(U(s), \bar{X}_s) \Delta s$ is a deterministic function $d(U(s), \bar{X}_s) \Delta s$ of \bar{X}_s and $U(s)$. We will refer to $d(U(s), \bar{X}_s)$ as a rank preserving (RP) instantaneous blip function. If a subject is off treatment at s , obviously the effect $U(s + \Delta s) - U(s)$ of his treatment at s is 0.

To help understand the meaning of the instantaneous blip function, consider as an example a setting in which (i) failure is death from a particular infectious disease, (ii) if death from this disease occurs, it always occurs within five weeks from the time of initial unrecorded subclinical infection, and (iii) treatment at s is a preventive antibiotic which is never harmful

but is of no benefit unless the study subject is already infected by s . The restriction on the instantaneous blip function $d(U(s), \bar{X}_s)$ implied by the biological knowledge that treatment at time s is only harmful or beneficial to those who are already infected; that is, to those who were destined to fail by $s + 5$ if they had received no further treatment (that is, to those with $U(s) < s + 5$) is that $d(U(s), \bar{X}_s) = 0$ if $U(s) - s > 5$. A instantaneous rate RP SNFTM is a parametric model $d(U(s), \bar{X}_s, \psi)$ for the unknown RP blip function $d(U(s), \bar{X}_s)$ satisfying $d(U(s), \bar{X}_s, 0) = 0$ so $\psi = 0$ represents the null hypothesis of no treatment effect.

We next describe through an example how to estimate the parameters of a RP SNFTM under the assumption of no unmeasured confounders for the potential survival time. The key is to calculate a subject's failure time $U \equiv U(0)$ when never treated which we will see can be obtained by solving a differential equation. Consider the instantaneous rate RP SNFTM $d(U(s), \bar{X}_s, \psi) = 1 - \exp\{\psi_1 + \psi_2 X_s^*\}$ where X_s^* is the most recently recorded level of depressive symptoms prior to s . This model states that the effect of a final instantaneous brief blip of treatment of length Δs is to add or subtract $[1 - \exp\{\psi_1 + \psi_2 X_s^*\}] \Delta s$ to a subject's survival time regardless of $U(s)$. Hence $\psi_1 = \psi_2 = 0$ implies no effect of treatment on survival and ψ_2 will be negative if the effect of treatment is more beneficial for those currently depressed. Write $D\{U(s), s\}$ as shorthand for $Z_s d(U(s), \bar{X}_s, \psi^*)$ where ψ^* is the true value of ψ and consider the solution $U(s)$ to the differential equation

$$dU(s)/ds = D\{U(s), s\}$$

through the known point known $(T, U(T)) = (T, T)$ corresponding to a subject's actual failure time. We obtain $U(s) = s + \int_s^T \exp\{Z_u(\psi_1^* + \psi_2^* X_u^*)\} du$, where the integration is over the u argument in X_u^* and Z_u . Hence $U \equiv U(0) = \int_0^T \exp\{Z_u(\psi_1^* + \psi_2^* X_u^*)\} du$.

If every subject's failure time T is observed, a G -estimate of $\psi^* = (\psi_1^*, \psi_2^*)$ is obtained as follows. Define $H(\psi) = \int_0^T \exp\{Z_u(\psi_1 + \psi_2 X_u^*)\} du$ so $U = H(\psi^*)$. Separately, for each value of the two-vector ψ on a two-dimensional grid, add $\theta_1 H(\psi) + \theta_2 H(\psi) X_j^*$ to logistic models $\text{expit}(\gamma W_j)$ for the propensity score $e_{j|j-1} = \text{pr}[Z_j = 1 | Z_{j-1} = 1, \bar{X}_j, Y_j = 0]$ at j to obtain the model $\text{expit}(\gamma W_j + \theta_1 H(\psi) + \theta_2 H(\psi) X_j^*)$ for $j = 0, 1, \dots, T$. Because under the assumption of no unmeasured confounders $U = H(\psi^*)$ is independent of Z_j given $(Z_{j-1} = 1, \bar{X}_j, Y_j = 0)$, the values of ψ for which a χ_2^2 logistic regression score test of the joint hypothesis $\theta_1 = \theta_2 = 0$ does not reject form a large sample 95 per cent confidence interval for ψ^* . The value of ψ where the χ_2^2 test statistic is zero serves as a point estimate, referred to as a G -estimate. The parameter ψ is treated as fixed in performing the χ_2^2 test. References [18] and [21] describe how to modify this procedure when some failure times are missing due to censoring and how to use the G -estimates of ψ^* to obtain estimates of the survival distribution under both dynamic and non-dynamic regimes. In reference [7], methods for obtaining doubly robust G -estimates are described.

4.2.2. SNFTMs without rank preservation. Consider two subjects, say i and j , who have identical survival times and covariate and treatment histories. Under the assumption of local rank preservation, the two subjects would have identical survival times U if treatment had been withheld because the two subjects satisfy the same U -defining differential equation. This assumption is biologically implausible. To see why, again consider the infectious disease example. Suppose treatment has a beneficial biological effect; subjects i and j are both infected and on treatment at time s , subject i fails to absorb his/her dose at due to

gastrointestinal difficulties, while subject j successfully absorbs his/her dose. Then if $T_i = T_j$, we would expect $U_i = T_i = T_j > U_j$ since subject j but not subject i experiences the benefit of treatment. Dependence of the magnitude of the treatment effect on unmeasured factors such as bioabsorption and genetic endowment is the rule.

Suppose, for infinitesimal positive Δs , conditional on \bar{X}_s and $Z_s = 1$, that u is, say, the conditional p th quantile of (the distribution of) $U(s)$. We then define the non-RP blip function $d(u, \bar{X}_s)$ by the requirement that $d(u, \bar{X}_s)Z_s\Delta s + u$ is the conditional p th quantile of $U(s + \Delta s)$ for infinitesimal Δs . This requirement implies that a blip of treatment of brief duration Δs can only make a small difference proportional to Δs between any quantiles of $U(s + \Delta s)$ and $U(s)$. Regardless of local rank preservation, references [18], [19] and [21] show that one can estimate the parameters of a parametric non-RP SNFTM model $d(u, \bar{X}_s, \psi)$ for $d(u, \bar{X}_s)$ by G -estimation and use these G -estimates of ψ to obtain estimates of the survival distribution under both dynamic and non-dynamic regimes exactly as for RP SNFTMs. In the absence of RP, the observed random variable $H(\psi)$ has the same distribution as, but is not equal to, the counterfactual variable $U = U(0)$ when $\psi = \psi^*$. However, the non-RP SNFTMs include the RP SNFTMs as a special case, because, in those rare instances where local rank preservation actually holds, the new definition of $d(u, \bar{X}_s)$ in terms of quantile-quantile functions reduces to the RP definition of $d(u, \bar{X}_s)$ and thus $H(\psi^*) = U$.

DL's interesting and biologically plausible unaltered naturalistic course (UNC) assumption provides an excellent case study of the subtle but crucial distinction between rank preserving and non-rank preserving SNFTMs. The UNC assumption is that antidepressant therapy prevents clinically diagnosed depression only by suppressing the clinical symptoms of an underlying biological depressive episode; as a consequence the failure time (i.e. time to clinical depression) of a survivor currently off treatment does not depend on the duration of previous treatment. We will show that the UNC assumption by itself implies the biological implausibility of the rank preserving assumption that $U(s)$ is differentiable, (that is, $U(s + \Delta s) - U(s)$ divided by Δs goes to a finite limit). Specifically, we prove in Appendix A that if $U(s)$ is differentiable and DL's UNC assumption both hold, then any subject whose survival time is prolonged by his having received treatment must fail at the very instant that treatment is discontinued. As this implication is biologically implausible (and easily rejected by the data, assuming a moderately large beneficial treatment effect), we conclude the UNC assumption implies the implausibility of the differentiability of $U(s)$. Let us try to understand the substantive basis for this result. Differentiability of $U(s)$ essentially means that each very short blip Δs of treatment can only prolong a subject's survival by a very small amount, proportional to Δs . The wearing of seat belts is an example of a treatment for which this assumption would fail, since, by wearing a seat belt just at the moment of a high speed crash, survival can be prolonged for years. Analogously, if, as under the UNC assumption, the effect of antidepressant therapy is only to suppress depressive symptoms, then $U(s)$ cannot be differentiable. To see why consider a subject with an underlying biological depressive episode beginning at time s^* and ending at time s who took symptom suppressing antidepressives until just after s and who then developed a second depressive episode with overt clinical depression seven weeks later so $T = s + 7$ and $U(s) = s + 7$. However, if the subject had stopped antidepressives at time $s - \Delta s$, he would have become clinically depressed at s so $U(s - \Delta s) = s$. Because $U(s)$ is not differentiable, we cannot use a RP SNFTM when the UNC assumption holds.

However, treatment from time $s - \Delta s$ to s will have only a small effect on the difference in conditional quantiles of $U(s - \Delta s)$ and $U(s)$ since, under the realistic assumption that the times

at which episodes of underlying depression end have a continuous distribution, the fraction of the subpopulation (defined by the conditioning event $(\bar{X}_s, Z_s = 1)$) whose underlying depressive episodes will end in any infinitesimal interval $(s - \Delta s, s)$ is proportional to the length Δs , which implies a difference in conditional quantiles proportional to Δs . For such treatments we can always validly specify and fit a non-rank-preserving SNFTM under DL's UNC assumption.

5. DIFFICULTIES WITH DL'S POOLED RISK SET METHODOLOGY

In their Section 4.1, DL chose not to address the problem of the small sample sizes associated with many regimes by specifying either an MSM or SNFTM, but rather by applying their pooled risk set methodology under their unaltered naturalistic course (UNC) assumption. We will show that even when their sequential ignorability and UNC assumptions both hold, DL's pooled risk set methodology will fail whenever Z_{s-1} is a direct cause of a component V_s of X_s .

5.1. Difficulties in estimation of $\lambda_{d_s}(t|\bar{X}_s)$

DL's sequential ignorability assumption alone implies that conditional on covariate history \bar{X}_s up to s , the hazard at t had everyone followed the regime d_s (by stopping treatment after time $s - 1$) equals the observed hazard at t among those untreated at s but treated at $s - 1$ (that is, $Z_s = 0, Z_{s-1} = 1$). That is

$$\lambda_{d_s}(t|\bar{X}_s) = \lambda(t|Z_s = 0, \bar{X}_s, Z_{s-1} = 1) \quad \text{for } t \geq s$$

My understanding of a chief claim made in DL's Section 4.1 is that if we additionally impose DL's UNC assumption, then \bar{Z}_{s-1} history does not predict the hazard of failure given $(Z_s = 0, \bar{X}_s)$ so that

$$\lambda(t|Z_s = 0, \bar{X}_s, Z_{s-1} = 1) = \lambda(t|Z_s = 0, \bar{X}_s), \quad t \geq s \quad (1)$$

and thus

$$\lambda_{d_s}(t|\bar{X}_s) = \lambda(t|Z_s = 0, \bar{X}_s) \quad \text{for } t \geq s$$

Since there will be many more individuals with $Z_s = 0$ than with $(Z_s = 0, Z_{s-1} = 1)$, were DL's claim correct, modelling of $\lambda_{d_s}(t|\bar{X}_s)$ would be greatly facilitated by imposing the UNC assumption as one would then only need to model $\lambda(t|Z_s = 0, \bar{X}_s)$ by pooling risk sets as DL do in their equation (5). However, I will demonstrate DL's claim to be incorrect. A formal argument is provided in Appendix B using causal directed acyclic graphs. A less formal discussion is provided in the text below.

One might naively think that under the UNC assumption, equation (1) must be true; but, as DL point out, this need not be so, owing to selection effects. For example, if there is a direct cause (such as an underlying propensity for biological depression) of repeated episodes of clinical depression that is not recorded in \bar{X}_s , then (1) will be false. DL recognize this difficulty when they remark in the paragraph after their equation (5) that the recorded covariates 'must include...factors determining the loss of more vulnerable subjects'. However, it is not apparent that they ever elevate this remark to a formal assumption. More importantly, it is difficult to

imagine that this assumption would hold, owing to the existence of unrecognized and thus unmeasured genetic causes of biological depression.

Further, DL do not recognize that even if their assumptions hold and all the direct causes of repeated episodes of depression have been recorded in \bar{X}_s , none the less (1) will still generally be false if (i) an unrecorded variable U_{s-1} at time $s-1$ is a direct cause of both Y_{s+1} and of a component V_s of X_s , and (ii) Z_{s-1} is a direct cause of V_s .

As a simple example, suppose, conditional on past recorded covariate history \bar{X}_{s-1} and depression status Y_s at s , subjects on antidepressives at $s-1$ ($Z_{s-1}=1$) are more likely than subjects off medication ($Z_{s-1}=0$) to attend their next clinic visit V_s (because many subjects off medication feel there is little reason to attend clinic and thus must have their clinical status Y_s evaluated at home by outreach workers). Further suppose subjects who had an unrecorded dispute U_{s-1} with a superior at work (or other authority figure) at time $s-1$ are more likely both to skip their next clinic visit V_s (because the dispute U_{s-1} has made them adverse to engaging authority figures, including physicians) and to become depressed at time $s+1$ (but not at time s , as it takes some time for the effect of the dispute to influence symptoms). Then (1) will be false at $t=s$, because, among those who are untreated at s , undepressed at s , and fail to attend clinic at s , those with $Z_{s-1}=1$ likely had a dispute (in order to explain their non-attendance) and thus are more likely to be depressed at time $s+1$, than are those with $Z_{s-1}=0$ (as they already have a reason for not attending clinic and so are not at increased risk of depression, owing to a dispute).

Now equation (1) is testable from the data. Therefore, rather than imposing *a priori* assumptions sufficient to imply (1), one could simply fit a logistic model such as $\text{logit}(\lambda(t|Z_s=0, \bar{X}_s, \bar{Z}_{s-1})) = \alpha_{0t} + \alpha_1 W_s + \phi \text{cum}(\bar{Z}_{s-1})$ to the pooled risk sets and test the hypothesis that $\phi=0$. Here $\text{cum}(\bar{Z}_{s-1}) = \sum_{j=0}^{s-1} Z_j$ and W_s is a covariate vector of functions of \bar{X}_s . Indeed, as pointed out in reference [14], page 279, even if the test of $\phi=0$ rejects, one obtains a much more efficient estimate of the hazard $\lambda_{d_s}(t|\bar{X}_s) = \lambda(t|Z_s=0, \bar{X}_s, Z_{s-1}=1)$ of interest by using the pooled risk set MLE $\text{expit}(\hat{\alpha}_{0t} + \hat{\alpha}_1 W_s + \hat{\phi}s)$ of $\lambda(t|Z_s=0, \bar{X}_s, Z_{s-1}=1)$.

5.2. Difficulties in estimation of $\lambda_{d_s}(t)$ using modified inverse weights

Suppose that assumptions, including the UNC assumption, sufficient to imply (1) are true so that $\lambda_{d_s}(t|\bar{X}_s) = \lambda(t|Z_s=0, \bar{X}_s)$. DL's final target of inference remains the marginal hazard $\lambda_{d_s}(t)$. They propose using an inverse probability weighted estimator with the modified inverse probability weights given in their Appendix A. In Appendix C, we prove that this estimator is generally inconsistent for $\lambda_{d_s}(t)$ whenever Z_{s-1} and a component V_s of X_s are conditionally dependent given $(\bar{X}_s, Y_s=0)$, as will be the case if Z_{s-1} is a direct cause of V_s (even if there is no unrecorded variable U_{s-1} that is a direct cause of both Y_{s+1} and of V_s).

In summary, I believe that the methods proposed by DL in their Sections 4 and 5 should be avoided, except in the exceptional case in which one has a strong substantive belief that none of the conditions which invalidate equation (1) holds and that Z_{s-1} is not a direct cause of any component V_s of X_s .

6. EVALUATION OF DL'S CRITIQUE

DL's comment in their discussion section that the structural parameter in the 'G-estimation' model relates different outcomes for the same subject under different treatments is true of

rank preserving SNFTMs but is not true of non-rank preserving SNFTMs (which instead relate quantiles of survival distributions under different treatments); thus, for the reasons summarized in Section 4.2.2, non-rank preserving SNFTMs models offer important conceptual advantages over those structural models that have been developed by economists and others that are rank preserving. DL go on to say that, like sequential ignorability, structural models for potential outcomes are not directly testable from the data. This statement is true for RP SNFTMs in the sense that we can never obtain a consistent test (that is, a test with power converging to 1 in large samples) of the hypothesis implied by the RP SNFTM model $d(U(s), \bar{X}_s, \psi) = 1 - \exp\{\psi_1 + \psi_2 X_s^*\}$ of Section 4.2.1 that a subject's outcome $U(0)$ when never treated is equal to the function $\int_0^T \exp\{Z_u(\psi_1^* + \psi_2^* X_u^*)\} du$ of their observed treatment, covariate, and failure history, since we never observe $U(0)$ for treated subjects. However, DL's statement is false for non-rank preserving SNFTMs and MSMs in the sense that, given sequential ignorability (or no unmeasured confounders for the potential survival outcomes), at sufficiently large sample sizes, an omnibus goodness-of-fit test of any parametric MSM or non-RP SNFTM that has power arbitrarily close to one can be constructed (that is, there exist consistent tests). (This result follows from the fact that MSMs and non-RP SNFTMs are models relating distributions under different regimes rather than individual outcomes.) Of course, we agree with DL that with the sample sizes occurring in practice, it will be impossible to construct omnibus tests that have good power in most directions. In applied [23–25], as opposed to theoretical papers [19, 21], co-workers and I have stressed RP SNFTMs over non-RP SNFTMs wholly for pedagogic purposes; the rank preserving models are easier to explain. However, as discussed earlier, the RP models are almost always biologically implausible and thus, from a philosophical and conceptual perspective, the non RP models should really be used instead. However, the RP and non-RP models lead to exactly the same estimates of regime specific hazards and survival curves, which is why I and co-workers could afford to restrict consideration to RP models in our applied work.

In their discussion, DL are incorrect when they state that the approach to sensitivity analysis developed by myself and co-authors allows treatment choices to depend on (future) potential outcomes, while assuming that there are no unmeasured cofounders. We, like DL and others, assume that unmeasured confounders exist. However, in our approach [26, 27], we implicitly marginalize over the unmeasured confounders U and quantify the net confounding by means of a selection bias function that encodes the degree of dependence between the treatment and counterfactual outcome within levels of measured covariates after marginalization. Our approach is in contrast to approaches in which one explicitly varies the association of U with the outcome Y (within levels of treatment and measured confounders) and/or the association of U with the treatment (within levels of measured confounders) [28–31]. The advantage of our approach is that (i) there are many fewer sensitivity parameters to vary, and (ii) the (essentially impossible) decision as to whether to view U as univariate or multivariate, continuous or discrete, is done away with. A link between the two approaches is that the counterfactual variables can be considered the ultimate unmeasured confounder U . This reflects the fact that, given the counterfactuals and treatment, other unmeasured covariates U fail to predict the observed outcome (and thus are superfluous and can be dispensed with), since the observed outcome variable is a deterministic function of the treatment and the counterfactual outcome.

In my opinion, the standard unmeasured confounder U approach should be generally preferred to our counterfactual approach only in circumstances, where (i) U represents a known confounder (for example, cigarette smoking) that for logistical reasons was not measured in a


particular study, and furthermore, (ii) there exists reasonable historical knowledge about the magnitude of association of U with both the outcome (conditional on treatment and measured confounders) and the treatment (conditional on measured confounders). In contrast, when U is to represent all possible unmeasured factors, we believe that it is often easier for subject-matter experts to give their opinions about the plausible magnitude of the association of the counterfactual outcome with treatment than about the question of whether any unmeasured confounders U are continuous or discrete, single or multi-dimensional, and the associations of such confounders with treatment and/or outcome. Our counterfactual approach also leads to extremely simple computations that can be carried out with standard software. In contrast, as discussed in references [30] and [31], there can be formidable computational issues associated with the approach based on positing an unmeasured covariate.

In their final paragraph, DL contrast their local approach, which restricts the number of potential outcomes being modelled by restricting consideration to monotone treatment patterns, to a more global approach, associated with structural models like MSMs and SNFTMs, in which non-monotone treatment patterns are included. I believe that to a large extent this is a false dichotomy. First, as shown in Section 4 above and in reference [14], both MSMs and SNFTMs are applicable to and greatly benefit the analysis of monotone treatment data. Second, for such data, the number of potential outcome distributions that can be modelled with those MSMs which do not include baseline covariates X_0 is precisely equal to the number of potential outcomes distributions considered by DL in their approach, that is, the potential outcomes distributions associated with the non-dynamic regimes $(c, d_s; s = 0, \dots, T)$. The number of potential outcomes distributions modelled by SNFTMs and MSMs which include interactions between treatment and time independent covariates (MSMs) and time dependent covariates (SNFTMs) will typically be much greater in number because for example, potential outcomes distributions associated with dynamic regimes can be included; however this is a point in favour of rather than against these models, as it would constitute malpractice to fail to discover a qualitative interaction between treatment and a covariate (for example, bone marrow suppression) representing time-dependent treatment-induced toxicity. Finally, even with nonmonotone treatment patterns, one can follow DL and chose to restrict consideration to potential outcome regimes outcomes associated with monotone treatment patterns by artificially censoring subjects at the time they reinitiate treatment. However, such censoring must be regarded as informative and analysed appropriately, say using IPTW estimators. However, from a medical perspective, it is probably unwise to restrict the analysis to potential outcome distributions associated with monotone treatment regimes. For example, cancer patients on chemotherapy often must stop therapy temporarily due to bone-marrow toxicity; usually the optimal regime will turn out to be a non-monotone dynamic regime which prescribes that therapy be restarted at some appropriate time following recovery of bone-marrow function.


APPENDIX A

A1. Theorem

Suppose the following are true: (i) the UNC assumption that if $t > s$, (a) $U(t) \geq U(s)$ and (b) if $Z(t) = 0$ and $U(s) \geq t$ then $U(t) = U(s)$, and (ii) the derivative $\dot{U}(t)$ exists and is finite for all t such that $U(t) \geq t$. Then $U(t) > U(0)$ implies $U(t) = t$.



DAG 1



DAG 2

Figure A1.

A2. Proof by contradiction

Assume $U(t) > U(0)$ but $U(t) \neq t$. Then by $U(t) > t$, there exists $s \in (U(0), t)$ and $\varepsilon > 0$ such that $U(s) - s = \varepsilon$ and $\dot{U}(s) > 0$. This in turn implies that there exists $\sigma > 0$ and $s^* \in [s - \sigma, s]$ such that $U(s - \sigma) = U(s) - \dot{U}(s^*)\sigma > U(s) - \varepsilon = s$, where the first equality is by the mean value theorem and the second by the fact that σ can be chosen as small as we wish and $\dot{U}(s^*)$ is finite. Further note $\dot{U}(s^*)$ is positive because $U(s - \sigma)$ is strictly less than $U(s)$ as $\dot{U}(s)$ is positive. Hence, since $s > s - \sigma$ and $U(s - \sigma) > s$, it follows from the UNC supposition (i) that $U(s - \sigma) = U(s)$ which, substituting, implies $U(s) = U(s) - \dot{U}(s^*)\sigma$ which is false since both $\dot{U}(s^*)$ and σ are positive.

APPENDIX B

The causal DAG 1 in Figure A1 represents a non-parametric structural equation model (NPSEM) as defined in Pearl [32]. Arrows on the graph represent direct causal relations. I [33] show that a NPSEM has a nearly equivalent interpretation as a member of a class of counterfactual causal models proposed in reference [3] [36]. With the exception of U , all variables on the DAG are assumed to have been recorded for data analysis. U is an unmeasured common cause of survival at time 1 and 2. Because there are no arrows from any unmeasured variable U directly into treatments Z_0 or Z_1 , DL's sequential ignorability assumption holds [34]. Further, the absence of an arrow directly from Z_0 to Y_2 is equivalent to DL's UNC hypothesis. Pearl's [32] graphical d -separation procedures can be used to determine all the conditional and unconditional independencies implied by the underlying causal structure that obtain in the observed data. In particular, Z_0 is not d -separated from Y_2 conditional upon the observed variables Y_1, Z_1, X_1 , because, in Pearl's terminology, conditioning on the 'collider' Y_1 opens the path $Z_0 - Y_1 - U - Y_2$. This implies that Z_0 and Y_2 will generally be dependent conditional on (Y_1, Z_1, X_1) , and thus equation (1) will be false. An exception to

such conditional dependence would be the unlikely and empirically untestable special case in which the effect of the unmeasured U on the underlying continuous time hazard is additive; that is, $\lambda(t | Z(t), U = u) - \lambda(t | Z(t), U = 1)$ is a function, say $r(t, u)$, that does not depend on $Z(t)$.

On the causal DAG 2 of Figure A1, there is no longer an unmeasured common cause U of Y_1 and Y_2 . Rather, there is now an unmeasured common cause of X_1 and Y_2 . Further, Z_0 is a direct cause of X_1 . The UNC assumption holds because there is no arrow directly from either Z_0 or X_1 to Y_2 . Again, Z_0 is not d -separated from Y_2 given (Y_1, X_1, Z_1) , because conditioning on the collider X_1 opens the path $Z_0 - X_1 - U - Y_2$. Thus again equation (1) will generally be false.

APPENDIX C

Without loss of generality, we consider the special case where $s = t = 1$, there are no baseline covariates X_0 , and X_1 is discrete. It follows from DL's sequential ignorability assumption, the G -computation algorithm formula and $\lambda(1 | Z_1 = 0, Z_0 = 1, x_1) = \lambda(1 | Z_1 = 0, x_1)$ from equation (1) that

$$\begin{aligned} \lambda_{d_1}(1) &= \sum_{x_1} \lambda(1 | Z_1 = 0, x_1) f(x_1 | Y_1 = 0, Z_0 = 1) \text{pr}[Y_1 = 0 | Z_0 = 1] / \text{pr}[Y_1 = 0 | Z_0 = 1] \\ &= \sum_{x_1} \lambda(1 | Z_1 = 0, x_1) f(x_1 | Y_1 = 0, Z_0 = 1) \end{aligned}$$

Now DL's estimator with modified inverse weights is easily seen to be converging to $\lambda_{d_1}^*(1) = \sum_{x_1, z_0} \lambda(1 | Z_1 = 0, x_1) f(x_1 | Y_1 = 0, z_0) \text{pr}[Y_1 = 0 | Z_0 = 1] f(z_0) / \text{pr}[Y_1 = 0 | Z_0 = 1] = \{ \text{pr}[Z_0 = 1] \lambda_{d_1}(1) + \text{pr}[Z_0 = 0] \sum_{x_1} \lambda(1 | Z_1 = 0, x_1) f(x_1 | Y_1 = 0, Z_0 = 0) \}$. Hence, in general, $\lambda_{d_1}^*(1)$ will equal $\lambda_{d_1}(1)$ only if $X_1 \perp\!\!\!\perp Z_0 | Y_1 = 0$. Here we have assumed that in the last sentence of the first paragraph of their section 'Hazard calculations from DTSM' DL meant either to have $R_{s_t-1} = R_{s_{s-1}} = \{i | Y_{s-1,i} = 0\}$ when $s = t$ or to have instead $N_{s-1}^s = W_{s-1} / e_{s|s-1}^s$.

APPENDIX D

Our proof that $\hat{\lambda}_c(t)$ is CAN if the model for the propensity score is correct is based on the fact that $\hat{\lambda}_c(t)$ is also the solution $\hat{\mu}_{\text{eff}}$ to the augmented IPTW estimating equation $0 = \sum_i \hat{U}_{\text{eff},i}(\mu)$ where $\hat{U}_{\text{eff}}(\mu)$ is an estimate of the efficient score $U_{\text{eff}}(\mu) = (Y_{t+1} - I(Y_t = 0)\mu)Z_t/e_t - [\sum_{m=0}^t (Z_m - Z_{m-1}e_{m|m-1})(G_{m0} - G_{m1}\mu)/e_m] = \{ \sum_{m=1}^{t+1} \{ (G_{m0} - G_{(m-1)0}) - \mu(G_{m1} - G_{(m-1)1}) \} Z_{m-1}/e_{m-1} \} + G_{00} - G_{01}\mu$ for $\mu = \lambda_c(t)$ in the non-parametric model that only assumes no unmeasured confounders for potential survival outcomes. Specifically the estimate $\hat{U}_{\text{eff}}(\mu)$ has \hat{e}_m and $\hat{e}_{m|m-1}$ substituted for e_m and $e_{m|m-1}$ and \hat{G}_{mj} substituted for G_{mj} . It follows from the second representation for $U_{\text{eff}}(\mu)$ and the fact that $(\hat{\alpha}_{mj}, \hat{\kappa}_{mj})$ satisfies the linear logistic score equation that the estimate of the sample average of the term in the set braces is precisely 0 and thus $\sum_i \hat{U}_{\text{eff},i}(\mu) = \sum_i \hat{G}_{00,i} - \hat{G}_{01,i}\mu$ so $\hat{\mu}_{\text{eff}}$ solving $0 = \sum_i \hat{U}_{\text{eff},i}(\mu)$ equals $\hat{\lambda}_c(t) = \sum_i \hat{G}_{00,i} / \sum_i \hat{G}_{01,i}$. However, the first representation for $U_{\text{eff}}(\mu)$ implies that $\hat{\mu}_{\text{eff}}$ is CAN

if the propensity model is correctly specified since, each term of $U_{\text{eff}}(\mu)$ has mean zero when $\mu = \lambda_c(t)$ and $e_{m|m-1}$ is the true propensity score.

Note that if we left the term $\kappa_{mj}\hat{e}_m^{-1}$ out of the model for the conditional expectation of $\hat{G}_{(m+1)j}$, the estimator $\hat{\lambda}_c(t) = \sum_i \hat{G}_{00,i} / \sum_i \hat{G}_{01,i}$ would remain CAN if the logistic model $\text{expit}(\alpha_{mj}W_{mj})$ were correct but would no longer be doubly robust and would differ from the estimator $\hat{\mu}_{\text{eff}}$ solving $0 = \sum_i \hat{U}_{\text{eff},i}(u)$. However, $\hat{\mu}_{\text{eff}}$ would still be doubly robust in the sense that it would be CAN if either the models $\text{expit}(\alpha_{mj}W_{mj})$ or the propensity models were correct, but $\hat{\mu}_{\text{eff}}$ now would have the augmented regression estimator form since the estimate of the sample average of the term in the set braces will no longer be zero.

Finally, the IPTW $\hat{\mu}_{\text{IPTW}} = \{\sum_i Z_{t,i} Y_{(t+1),i} / \hat{e}_{t,i}\} / \{\sum_i Z_{t,i} / \hat{e}_{t,i}\}$ will equal $\hat{\mu}_{\text{eff}}$ and thus be doubly robust if the estimate of the sample average of the term in brackets in the first representation of U_{eff} is identically zero. This can be accomplished with standard software if we take \hat{e}_t to be the value at convergence as $k \rightarrow \infty$ of \tilde{e}_{tk} where \tilde{e}_{tk} is obtained by estimating the parameters $(\gamma_{mk}, \theta_{mk})$ recursively for $k = 1, 2, \dots$, of the extended logistic models $\text{expit}(\gamma_{mk}W_m + \theta_{mk}(\hat{G}_{m0} - \hat{G}_{m1}\mu) / \tilde{e}_{m(k-1)})$ for the outcome Z_m among subjects with $Z_{m-1} = 1, m = 0, \dots, T$, to obtain estimates $\tilde{e}_{m|m-1,k} = \tilde{e}_{m|m-1,k}(\tilde{\gamma}_{mk}, \tilde{\theta}_{mk})$ and $\tilde{e}_{mk} = \prod_{l=0}^m \tilde{e}_{l|l-1,k}$.

REFERENCES

1. 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**(9–12):923–945.
2. Robins JM. Causal inference from complex longitudinal data. In *Latent Variable Modeling and Applications to Causality*. Lecture Notes in Statistics (120), Berkane M (ed.). Springer-Verlag: New York, 1997; 69–117.
3. Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods – application to control of the healthy worker survivor effect. *Mathematical Modelling* 1986; **7**:1393–1512.
4. Scharfstein DO, Rotnitzky A, Robins JM. Adjusting for non-ignorable drop-out using semiparametric non-response models. (with Discussion and Rejoinder) *Journal of the American Statistical Association* 1999; **94**:1096–1120 (1121–1146).
5. Robins JM, Rotnitzky A, Bonetti M. Discussion of a Paper by Frangakis and Rubin. *Biometrics* 2001; **57**(2):343–347.
6. Robins JM, Rotnitzky A, van der Laan M. Comment on 'On profile likelihood' by S.A. Murphy and A.W. van der Vaart. *Journal of the American Statistical Association – Theory and Methods* 2000; **95**(450):431–435.
7. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. *Proceedings of the Journal of the American Statistical Association*, Section on Bayesian Statistics 2000, pp. 6–10.
8. Rosenbaum PR. Model-based direct adjustment. *Journal of the American Statistical Association* 1987; **82**:387–394.
9. Robins JM, Rotnitzky A. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology – Methodological Issues*. Jewell N, Dietz K, Farewell V (eds). Birkhäuser, Boston, MA, 1992; 297–331.
10. Heyting A, Tolboom JTB, Essers JGA. Statistical handling of drop-outs in longitudinal clinical trials. *Statistics in Medicine* 1992; **11**:2043–2062.
11. Robins JM, Wang N. Inference for imputation estimators. *Biometrika* 2000; **87**(1):113–124.
12. 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, pp. 24–33.
13. Robins JM, Rotnitzky A, Zhao L-P. Analysis of semiparametric regression models for repeated outcomes in the presence of missing data. *Journal of the American Statistical Association* 1995; **90**:106–121.
14. Robins JM. Correction for non-compliance in equivalence trials. *Statistics in Medicine* 1998; **17**:269–302.
15. Robins JM. Marginal structural models. 1997 *Proceedings of the American Statistical Association*, Section on Bayesian Statistical Science, 1998, pp. 1–10.
16. Robins JM. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical Models in Epidemiology: The Environment and Clinical Trials*. Halloran ME, Berry D (eds). Springer-Verlag: New York, 1999, 95–134.

17. Hernán M, Brumback B, Robins JM. Marginal structural models to estimate the joint causal effect of non-randomized treatments. *Journal of the American Statistical Association – Applications and Case Studies* 2001; **96**(454):440–448.
18. Robins JM, Blevins D, Ritter G, Wulfsohn M. G-estimation of the effect of prophylaxis therapy for pneumocystis carinii pneumonia on the survival of AIDS patients. *Epidemiology* 1992; **3**:319–336.
19. Robins JM. Structural nested failure time models. In *The Encyclopedia of Biostatistics*, Armitage P, Colton T (eds). Wiley: Chichester, 1998; 4372–4389.
20. Murphy S, van der Laan M, Robins JM. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association* 2001; **96**(456):1410–1423.
21. Robins JM. Analytic methods for estimating HIV treatment and cofactor effects. In *Methodological Issues of AIDS Mental Health Research*. Ostrow DG, Kessler R (eds). Plenum Publishing: New York, 1993, 213–290.
22. Robins JM. The analysis of randomized and non-randomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. In *Health Service Research Methodology: A Focus on AIDS*. Sechrest L, Freeman H, Mulley A (eds). U.S. Public Health Service, National Center for Health Services Research: Washington, D.C., 1989; 113–159.
23. Mark SD, Robins JM. Estimating the causal effect of smoking cessation in the presence of confounding factors using a rank preserving structural failure time model. *Statistics in Medicine* 1993; **12**:1605–1628.
24. Witteman JC, d’Agostino RB, Stijnen T, Kannel WB, Cobb JC, deRidder MAJ, Hoffman A, Robins JM. G-estimation of causal effects: isolated systolic hypertension and cardiovascular death in the Framingham Study. *American Journal of Epidemiology* 1998; **148**:390–401.
25. Keiding N, Filiberti M, Esbjerg S, Robins JM, Jacobsen N. The graft-versus leukemia effect after bone marrow transplantation: A case study using structural nested failure time models. *Biometrics* 1997; **55**(1):23–28.
26. Robins JM, Scharfstein D, Rotnitzky A. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical Models in Epidemiology: The Environment and Clinical Trials*. Halloran ME, Berry D (eds). Springer-Verlag: New York, 1999; 1–94.
27. Robins JM. Association, causation, and marginal structural models. *Synthese* 1999; **121**:151–179.
28. Schlesselman JJ. Assessing effects of confounding variables. *American Journal of Epidemiology* 1978; **108**: 3–8.
29. Rosenbaum PR, Rubin DB. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society, Series B* 1983; **11**:212–218.
30. Lin DY, Psaty BM, Kronmal RA. Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics* 1998; **54**:948–963.
31. Rosenbaum PR. *Observational Studies*. Springer-Verlag: New York, 1995.
32. Pearl J. Causal diagrams for empirical research. *Biometrika* 1995; **82**(4):669–688.
33. Robins JM. Discussion of ‘Causal diagrams for empirical research by J. Pearl’. *Biometrika* 1995; **82**(4): 695–698.
34. Pearl J, Robins JM. Probabilistic evaluation of sequential plans from causal models with hidden variables. Uncertainty in Artificial Intelligence. *Proceedings of the 11th Conference*, 1995, pp. 444–453.
35. Robins JM, Rotnitzky A. Comment on the Bickel and Kwon article, ‘Inference for semiparametric models: Some questions and an answer’. *Statistica Sinica*, 2001; **11**(4):920–936.
36. Robins JM. Semantics of causal DAG models and the identification of direct and indirect effects. In: *Highly Structured Stochastic Systems*, P. Green, N. Hjort, S. Richardson (Eds), Oxford University Press 2002 (to appear).
37. Robins JM. Optimal-regime structural nested models. *Proceedings of the Second Seattle Symposium on Biostatistics*. Lin, DY, Haegerty, P (Eds), 2002 (to appear).
38. Murphy, SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society B*, 2002 (to appear).