Statistics
in Medicine

Special Issue Paper

Received 27 November 2012, Accepted 14 June 2013 Published online 2 August 2013 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.5907

Estimation of the effect of interventions that modify the received treatment

S. Haneuse[a] and A. Rotnitzky[b]*†

Motivated by a study of surgical operating time and post-operative outcomes for lung cancer, we consider the estimation of causal effects of continuous point-exposure treatments. To investigate causality, the standard paradigm postulates a series of treatment-specific counterfactual outcomes and establishes conditions under which we may learn about them from observational study data. While many choices are possible, causal effects are typically defined in terms of variation of the mean of counterfactual outcomes in hypothetical worlds in which specific treatment strategies are 'applied' to all individuals. For example, one might compare two worlds: one where each individual receives some specific dose and a second where each individual receives some other dose. For our motivating study, defining causal effects in this way corresponds to (hypothetical) interventions that could not conceivably be implemented in the real world. In this work, we consider an alternative, complimentary framework that investigates variation in the mean of counterfactual outcomes under hypothetical treatment strategies where each individual receives a treatment dose corresponding to that actually received but modified in some pre-specified way. Quantification of this variation is defined in terms of contrasts for specific interventions as well as in terms of the parameters of a new class of marginal structural mean models. Within this framework, we propose three estimators: an outcome regression estimator, an inverse probability of treatment weighted estimator and a doubly robust estimator. We illustrate the methods with an analysis of the motivating data.

Keywords: causal inference; observational study; marginal structural mean model; double robustness

1. Introduction

Lung cancer is the most common cancer in the world, with early-stage non-small-cell lung cancer (NSCLC) accounting for 80–90% of all cases. While surgical resection is the current treatment of choice for NSCLC, patients and physicians must nevertheless weigh the risks and benefits of surgery as well as manage expectations. Investigating surgery-specific factors is particularly important in that an understanding of their impact on patient outcomes may help in the development of strategies that improve efficiency and, eventually, reduce healthcare costs. With this in mind, at a recent collaborative consultation, we were posed with the following question: What is the impact of operating time (i.e., the time spent in surgery) on the risk of post-operative outcomes among patients undergoing surgical resection for NSCLC?

As posed, this consultation question seeks to characterize the causal effect of a continuous point-exposure treatment on an outcome. For the evaluation of causal effects of point-exposure treatments, the so-called counterfactual or potential outcomes model is often adopted [1]. This model conceptualizes a series of hypothetical worlds across which different treatment strategies are implemented. The model also conceptualizes a set of counterfactual outcomes for each individual in the population, each corresponding to the value taken by the outcome within each hypothetical world. Causal analysis then proceeds by defining the treatment effect in terms of some target parameter that summarizes differences between the distributions of the distinct counterfactuals and then establishing conditions under which the parameter is estimable. Finally, estimators of the target parameter are developed whose consistency often relies on the validity of the identifying conditions as well as some dimension-reducing assumptions.

[a]*Department of Biostatistics, Harvard School of Public Health, Boston, MA, U.S.A.*
[b]*Department of Economics, Universidad Di Tella, Buenos Aires, Argentina*
**Correspondence to: A. Rotnitzky, Department of Economics, Universidad Di Tella, Buenos Aires, Argentina.*
†*E-mail: arotnitzky@utdt.edu*

Based on this general paradigm, in recent years, there has been an explosion of methods for causal inference. Whereas statistical developments for point-exposure causal effects have mostly focused in settings where treatments take on one of a finite number of values, causal inference with continuous treatments has also been considered [2, 3]. The direct application of this paradigm to our collaborative setting, however, is problematic. In particular, while one could, in theory, consider a series of hypothetical worlds where individuals receive some pre-specified length of surgery, perhaps according to their baseline covariates, one could not reasonably implement an intervention that imposed such 'treatments' in the real world. As such, causal contrasts defined in the standard manner do not have any practical utility. To see this more concretely, suppose that for some subpopulation, under current medical practice, operating times range between 50 and 120 min. The standard paradigm would consider different hypothetical worlds, each being a world in which the operating times of all subjects in the subpopulation are the same. Consider, for example, the world where each operation is 100 min. While one could (at least in theory) investigate the distribution of the counterfactuals in this world, it is hard to conceive of an intervention that simultaneously prolongs the surgery of those individuals whose operation actually lasted 50 min and shortens the operation of those individuals whose operation actually lasted 120 min. Consequently, understanding the distribution of counterfactuals in a world where all operations take 100 min does not provide information that can be translated into real-world practice.

An alternative paradigm, recently proposed by Diaz Munoz and van der Laan [4], considers hypothetical worlds in which the treatment is randomly assigned according to mechanisms that possibly depend on baseline covariates and the distribution of the treatment received in the factual world. For instance, suppose that in a given subpopulation defined by the level of baseline covariates, surgery length was distributed according to a uniform(50, 120) distribution. The Diaz Munoz and van der Laan [4] setting would allow one to estimate the effect of random policies in which the surgery length received by each subject in the subpopulation is drawn from a law that is some transformation of the uniform(50, 120) distribution, for instance, uniform($50+\delta_1$, $120-\delta_2$) with δ_1 and δ_2 positive constants. This setting remains unsatisfactory in our present application because the hypothetical world still contemplates the possibility that subjects whose actual surgery length was, say, 50 min would have a surgery length close to $120-\delta_2$ min.

In the context of our collaborative setting, a more plausible and, arguably, substantively relevant hypothetical world is one in which the operating time that each subject actually experienced, say, A minutes, is reduced by some amount that depends on A, say $q(A)$. We refer to $q(\cdot)$ as the *modified treatment policy (MTP)*. As an example, consider the hypothetical world in which everyone's operating time was reduced by 5 min; that is, $q(A) = A-5$. Arguably, a treatment modification of this magnitude is conceivable and could, for example, be achieved by improving the speed at which the stitching is carried out. However, even this hypothetical world may not always be plausible. For example, certain surgeons may have already reached their maximum stitching speed. This suggests that one consider an intervention that yields a reduction of operative time that depends not only on A but also on pre-treatment covariates L. A causal analysis could then be conducted by investigating the (estimated) contrast between the rate of post-surgical complications under (i) current practice and (ii) in a hypothetical world in which everybody's operative times was reduced to, say, $q(A, L)$. Alternatively, one could also study how the rates of post-surgical complications vary with different choices of $q(A, L)$, for instance, how they vary with different modest amounts of δ where $q(A, L) = A - \delta$.

A modified treatment policy is different from a dynamic treatment regime (DTR) [5–8]. The latter is a sequence of time-varying decision rules that specify which treatment, among a set of available ones, should be offered next based on patient covariate and responses collected during prior therapy. When treatment decisions can be made only once, a DTR is simply a rule that specifies the treatment assignment as a function of baseline covariates. Thus, a modified treatment policy $q(A, L)$ is a generalization of a point exposure DTR, in that the latter is a function $q(L)$ of just the baseline covariates L whereas the former is a function $q(A, L)$ of treatment actually received and, possibly, of baseline covariates L. This extension involves conceptual and technical nuances.

Conceptually, MTPs are appealing in settings like the one that motivated this work where the set of feasible treatments for each subject depends on attributes that are not fully captured by baseline covariates but which are reasonably captured by the treatment actually received. However, unlike DTRs, in general, MTPs are rules that cannot be implemented in practice because they depend on the treatment level A actually received. For instance, it is impossible to force a surgery to last 5 min less than its natural duration as this would require to first know what such natural duration would be. As a consequence, causal effects of MTPs, unlike those of DTRs, are not even in principle experimentally testable. That is,

it is impossible to run randomized controlled trials to compare different MTPs. Nevertheless, MTPs are valuable insofar they can help generate hypothesis about mechanisms that explain their causal effects and thus formulate promising interventions that can be tested in randomized experiments. For instance, suppose that we find that for subjects undergoing thoracotomy, a reduction of 5 min in the duration of surgery reduces the rate of a major post-surgical complication during hospital stay. This result would suggest that it would be worthwhile running a randomized study that examines interventions that could potentially reduce thoracotomy length by 5 min. Examples of interventions might include training programs to improve the speed/efficiency of surgical stitching skill, the introduction of a pre-operative checklist to enhance efficiency during surgery or the use of some new piece of technology that increases efficiency in some aspect of the surgery.

Technically, there are important subtleties about identification and estimation of causal effects of MTPs. As we will argue in Section 2, identification requires a so-called positivity condition that for every possible level that the modified treatment $q(A, L)$ can take, there exists in the target population at least one subject that would naturally received that level, a demand that we will argue is not realistic in our example and hence forces us to redefine the target population on which the effects of MTPs are quantified. Furthermore, it requires the exchangeability across all possible treatment levels of just two subpopulations, one receiving a treatment level and the second receiving the modified treatment level. This assumption differs also in subtle ways from the exchangeability condition required for identification of DTRs. Finally, contrasts that quantify the causal effects of MTPs, unlike those of DTRs, depend on the treatment assignment probabilities. This distinction has implications for inference. As we will see in Section 3, just as for DTRs, it is possible to construct inverse probability weighted (IPW) and doubly robust estimators of causal contrasts for MTPs; however, unlike DTRs, standard error estimators that ignore the uncertainty in the treatment probabilities are not guaranteed to be conservative.

To the best of our knowledge, MTPs were first introduced in [9] and further discussed in [10]. A conceptual distinction between the applications that motivated these articles and our work is that they considered settings in which one can conceive, at least in principle, interventions that instantaneously change the naturally occurring level A of treatment to a new exposure level $q(A, L)$. For instance, Taubman *et al.* [10] considered a rule that stipulates that daily physical activity over a period, say 1 month, should last as long as it would in the absence of intervention provided this natural duration is at least 30 min; otherwise, it should last 30 min. This rule is indeed an MTP as it is defined by the function $q(A, L) = AI(A \geqslant 30) + 30I(A \leqslant 30)$ that depends on the natural duration A of daily physical activity. Yet, for this MTP, one can at least in principle conceive an intervention where any subject that exercises less than 30 min on any given day is forced to exercise for 30 min on that day. In fact, these authors considered not just point exposure but also time-dependent dynamic interventions at time points $t_1, \ldots, t_J$ in which the treatment to be administered at time t_j is allowed to depend on the treatment that would be received at time t_j if the planned interventions were made through t_{j-1} but no intervention was made at t_j. Robins *et al.* [9] and Taubman *et al.* [10] derived an extension of the g-computation formula that yields the survival distribution of a time-to-event endpoint in a hypothetical world in which one such intervention is implemented, provided certain identifiability assumptions hold. They also described survival estimators that rely on estimates of the conditional distributions that enter into the extended g-formula. However, these authors did not discuss the positivity requirements for identifiability, perhaps because for the applications that motivated their work, these are likely trivially satisfied.

As far as we know, the only other existing article in the literature that defines MTPs for point exposures and discusses identifiability conditions is Shpitser and Pearl [11]. Unlike the work of Robins and colleagues and our work, that article assumes graphical, rather than counterfactual, models. The article discusses neither the subtle positivity requirements for identification nor inference.

This article contributes to the literature on MTPs in the following ways. First, supported by our motivating example, it raises the important point not previously noticed that MTPs can be a useful analytic concept even when they refer to policies that cannot be possibly implemented. Second, it explicitly formulates conditions under which causal effects of MTPs can be identified from observational data. Third, it discusses analytic strategies that can be followed when, as in our motivating example, these conditions are unrealistic. Fourth, it shows that the g-formula that identifies the causal contrast of interest agrees with the identifying formula of the contrast quantifying the effect of the random intervention of Munoz and van der Laan [4]. Fifth, it provides three distinct estimation procedures: a so-called outcome regression estimator whose consistency relies on the correct specification of a model for the conditional mean of the outcome given baseline covariates and received treatment and which is the point-exposure/continuous-outcome equivalent of the parametric g-formula estimators discussed in [9, 10]; an

IPW estimator whose consistency relies on the correct specification of a model for the treatment assignment probabilities, which differs from a similar estimator proposed by Munoz and van der Laan [4] in that it is ensured to fall in the parameter space; and a doubly robust estimator whose consistency relies on the correct specification of either the outcome mean model or the treatment model but not necessarily both. Munoz and van der Laan [4] also described a doubly robust estimator, the so-called targeted maximum likelihood (TML) estimator that, like the doubly robust estimator proposed here, takes an outcome regression form. However, unlike the TML estimator, our estimator is non-recursive and can be implemented using standard regression software. Sixth and lastly, it describes a novel marginal structural mean model whose parameters quantify the effects of different modified treatment policies and discuss doubly robust estimation of the model parameters.

The remainder of this article is as follows. In Section 2, we outline the proposed framework, introduce the causal contrasts of interest and develop the conditions under which they are identified. Section 3 proposes three estimators for the contrasts defined within our proposed framework; extensions to the estimation of components of a marginal structural model are outlined in Section 4. The methods are applied to our motivating dataset in Section 5. Finally, Section 6 concludes the paper with a brief discussion.

2. Causal contrasts for continuous treatments

Consider an observational study with n subjects drawn at random from the target population. Suppose we observe (L_i, A_i, Y_i), independent and identically distributed across $i = 1, \ldots, n$; Y_i is a scalar outcome, A_i is a treatment variable with support on the set $\mathcal{A}$ and $L_i = (L_{i,1}, \ldots, L_{i,K})$ is a vector of pre-treatment variables known or hypothesized to confound the effect of A_i on Y_i. In the context of our consultation study, A_i is the operating time, and $\mathcal{A}$ is the set of all durations that surgery can possibly take in the target patient population, Y_i is an indicator of the occurrence of a post-surgery complication and L_i includes pre-surgical covariates known (or hypothesized) to be associated with both operating time and risk of post-surgical complications (see Section 5).

As indicated in Section 1, when seeking to establish causation from observational data, the counterfactual model is often adopted. This model assumes that for each a in a set $\mathcal{A}'$ of candidate doses (often $\mathcal{A}' = \mathcal{A}$) and each subject ω of the target population Ω, there exists a counterfactual outcome $Y_a(\omega)$ defined as the value that the outcome for subject ω would take on in the hypothetical world in which one could intervene and set the level of treatment to a. Implicit in this notation is the so-called stable unit treatment value assumption (SUTVA) [12], which postulates that (i) the treatment status of any unit does not affect the counterfactual outcomes of the other units [13] and (ii) all means at arriving at a given treatment a result in the same outcome $Y_a(\omega)$ [14]. SUTVA implies the following assumption:

C0. *Consistency*: If $A(\omega) = a$, then $Y(\omega) = Y_a(\omega)$, $\forall \omega \in \Omega$.

In the context of our collaboration, condition (i) is realistic. The validity of condition (ii) is less clear. In particular, there may be many ways by which one can arrive at an operating time, each resulting in a different post-operative outcome. Nevertheless, throughout this paper, we assume SUTVA and, in particular, that the consistency assumption C0 holds; a discussion of the consequences of the violation of condition (ii) in our application is postponed to Section 5.4.

Aside from potential violation of condition (ii), in our motivating application, the possibility of being able to set the treatment level for every subject to any of the $a \in \mathcal{A}'$ is unrealistic. For instance, as indicated in Section 1, for a subject ω whose surgery took 120 min, it is unrealistic to assume that one could reduce his or her operative time to, say, $a = 60$ min. More realistic in our application is to postulate a relaxed assumption that states that for each subject ω, there exists a set $\mathcal{A}_\omega$ of enforceable treatment levels. Formally, this is tantamount to assuming that the random variables Y_a, for $a \in \mathcal{A}'$, are defined on distinct probability spaces, the domain of Y_a being $\Omega_a = \{\omega \in \Omega : a \in \mathcal{A}_\omega\}$. Clearly, the set $\mathcal{A}_\omega$ contains the value $A(\omega)$ of the treatment actually received by subject ω. The remaining elements of $\mathcal{A}_\omega$ depend on the specific application. For instance, if it is conceivable that one could have shortened subject ω's operative time by 5 min, then $A(\omega) - 5 \in \mathcal{A}_\omega$.

2.1. *Defining causal contrasts*

The lack of a unique domain for $Y_a(\cdot)$ for all $a \in \mathcal{A}'$ questions the scientific relevance of characterizing causal effects by examining variation of $\mathrm{E}[Y_a]$ as a function of $a \in \mathcal{A}'$. In particular, we view this

approach as uninformative for treatment effects because $\mathrm{E}[Y_a]$ is the mean of Y_a over subpopulations Ω_a that vary with a.

When, as in our application, $\mathcal{A}_\omega$ changes with ω, a reasonable alternative target of analysis is an extension of the so-called average treatment effect on the treated (ATT) causal contrast for binary treatments. Specifically, for a binary treatment A, ATT is defined as $\mathrm{E}[Y_1|A=1] - \mathrm{E}[Y_0|A=1]$. This contrast quantifies the average change on the outcome that switching from treatment $A=1$ to treatment $A=0$ would have on those that were actually treated with $A=1$. This contrast is particularly relevant when it is conceivable to enforce $A=0$ on those that actually received $A=1$, but no action can be conceived that would enforce treatment $A=1$ on subjects that actually received $A=0$. That is, when $\mathcal{A}_\omega = \{0,1\}$ for subjects ω with $A(\omega)=1$ and $\mathcal{A}_\omega = \{0\}$ for subjects with $A(\omega)=0$.

For a continuous treatment, assume that the same treatment doses could be enforced on all subjects that actually received the same treatment dose and had the same baseline covariates; that is, assume that if $(A(\omega), L(\omega)) = (A(\omega'), L(\omega'))$, then $\mathcal{A}_\omega = \mathcal{A}_{\omega'}$. Then, the ATT contrast extends in the following natural way. For the group that actually received dose $A=a$ and had baseline covariates $L=l$, consider a scientifically relevant and enforceable dose $q(a,l)$, that is, $q(a,l) \in \mathcal{A}_\omega$ for all ω such that $(A(\omega), L(\omega)) = (a,l)$. Then, $\mathrm{E}[Y_a|A=a, L=l] - \mathrm{E}\left[Y_{q(a,l)}|A=a, L=l\right]$ quantifies the effect of switching from the dose actually received, that is, a, to the enforceable dose $q(a,l)$ among those who received dose a and had baseline covariates l. By noticing that $\mathrm{E}[Y_a|A=a, L=l] = E(Y|A=a, L=l)$ and averaging over all possible combinations of received doses and covariate levels, we arrive at the following contrast that quantifies the average effect in the studied population of a policy that switches each received dose A to the modified dose $q(A,L)$:

$$\mathrm{E}[Y] - \int \mathrm{E}\left[Y_{q(a,l)}|A=a, L=l\right] \mathrm{d}F_{A,L}(a,l) \tag{1}$$

For notational convenience, we denote the second term by $\mathrm{E}\left[Y_Q\right]$, with $Y_Q \equiv Y_{q(A,L)}$.

Note, strictly speaking, contrast (1) is not well defined because conditional distributions can be defined arbitrarily on events $(A=a, L=l)$ of zero probability; as such, the integral in (1) may not be unique. To resolve this, we assume that for each (a,l) in the support of (A,L), the law of $Y_{q(a,l)}|A=a', L=l'$ is (weakly) continuous in all the arguments of (a',l') that represent realizations of absolutely continuous random variables. Next, we assume that the laws entering the integral in (1) are the unique weakly continuous laws. Likewise in the rest of this article, we assume that the law$(Y|A=a, L=l)$ and the law$(A|L=l)$ are (weakly) continuous in all the arguments of (a,l) that represent realizations of absolutely continuous random variables, and throughout, we take the conditional laws as the unique ones prescribed by continuity.

2.2. *Identifiability conditions*

Extensions of estimands such as $\mathrm{E}\left[Y_Q\right]$ to time-dependent intervention plans were considered by Robins *et al.* [9] and Taubman *et al.* [10]. Also, Shpitser and Pearl [11] considered the estimand $\mathrm{E}[Y_Q]$ when L is empty and A is discrete. These three articles arrived at a formula that identifies $\mathrm{E}[Y_Q]$ as a function of the observed data distribution under certain exchangeability conditions. The formula is identical to the formula (2) that we derive later. However, these papers assumed that $Y_a(\omega)$ is well defined for all $a \in \mathcal{A}$ and all $\omega \in \Omega$. What exactly must be assumed for $\mathrm{E}[Y_Q]$ to be identified when this is not the case is subtle and is discussed next.

In the Supporting information, we show that $\mathrm{E}[Y_Q]$ is identified; that is, it can be computed as a function of the law of the factual variables (Y, A, L), if the following two conditions hold. In the sequel for any random vector V, supp$V$ denotes the support of the probability distribution of the vector V.

C1. *Positivity*: If $(a,l) \in \mathrm{supp}(A,L)$, then $(q(a,l),l) \in \mathrm{supp}(A,L)$.
C2. *Conditional exchangeability of policy related subpopulations*: For each $(a,l) \in \mathrm{supp}(A,L)$ and $a' \equiv q(a,l)$, law$(Y_{a'}|A=a, L=l) =$ law$(Y_{a'}|A=a', L=l)$.

Positivity states that if in the target population there is a positive probability of finding a subject, say ω, with covariates $L=l$ who in the absence of intervention would be administered dose $a \in A$, then there is also a positive probability of finding in the same population a subject, say ω', with the same covariates $L=l$ and who would be administered dose $q(a,l)$ in the absence of intervention. Note that positivity alone does not imply that dose $q(a,l)$ is necessarily feasible for subject ω. For instance, the subpopulation with covariates $L=l$ may be composed of subjects with surgery lengths from 50 to

120 min. If ω had a surgery that lasted 110 min, and $q(a,l) = a - 50$, then positivity holds for $a = 110$ and $L = l$, but the surgery length of $q\,(110, l) = 60$ min is unfeasible for subject ω. However, positivity combined with the exchangeability condition C2 essentially implies that dose $q(a,l)$ is feasible. This is because condition C2 is essentially the assumption that those that received dose a could have received dose $q(a,l)$ and that the assignment to doses a and $q(a,l)$ was randomized (by nature) according to a, possibly, L-dependent mechanism. Note that condition C2 is weaker than the usual no-unmeasured confounders assumption that postulates that in the subpopulation with $L = l$, any subject could have received any dose a observed for that group and that assignment to any such doses was randomized. Such assumption is too strong and unrealistic in our setting because it implies that all doses observed in a given subpopulation are feasible for all members of the subpopulation.

In the Supporting information, we show that under assumptions C0, C1 and C2, $\mathrm{E}[Y_Q]$ is identified and equal to the extended g-computation formula of Robins *et al.* [9]:

$$\mathrm{E}\left[Y_Q\right] = \int m(q(a,l),l)\,\mathrm{d}F_{A,L}(a,l) \tag{2}$$

where $m(a,l) \equiv \mathrm{E}\,[Y\,|A = a, L = l]$. Note, in the right-hand side of (2), $m(q(a,l),l) = \mathrm{E}\,[Y\,|A = q(a,l), L = l]$ is the mean of the factual outcome Y among those that received treatment dose $q(a,l)$ and had covariates l. We emphasize that the positivity condition C1 ensures that such a group is not empty and, hence, that $m(q(a,l),l)$ is well defined for all $(a,l) \in \mathrm{supp}(A,L)$.

When $q(a,l) = a - \delta$ for a given δ, the right-hand side of (2) agrees with the identifying formula derived by Diaz Munoz and van der Laan [4] under stronger positivity and exchangeability assumptions for a causal contrast quantifying the effects of a stochastic intervention in which subjects in the subpopulation with covariates $L = l$ are randomly assigned to a (discrete) treatment A with randomization probability $\Pr(A = a|L = l) = f_{A|L}(a + \delta|l)$. Note that we interpret the causal parameter identified by formula (2) differently, namely as quantifying the effect of an intervention that deterministically decreases by δ the dose actually received by each subject in the subpopulation with covariates $L = l$. As discussed earlier, in the context of our consultation, for given choices of δ, the latter 'deterministic' intervention is conceivable whereas the former 'stochastic' intervention is not. For instance, suppose that for the subpopulation with covariates $L = l$, the actual operative times are uniformly distributed in the interval (50,120), and $\delta = 5$, then an intervention that reduces by 5 min each subject's surgery length is conceivable whereas an intervention that randomly assigns each subject in the subpopulation to a surgery of a length between 45 and 115 min according to a uniform distribution is not.

2.3. *Ensuring positivity*

If the support of the conditional distribution of A given covariates L is $(b(L), b^*(L))$ for some $-\infty < b(L) \leqslant b^*(L) < \infty$, as will be the case if A is operative time, policies $q(a,l)$ will have to be carefully defined so as to ensure that the positivity assumption holds. Suppose for example that we are interested in determining the effect of a reduction of δ minutes in operative time. To ensure that positivity holds, we could define

$$q(a,l) = \begin{cases} a - \delta & \text{if } a > b(l) + \delta \\ a & \text{if } a \leqslant b(l) + \delta \end{cases} \tag{3}$$

Alternatively, we could redefine the target of inference to be $\mathrm{E}\,[Y\,|A > b(L) + \delta]$–$\mathrm{E}\left[Y_Q|A > b(L) + \delta\right]$, that is, the difference in the means of Y and Y_Q in the subpopulation that satisfies $A > b(L) + \delta$.

Both causal contrasts are special cases of a general causal target contrast defined as follows. Let $c(\cdot)$ and $d(\cdot)$ be investigators' chosen functions with domain the support of L, satisfying $-\infty \leqslant c(l) < d(l) \leqslant +\infty$, and such that for every a in the interval $(c(l), d(l))$, both (a,l) and $(q(a,l),l)$ belong to $\mathrm{supp}(A,L)$.

We redefine the target causal parameter of interest to be as follows:

$$\Delta = \mu - \mu_Q, \tag{4}$$

where $\mu = \mathrm{E}\,[Y\,|I_L(A) = 1]$, $\mu_Q = \mathrm{E}\left[Y_Q|I_L(A) = 1\right]$ and, for notational convenience, $I_L(A) = 1$ if $c(L) < A < d(L)$ and 0 otherwise. This contrast reduces to the earlier contrast (1) if for each l, $(c(l), d(l))$ is chosen to be the support of the distribution of A given $L = l$. Otherwise, it refers to the effect on a specific subset of the population where a treatment, possibly different, than the one actually

received could be realistically implemented. For instance, in the context of our example, if $q(a,l)$ is the modified treatment policy (3) and $(c(l),d(l))$ is the entire support of A given $L=l$, then Δ quantifies the reduction in risk of post-surgical complications *in the entire patient population*, of reducing the operative time by δ minutes only on a subset where such reduction is feasible, that is, on those for which $A > b(L)+\delta$. On the other hand, if $c(l)$ is chosen to be $b(l)+\delta$, then Δ quantifies the reduction in risk *just in the subpopulation satisfying* $A > b(L)+\delta$, that is, in those for which a reduction of δ minutes is regarded as feasible.

3. Estimation

In this section, we consider estimation and inference for Δ. The first component of the contrast, that is, μ, can be consistently estimated by the sample mean of Y among those subjects who received a treatment dose A_i in the interval $(c(L_i),d(L_i))$:

$$\widehat{\mu} = \frac{\sum_{i=1}^{n} I_{L_i}(A_i)Y_i}{\sum_{i=1}^{n} I_{L_i}(A_i)} \tag{5}$$

Note that if for each l, $(c(l),d(l))$ is chosen equal to the support of A given $L=l$, then the actual values of $c(l)$ and $d(l)$ need not be known in order to compute $\widehat{\mu}$ as $I_{L_i}(A_i)=1$ with probability 1.

Towards estimating the second component of Δ, that is, μ_Q, we first provide two key identifying formulas. The first is a straightforward extension of formula (2), namely

$$\mu_Q = \frac{\mathrm{E}\left[I_L(A)m\left(Q,L\right)\right]}{\mathrm{E}\left[I_L(A)\right]} \tag{6}$$

which, as shown in the Supporting information, holds under assumptions C0, C1 and C2. The second requires the validity of the following additional technical condition, which ensures that one can partition $\mathcal{I}(l) = (c(l),d(l))$ so that on each partition interval, say $\mathcal{I}_j(l)$, the integral $\int_{\mathcal{I}_j(l)} m(q(a,l),l)\,\mathrm{d}F_{A|L}(a|l)$ can be computed by the change of variables formula:

C3. *Piecewise smooth invertible policy*: Suppose that for each l, there exists a partition of the interval $\mathcal{I}(l)=(c(l),d(l))$ into intervals $\mathcal{I}_j(l)$, $j=1,\ldots,J(l)$, such that on $\mathcal{I}_j(l)$, $q(\cdot,l)$ is equal to $q_j(\cdot,l)$ and on the interior of $\mathcal{I}_j(l)$, $q_j(\cdot,l)$ has differentiable inverse function $h_j(\cdot,l)$ with derivative denoted by $h'_j(\cdot,l)$.

Condition C3 does not require full differentiability of $q(\cdot,l)$ but rather it makes the less stringent requirement of piecewise differentiability of $q(\cdot,l)$. This flexibility is important as many interesting MTPs such as (3) or the modified treatment policy of Taubman *et al.* [10] $q(a,l) = aI(a \geqslant 30) + 30I(a \leqslant 30)$ discussed in Section 1, are differentiable piecewise but not on their entire domain.

Under C0, C1, C2 and C3, μ_Q can be re-expressed as

$$\mu_Q = \frac{\mathrm{E}\left[\lambda(A,L)Y\right]}{\mathrm{E}\left[\lambda(A,L)\right]} \tag{7}$$

where

$$\lambda(a,l) = \sum_{j=1}^{J(l)} I_{j,l}(h_j(a,l))w_j(a,l)h'_j(a,l),$$

$I_{j,l}(u)=1$ if $u \in \mathcal{I}_j(l)$ and 0 otherwise, and

$$w_j(a,l) = \frac{f_{A|L}(h_j(a,l)|l)}{f_{A|L}(a|l)}$$

where $f_{A|L}(\cdot|\cdot)$ denotes the density of the (weakly continuous version of the) conditional distribution of A given L. Proofs are provided in the Supporting information[‡].

[‡]*Supporting information may be found in the online version of this article.*

We thus see that μ_Q can be expressed as a weighted mean of the outcome Y, with weights $\lambda(A, L)$. The weight λ is computed in such a way so that the outcome of a subject ω' in the population who, under the modified treatment policy, would have $q\left(A\left(\omega'\right), L\left(\omega'\right)\right) = a$ is equally represented by the observed outcome $Y(\omega)$ of every subject, say ω, in the population whose treatment $A(\omega)$ in the absence of intervention is equal to a and who has $L\left(\omega'\right) = L(\omega)$.

For instance, as indicated earlier, condition C3 holds when $q(a, l)$ is as defined in (3). In such case if $\mathcal{I}(l) = (c(l), d(l))$ denotes the support of A given L, and $b(l) = c(l)$, then $\mathcal{I}_1(l) = (b(l), b(l) + \delta]$, $\mathcal{I}_2(l) = (b(l) + \delta, d(l))$, $q(a, l)$ is equal to $q_1(a, l) = a$ on $\mathcal{I}_1(l)$ and $q_2(a, l) = a - \delta$ on $\mathcal{I}_2(l)$. In this example, $h_1(a, l) = a$ and $h_2(a, l) = a + \delta$, so

$$\lambda(a, l) = I_{1,l}(a) + I_{2,l}(a + \delta)\frac{f_{A|L}(a + \delta|l)}{f_{A|L}(a|l)}$$

Here, a subject ω with $A(\omega) \in (b(l), b(l)+\delta] \cap (b(l), d(l)-\delta)$ has $I_{1,l}(A(\omega)) = 1$ and $I_{2,l}(A(\omega)) = 1$, so $\lambda\left(A(\omega), L(\omega)\right) = 1 + \frac{f_{A|L}(A(\omega)+\delta|L)}{f_{A|L}(A(\omega)|L)}$. The constant 1 in $\lambda(A(\omega), L(\omega))$ appears because the subjects own observed outcome stands to represent his counterfactual outcome under the MTP since according to (3), $q\left(A(\omega), L(\omega)\right) = A(\omega)$ for $A(\omega) \in (b(l), b(l) + \delta]$. The ratio $\frac{f_{A|L}(A(\omega)+\delta|L)}{f_{A|L}(A(\omega)|L)}$ in $\lambda\left(A(\omega), L(\omega)\right)$ appears because the outcome of subject ω and those of all other subjects that have the same values of (A, L) as subject ω, by virtue of $A(\omega) \in (b(l), d(l) - \delta)$, account equally to represent the counterfactual outcome Y_Q of subjects ω' who have $q\left(A\left(\omega'\right), L\left(\omega'\right)\right) = A\left(\omega'\right) - \delta = A(\omega)$ and $L\left(\omega'\right) = L(\omega)$.

The expressions on the right-hand side of (6) and (7) suggest two estimators of μ_Q in the spirit of the familiar outcome regression and IPW estimators of counterfactual means $E(Y_a)$ under a fixed treatment a. Further, just as for estimation of $E(Y_a)$, it is also possible to construct doubly robust estimators of μ_Q. The following subsections describe the three estimation approaches.

3.1. *Outcome regression estimation*

Suppose that the conditional mean $\mathrm{E}\left[Y|A, L\right]$ is modeled as

$$\mathrm{E}\left[Y|A, L\right] = m(a, l; \tau^*) \tag{8}$$

where

$$m(a, l; \tau^*) = \Phi^{-1}\left\{r\left(a, l; \tau^*\right)\right\} \tag{9}$$

Φ^{-1} is an inverse link function with range in the parameter space of μ_Q, $r(\cdot, \cdot; \cdot)$ is a known function and τ^* is an unknown finite-dimensional vector. For example, if Y is binary, a common choice is the logistic regression $\mathrm{E}\left[Y|A, L\right] = \mathrm{expit}\left\{\tau_0 + \tau_{\mathrm{A}} A + \tau_{\mathrm{L}}^T L\right\}$, where $\tau = (\tau_0, \tau_{\mathrm{A}}, \tau_{\mathrm{L}})$ is a $(K + 2)$ vector of regression coefficients and $\mathrm{expit}\{u\} = \exp(u)/\left[1 + \exp(u)\right]$. Whatever the specification of $m(A, L; \tau)$, let $\widehat{\tau}$ denote a consistent estimator of τ^* under model (8). For example, if τ indexes the aforementioned logistic regression model, then $\widehat{\tau}$ is its maximum likelihood estimator (MLE); otherwise, $\widehat{\tau}$ is a weighted least squares estimator of τ.

Given $\widehat{\tau}$, the identity (6), suggests a so-called outcome regression estimator of μ_Q

$$\widehat{\mu}_{Q,OR} = \frac{\sum_{i=1}^n I_{L_i}(A_i) m(Q_i, L_i; \widehat{\tau})}{\sum_{i=1}^n I_{L_i}(A_i)} \tag{10}$$

where $Q_i \equiv q(A_i, L_i)$. Note $m(Q_i, L_i; \widehat{\tau})$ is the fitted conditional mean for the i^{th} individual when their treatment A_i is replaced by Q_i and their covariates are left as L_i.

Given the estimator $\widehat{\mu}_{Q,OR}$, we construct the following outcome regression estimator of Δ:

$$\widehat{\Delta}_{OR} = \widehat{\mu} - \widehat{\mu}_{Q,OR}. \tag{11}$$

In the Supporting information, we show, under C0, C1 and C2, correct specification of model (8) and standard regularity conditions for M estimators, $\widehat{\Delta}_{OR}$ is consistent and asymptotically Normally distributed. In particular,

$$\sqrt{n}\left(\widehat{\Delta}_{OR} - \Delta^*\right) \longrightarrow N(0,\ \mathcal{V}_{OR})$$

where Δ^* is the true value of Δ,

$$\mathcal{V}_{OR} = \mathrm{E}\left[I_L(A)\right]^{-2} \mathrm{Var}\left[I_L(A)(Y - m(Q, A; \tau^*) - \Delta^*) \; + \; \mathcal{D}_{OR}\mathcal{F}_{OR}^{-1} S(\tau^*)\right]$$
$$\mathcal{D}_{OR} = \frac{\partial}{\partial \tau^T}\mathrm{E}\left[I_L(A)m(Q, A; \tau)\right]\Big|_{\tau=\tau^*}$$
$$\mathcal{F}_{OR} = \frac{\partial}{\partial \tau^T}\mathrm{E}\left[S(\tau)\right]\Big|_{\tau=\tau^*},$$

and $S(\tau)$ is the estimating function used to compute $\widehat{\tau}$, that is, $\widehat{\tau}$ solves $0 = \sum_{i=1}^n S_i(\tau)$. For example, if $\widehat{\tau}$ is the MLE of τ^* under the aforementioned logistic model for (8), then $S_i(\tau) = X_i(Y_i - m(A_i, L_i; \tau))$. A consistent estimator of $\mathcal{V}_{OR}$ can be obtained by replacing in its formula the population variance and each of the population means by their sample counterparts and replacing each Δ^* and τ^* with $\widehat{\Delta}_{OR}$ and $\widehat{\tau}$. Alternatively, a bootstrap variance estimator can be used.

Note, because $\widehat{\mu}_{Q,OR}$ is a plug-in estimator of μ_Q, it follows that if $\widehat{\tau}$ is the MLE of τ under model (8), then $\widehat{\mu}_{Q,OR}$ is the semiparametric MLE of μ_Q and, hence, is semiparametric efficient under the model that assumes C0, C1, C2 and the outcome regression model (8). However, this does not imply that $\widehat{\Delta}_{OR}$ is a semiparametric efficient estimator of Δ because the non-parametric estimator $\widehat{\mu}$, given by expression (5), is not an efficient estimator of μ under model (8). Clearly, instead of $\widehat{\mu}$, we could use the MLE of μ under model (8), that is,

$$\widehat{\mu}_{ML} \; = \frac{\sum_{i=1}^n I_{L_i}(A_i)m(A_i, L_i; \widehat{\tau})}{\sum_{i=1}^n I_{L_i}(A_i)},$$

as a plug-in estimator of μ. Our philosophy, however, is that the sole purpose of model (8) is to overcome the curse of dimensionality associated with estimating $\mathrm{E}[Y|A = a, L = l]$ non-parametrically. As such, our preference is to minimize the dependence of the estimation of Δ on some assumed functional form for $E(Y|A = a, L = l)$ so as to minimize misspecification bias and, therefore, retain $\widehat{\mu}$ as the estimate of μ in expression (4).

3.2. *Inverse probability of treatment weighted estimation*

Suppose that the distribution of treatment A, conditional on covariates L, is parametrically modeled as having density

$$f_{A|L}(a|l) \; = \; f_{A|L}(a|l; \theta^*) \tag{12}$$

where θ^* is unknown finite-dimensional parameter vector θ^*, and for each θ, $f_{A|L}(a|l; \theta)$ is a known density. For example, if $A|L \sim \mathrm{Normal}(L^T\alpha, \sigma^2)$, then $\theta = (\alpha, \sigma)$.

Given some specification of $f_{A|L}(a|l; \theta)$, let $\widehat{\theta}$ denote the MLE of θ. Identity (7) suggests that given $\widehat{\theta}$, we compute an IPW estimator of μ_Q:

$$\widehat{\mu}_{Q,IPW} = \frac{\sum_{i=1}^n \lambda(A_i, L_i; \widehat{\theta})Y_i}{\sum_{i=1}^n \lambda(A_i, L_i; \widehat{\theta})}, \tag{13}$$

where, for any θ,

$$\lambda(A_i, L_i; \theta) = \sum_{j=1}^{J(l)} I_{j,l}\left\{(H_{j,i})\right\} w_{j,i}(\theta)h'_j(A_i, L_i),$$

and for notational convenience, $H_{j,i} = h_j(A_i, L_i)$ and

$$w_{j,i}(\theta) \; = \; \frac{f_{A|L}(H_{j,i}|L_i; \theta)}{f_{A|L}(A_i|L_i; \theta)}.$$

The form of identity (7) suggests that $\sqrt{n}(\widehat{\mu}_{Q,IPW} - \mu_Q^*)$ converges in law to a mean-zero Normal distribution when model (12) is correctly specified. However, for this to hold, we make a slightly more demanding positivity condition, which ensures that $\lambda(A_i, L_i)$ not only is well defined but also has finite variance:

C4. *Strengthened restricted positivity*: For some $\rho < \infty$,

$$\max_{j\in\{1,\ldots,J(L)\}} \sup_{a\in\mathcal{I}_j(L)} \frac{f_{A|L}\left(h_j(a,L)|L\right)}{f_{A|L}(a|L)} = \rho \text{ with probability 1}$$

If, for instance, $q(A|L)$ is as defined in (3), the strengthened positivity assumption holds if $f_{A|L}(\cdot|l) < \eta$ and $b(L)$ is strictly greater than the lower bound of the support of A given L so that $f_{A|L}(\cdot|l) > \kappa$ for some $\kappa > 0$.

Given the estimator $\widehat{\mu}_{Q,IPW}$, we construct the following IPW estimator of Δ:

$$\widehat{\Delta}_{IPW} = \widehat{\mu} - \widehat{\mu}_{Q,IPW}. \tag{14}$$

In the Supporting information, we show, under C0, C1, C3 and C4, correct specification of model (12) and standard regularity conditions for M estimators, $\widehat{\Delta}_{IPW}$ is consistent and asymptotically Normally distributed. In particular,

$$\sqrt{n}\left(\widehat{\Delta}_{IPW} - \Delta^*\right) \longrightarrow N(0,\ \mathcal{V}_{IPW})$$

where Δ^* is the true value of Δ,

$$\begin{aligned}
\mathcal{V}_{IPW} &= \mathrm{E}\left[I_L(A)\right]^{-2} \mathrm{Var}\left[I_L(A)(Y-\mu^*) - \lambda(A,L)(Y-\mu_Q^*) + \mathcal{D}_{IPW}\mathcal{F}_{IPW}^{-1}M(\theta^*)\right] \\
\mathcal{D}_{IPW} &= \frac{\partial}{\partial\theta^T}\mathrm{E}\left[\lambda(A,L;\theta)(Y-\mu_Q^*)\right]\Big|_{\theta=\theta^*} \\
\mathcal{F}_{IPW} &= \frac{\partial}{\partial\theta^T}\mathrm{E}\left[M(\theta)\right]\Big|_{\theta=\theta^*},
\end{aligned}$$

$M(\theta) = \partial \log f(A|L;\theta)/\partial\theta$ and μ^* and μ_Q^* are the true values of μ and μ_Q, respectively. A consistent estimator of $\mathcal{V}_{IPW}$ can be obtained by replacing in its formula the population variance and each of the population means by their sample counterparts and replacing each μ^*, μ_Q^* and θ^* with $\widehat{\mu}$, $\widehat{\mu}_{Q,IPW}$ and $\widehat{\theta}$. Alternatively, the bootstrap variance estimator can be used.

Note it is well known that IPW estimation of the counterfactual mean $E(Y_a)$ for a fixed treatment a is unstable due to the occasional presence of a unit i with small estimated $f_{A|L}(A_i|L_i)$ relative to the rest of the sample. Towards estimation of Δ, the IPW estimator (14) is intrinsically stabilized as it involves weighting by $\lambda(A_i, L_i)$, a sum of quantities $w_{j,i}(\widehat{\theta})$, each being a ratio of two (estimated) densities $f_{A|L}(\cdot|L_i;\widehat{\theta})$ evaluated at $h_j(A_i, L_i)$ and A_i, respectively. If $q(A_i, L_i)$ and A_i are close (i.e., if the policy $q(a, l)$ prescribes small changes of the received treatment), then the weights $\lambda(A_i, L_i)$ will not be large even if $f_{A|L}(A_i|L_i;\widehat{\theta})$ is large.

3.3. Doubly robust estimation

As they stand, neither $\widehat{\mu}_{Q,OR}$ nor $\widehat{\mu}_{Q,IPW}$ are entirely satisfactory estimators of μ_Q because their consistency depends on the validity of the dimension-reducing models (8) and (12). Towards providing some protection against misspecification of these models, we construct a doubly robust estimator. When Φ is a canonical link in a generalized linear model, the proposed doubly robust estimator can be easily implemented using software for fitting generalized linear models. The estimator is computed as the result of the following three steps:

Step 1: Compute the estimators $\widehat{\tau}$ and $\widehat{\theta}$ of the preceding subsections.
Step 2: For each unit i in the sample, compute $\widehat{\gamma}$ solving

$$\sum_{i=1}^{n} \lambda\left(A_i, L_i;\widehat{\theta}\right)\left[Y_i - \Phi^{-1}\left\{r\left(A_i, L_i;\widehat{\tau}\right) + \gamma\lambda\left(A_i, L_i;\widehat{\theta}\right)\right\}\right] = 0.$$

Note if Φ is a canonical link in a generalized linear model, the estimator $\widehat{\gamma}$ can be obtained with standard software by fitting a model with no intercept, a single covariate $X_{1,i} = \lambda\left(A_i, L_i;\widehat{\theta}\right)$ and an offset $X_{0,i} = r\left(A_i, L_i;\widehat{\tau}\right)$.

Step 3: Compute the estimator of μ_Q as

$$\widehat{\mu}_{Q,DR} = \frac{\sum_{i=1}^{n} I_{L_i}(A_i)\Phi^{-1}\left\{r\left(Q_i, L_i; \widehat{\tau}\right) + \widehat{\gamma}\lambda\left(Q_i, L_i; \widehat{\theta}\right)\right\}}{\sum_{i=1}^{n} I_{L_i}(A_i)}. \tag{15}$$

Note that formula (15) for $\widehat{\mu}_{Q,DR}$ is the same as formula (10) for the outcome regression estimator of μ_Q except that $m\left(Q_i, L_i; \widehat{\tau}\right) = \Phi^{-1}\{r\left(Q_i, L_i; \widehat{\tau}\right)\}$ is replaced by $\Phi^{-1}\left\{r\left(Q_i, L_i; \widehat{\tau}\right) + \widehat{\gamma}\lambda\left(Q_i, L_i; \widehat{\theta}\right)\right\}$. With $\widehat{\gamma}$ computed as in step 2, this replacement ensures that $\widehat{\mu}_{Q,DR}$ solves a doubly robust estimating equation (Theorem 4 of the Supporting information). This construction is analogous to that of double robust outcome regression estimators of parameters of conventional marginal structural models for point exposures (see, for example, [15]). The subtle and novel distinction is the form of the weights $\lambda\left(A_i, L_i; \widehat{\theta}\right)$.

The estimator $\widehat{\mu}_{Q,DR}$ can be used to construct the following estimator of Δ:

$$\widehat{\Delta}_{DR} = \widehat{\mu} - \widehat{\mu}_{Q,DR}. \tag{16}$$

In the Supporting information, we show, under C0, C1, C3 and C4, standard regularity conditions for M estimators and correct specification of model (8) or (12), but not necessarily both, $\widehat{\Delta}_{DR}$ is consistent and asymptotically Normally distributed. In particular,

$$\sqrt{n}\left(\widehat{\Delta}_{DR} - \Delta^*\right) \longrightarrow N(0, \mathcal{V}_{DR})$$

where Δ^* is the true value of Δ,

$$\mathcal{V}_{DR} = \mathrm{E}\left[I_L(A)\right]^{-2} \mathrm{Var}\left[U\left(\Delta^*, \tau^\dagger, \gamma^\dagger, \theta^\dagger\right) - \mathcal{D}_{DR,1}\mathcal{F}_{DR,1}^{-1}\widetilde{S}\left(\tau^\dagger, \gamma^\dagger; \theta^\dagger\right) - \mathcal{D}_{DR,2}\mathcal{F}_{DR,2}^{-1}M\left(\theta^\dagger\right)\right]$$

$$\mathcal{D}_{DR,1} = \frac{\partial}{\partial\left(\tau^T, \gamma^T\right)}\mathrm{E}\left[U(\Delta^*, \tau, \gamma, \theta^\dagger))\right]\Big|_{\tau=\tau^\dagger, \gamma=\gamma^\dagger}$$

$$\mathcal{F}_{DR,1} = \frac{\partial}{\partial\left(\tau^T, \gamma^T\right)}\mathrm{E}\left[\widetilde{S}\left(\tau, \gamma; \theta^\dagger\right)\right]\Big|_{\tau=\tau^\dagger, \gamma=\gamma^\dagger}$$

$$\mathcal{D}_{DR,2} = \mathcal{G}_{DR,2} - \mathcal{D}_{DR,1}\mathcal{F}_{DR,1}^{-1}\mathcal{G}_{DR,1}$$

$$\mathcal{G}_{DR,1} = \frac{\partial}{\partial\theta^T}\mathrm{E}\left[\widetilde{S}\left(\tau^\dagger, \gamma^\dagger, \theta\right)\right]\Big|_{\theta^\dagger}$$

$$\mathcal{G}_{DR,2} = \frac{\partial}{\partial\theta^T}\mathrm{E}\left[U\left(\Delta^*, \tau^\dagger, \gamma^\dagger, \theta\right)\right]\Big|_{\theta=\theta^\dagger}$$

$$\mathcal{F}_{DR,2} = \frac{\partial}{\partial\theta^T}\mathrm{E}\left[M(\theta)\right]\Big|_{\theta=\theta^\dagger},$$

with

$$\begin{aligned} U(\Delta, \tau, \gamma, \theta) &= I_L(A)\left[Y - \Phi^{-1}\left\{r(Q, L; \tau) + \gamma\lambda(Q, L; \theta)\right\} - \Delta\right] \\ &\quad - \lambda(A, L; \theta)\left[Y - \Phi^{-1}\left\{r(A, L; \tau) + \gamma\lambda(A, L; \theta)\right\}\right], \end{aligned}$$

$$\widetilde{S}(\tau, \gamma; \theta) = \begin{bmatrix} S(\tau) \\ \lambda(A, L; \theta)\left\{Y - \Phi^{-1}\left[r(A, L; \tau) + \gamma\lambda(A, L; \theta)\right]\right\} \end{bmatrix}$$

and $\tau^\dagger$, $\gamma^\dagger$ and $\theta^\dagger$ are the probability limits of $\widehat{\tau}, \widehat{\gamma}$ and $\widehat{\theta}$, respectively. Note $\tau^\dagger \equiv \tau^*$ and $\gamma^\dagger = 0$ if model (8) is correctly specified and $\theta^\dagger \equiv \theta^*$ if model (12) is correctly specified. A consistent estimator of $\mathcal{V}_{DR}$ can be obtained by replacing in its formula the population variance and each of the population means by their sample counterparts and replacing each Δ^*, $\tau^\dagger$, $\gamma^\dagger$ and $\theta^\dagger$ with $\widehat{\Delta}_{DR}$, $\widehat{\tau}$, $\widehat{\gamma}$ and $\widehat{\theta}$, respectively. Alternatively, the bootstrap variance estimator can be used.

Following [16], it is possible to show that when both models (8) and (12) are correctly specified, the asymptotic variance of $\widehat{\mu}_{Q,DR}$ is equal to the semiparametric variance bound for estimation of μ_Q in

the nonparametric model that imposes assumptions C0, C1, C3 and C4. It follows that when both (8) and (12) are correctly specified and $\widehat{\tau}$ is the MLE of τ under model (8), the asymptotic variance of $\widehat{\mu}_{Q,DR}$ is greater than or equal to that of the outcome regression estimator, $\widehat{\mu}_{Q,OR}$, because, as indicated in Section 3.1, $\widehat{\mu}_{Q,OR}$ is the semiparametric MLE of μ_Q under the smaller model that, in addition to C0, C1, C3 and C4, it also assumes the parametric form (8). The increase in the variance incurred by $\widehat{\mu}_{Q,DR}$, relative to $\widehat{\mu}_{Q,OR}$, is the price paid for the additional protection against misspecification of (8). An analytic comparison between the asymptotic variances of $\widehat{\mu}_{Q,IPW}$ and of $\widehat{\mu}_{Q,DR}$ is not possible because, even though $\widehat{\mu}_{Q,IPW}$ is consistent under correct specification of (12), it does not achieve the semiparametric variance bound for estimators of μ_Q under model that, in addition to C0, C1, C3 and C4, it also assumes (12). Note that, like the ATT contrast but unlike the ATE contrast, the asymptotic efficiency bound for estimation of μ_Q under this model is smaller than the bound under a model that does not impose (12) because, just like $\mathrm{E}[Y_{a=0}|A=1]$ but unlike $\mathrm{E}[Y_{a=0}]$, μ_Q is a functional that depends on $f_{A|L}(\cdot|\cdot)$ [17–19].

As indicated earlier, in the special case in which $(c(L), d(L))$ is equal to the support of A given L and $q(A|L) = A - \delta$, the identifying formula (2) for μ_Q agrees with the formula that identifies the causal estimand of Diaz Munoz and van der Laan [4]. The estimators proposed by these authors can therefore also be used to estimate μ_Q. In fact, our IPW estimator of μ_Q differs from the IPW proposed in that paper in that the denominator in formula (13) is replaced by n. By virtue of being a weighted average (with weights $\lambda(A_i, L_i; \widehat{\theta})$), our IPW estimators, unlike those in Diaz Munoz and van der Laan [4], are guaranteed to fall in the range of Y_i' and thus in the parameter space for μ_Q. Diaz Munoz and van der Laan [4] present two doubly robust estimators, one in the form of an augmented IPW estimator and another computed by the method of TML. Our choice of methodology for computing a doubly robust estimator is based on the following considerations: it is non-iterative, when $\Phi(\cdot)$ is a canonical link, it can be implemented easily with standard regression software, and because it has ultimately the form of an outcome regression estimator, it falls in the parameter space. The augmented IPW estimator is not guaranteed to fall in the parameter space and requires special programming. TML is a loss-based, recursive algorithm that produces an estimator of μ_Q, which has both the form of an outcome regression and IPW estimator. If the loss function is convex, the TML algorithm converges to a unique solution, and by virtue of being a substitution estimator, it falls in the parameter space. These two properties are shared by our non-iterative doubly robust estimator when Φ is a canonical link, which, unlike TML, is non-recursive. All doubly robust estimators have the same limiting distribution when both the treatment and the outcome regression models (8) and (12) are correctly specified. None dominates the other in terms of asymptotic efficiency when one of the models is incorrect.

Diaz Munoz and van der Laan [4] provide an estimator of the asymptotic variance of their IPW estimator, which they claim is conservative, citing van der Laan and Robins [20] to support their assertion. However, the theory in that reference is about inference on estimands that, like $\mathrm{E}[Y_a]$ but unlike μ_Q, do not depend on the treatment mechanism $f_{A|L}(\cdot|\cdot)$ and it does not apply to inference about μ_Q. Specifically, it is well known (e.g., Section 6.1 of Robins *et al.* [21]) that variance estimators of IPW estimators of $\mathrm{E}[Y_a]$ that do not account for the fact that the weights have been estimated are conservative. However, this result heavily relies on the fact that the statistical parameter that is being estimated in lieu of $\mathrm{E}[Y_a]$, namely $\int \mathrm{E}(Y|A=a, L=l)\,\mathrm{d}F_L(l)$, does not depend on $f_{A|L}(\cdot|\cdot)$ [22, 23]. In contrast, because the identifying formula for μ_Q does depend on $f_{A|L}(\cdot|\cdot)$ (Equation (2)), it is not true that failure to account for weight estimation yields conservative variance estimators of the IPW estimators of μ_Q. Yet, the variance estimators of Diaz Munoz and van der Laan [4] do not account for weight estimation, and thus, they are not ensured to be conservative. Likewise, their proposed variance estimator of their augmented IPW and TML estimators can be, contrary to what they claim, also anticonservative when model (12) is correct and model (8) is incorrect. We provide an example in the Supporting information.

4. Marginal structural models

So far, we have focused on the comparison between a single dosing policy $q(a, l)$ with current dosing standards. For the comparison of several candidate dosing policies in a candidate set $\{q_\delta(a, l) : \delta \in \Upsilon\}$, all satisfying C4 for the same $c(L)$ and $d(L)$, a model that parametrizes the dependence of $\mathrm{E}[Y_\delta|I_L(A)=1]$ as a function of δ is desirable where, for simplicity, throughout $Y_\delta = Y_{Q_\delta}$. More generally, if one wishes to examine how the dependence of the mean of Y_δ on δ varies across subpopulations defined by levels of a subset Z, possibly all, of the baseline covariates L, then a model parametrizing

the dependence of $\mathrm{E}[Y_\delta | I_L(A) = 1, Z]$ on δ and Z might be of interest, that is, a model of the following form:

$$\mathrm{E}[Y_\delta | I_L(A) = 1, Z] \;=\; \Phi^{-1}\{g(Z, \delta; \beta)\} \tag{17}$$

where Φ^{-1} is a given inverse link function, $g(\cdot, \cdot; \cdot)$ is a known function and β is an unknown parameter vector of dimension p. For instance, suppose that in the context of our consultation, A is operative time, L includes information on comorbid conditions and the surgery that jointly affect operative time and outcomes, $(c(l), d(l))$ is the support of A given $L = l$, $q_\delta(a, l)$ is defined like $q(a, l)$ in (3) for δ ranging in $\Upsilon = [0, 10]$ and $b(l) = c(l) + 10$. If Y is the indicator of a post-surgical complication, and Z is a binary indicator of gender (Z= 0/1 = female/male), we may take

$$\Phi^{-1}\{g(Z, \delta; \beta)\} = \frac{\exp\{\beta_0 + \beta_1\delta + \beta_2 Z + \beta_3 Z\delta\}}{1 + \exp\{\beta_0 + \beta_1\delta + \beta_2 Z + \beta_3 Z\delta\}}. \tag{18}$$

Then, $\exp(\beta_1 + \beta_3)$ is the proportionate change in the odds of a post-surgical complication in men caused by each additional reduction of 1 min in the surgery length for men for which such reduction is feasible. Model (17) is the generalization of a marginal structural model for dynamic regimes [24, 25] to modified treatment policies.

From the fact that each $\mathrm{E}[Y_\delta | I_L(A) = 1, Z]$ is identified when C0, C1 and C2 hold for each $q_\delta(\cdot, \cdot)$, we conclude that β is identified when these conditions hold simultaneously for all candidate policies. Just like for Δ, it is possible to construct three distinct estimators of β as we outline next. The three estimators will rely on an investigator's choice of a partition of the set Υ into intervals $[\delta_{t-1}, \delta_t)$ for $t = 1, \ldots, T$, where $\delta_1 < \delta_2 < \cdots < \delta_T$. Although not needed for consistency, to ensure estimators of β with good efficiency properties, we recommend that the partition be chosen so that the intervals $[\delta_{t-1}, \delta_t)$ have roughly the same length and T is at least 10. A similar issue arises in the estimation of marginal structural models for dynamic regimes [26]. In what follows, $\widehat{\tau}$ and $\widehat{\theta}$ are the estimators of subsections 3.1 and 3.2.

The first estimator, denoted $\widehat{\beta}_{OR}$, generalizes the outcome regression estimator of Section 3.1. It solves the estimating equation

$$\sum_{i=1}^{n}\sum_{t=1}^{T} s(Z_i, \delta_t) I_{L_i}(A_i) \left[m(Q_{\delta_t, i}, L_i, \widehat{\tau}) - \Phi^{-1}\{g(Z_i, \delta_t; \beta)\}\right] \;=\; 0 \tag{19}$$

where $Q_{\delta_t, i} = q_{\delta_t}(A_i, L_i)$ and $s(\cdot, \cdot)$ is a data analyst's chosen vector-valued function of the same dimension as β. For example, if $\Phi^{-1}\{g(Z_i, \delta_t; \beta)\}$ is as in (18), then we can choose

$$s(Z_i, \delta_t) = [1, \delta_t, Z_i, Z_i\delta_t]^T. \tag{20}$$

With this choice, $\widehat{\beta}_{OR}$ can be computed from any package that carries out logistic regression allowing continuous outcomes taking values in $(0, 1)$. To implement it, we create an extended dataset composed of $(T+1)\times n$ records, with $T+1$ records for each unit i in the sample. Record t of unit i contains copies of the variables L_i, A_i, Y_i, a pseudo covariate $D_{i,t} = \delta_t$ and a pseudo outcome $\widetilde{Y}_{i,t} = m(Q_{\delta_t, i}, L_i, \widehat{\tau})$. Then, $\widehat{\beta}_{OR}$ is the logistic regression estimator of outcome $\widetilde{Y}_{i,t}$ on covariates Z_i, D_{it} and $Z_i D_{it}$ and an intercept computed from the subset of the extended dataset corresponding to replicates of units i, which satisfy $c(L_i) < A_i < d(L_i)$.

In the Supporting information, we show that under regularity conditions, (19) has the solution $\widehat{\beta}_{OR}$ such that $\sqrt{n}(\widehat{\beta}_{OR} - \beta^*)$ converges to a mean-zero Normal random variable when C0, C1 and C2 and model (8) hold. In the Supporting information, we provide the expression for the variance of the limiting Normal distribution and a consistent estimator for it. The bootstrap can also be used as an alternative to the variance estimator given in the Supporting information.

The second estimator, denoted $\widehat{\beta}_{IPW}$, generalizes the IPW estimator of Section 3.2. It solves the estimating equation

$$\sum_{i=1}^{n}\sum_{t=1}^{T} s(Z_i, \delta_t)\lambda_{\delta_t}(A_i, L_i; \widehat{\theta}) \left[Y_i - \Phi^{-1}\{g(Z_i, \delta_t; \beta)\}\right] \;=\; 0 \tag{21}$$

where

$$\lambda_{\delta_t}(A_i, L_i; \widehat{\theta}) = \sum_{j=1}^{J(l)} I_{j,l}\left\{\left(H_{\delta_t,j,i}\right)\right\} w_{\delta_t,j,i}(\theta) h'_{\delta_t,j}(A_i, L_i)$$

$H_{\delta_t,j,i} = h_{\delta_t,j}(A_i, L_i)$ where for each l, $h_{\delta_t,j}(\cdot, l)$ is the inverse of the function $q_{\delta_t,j}(\cdot, l)$ and

$$w_{\delta_t,j,i}(\theta) = \frac{f_{A|L}(H_{\delta_t,j,i}|L_i;\theta)}{f_{A|L}(A_i|L_i;\theta)}.$$

If $\Phi^{-1}\{g(Z_i, \delta_t; \beta)\}$ is as in (18), and we choose $s(Z, \delta_t)$ as in (20), then we can compute $\widehat{\beta}_{IPW}$ using the aforementioned extended dataset by simply fitting a weighted logistic regression with outcome Y_i (the same for each replicate record of subject i), covariates Z_i, D_{it} and $Z_i D_{it}$ and weights $\widehat{\lambda}_{i,t} = \lambda_{\delta_t}(A_i, L_i; \widehat{\theta})$.

In the Supporting information, we show that under regularity conditions, (19) has a solution $\widehat{\beta}_{IPW}$ such that $\sqrt{n}(\widehat{\beta}_{IPW} - \beta^*)$ converges to a mean-zero Normal random variable when C0, C2, C3 and C4 hold for all $\delta \in \Upsilon$ and model (12) is correctly specified. In the Supporting information, we provide the expression for the variance of the limiting Normal distribution and a consistent estimator for it. The bootstrap can also be used as an alternative to the variance estimator given in the Supporting information.

Finally, the third estimator, denoted $\widehat{\beta}_{DR}$, generalizes the doubly robust estimator of Section 3.3, and it is computed as the result of the following three-stage procedure:

Step 1: Compute $\widehat{\gamma}$ by solving

$$\sum_{i=1}^{n}\sum_{t=1}^{T} s(Z_i, \delta_t)\lambda_{\delta_t}(A_i, L_i; \widehat{\theta})\left[Y_i - \Phi^{-1}\left\{r(A_i, L_i; \widehat{\tau}) + \gamma^T s(Z_i, \delta_t)\lambda_{\delta_t}(A_i, L_i; \widehat{\theta})\right\}\right] = 0.$$

Step 2: For each unit i in the sample and each $t = 1, \ldots, T$, define

$$\widetilde{m}_{i,t} = \Phi^{-1}\left\{r(Q_i, L_i; \widehat{\tau}) + \widehat{\gamma}^T s(Z_i, \delta_t)\lambda_{\delta_t}(Q_i, L_i; \widehat{\theta})\right\}.$$

Step 3: Compute $\widehat{\beta}_{DR}$ by solving

$$\sum_{i=1}^{n}\sum_{t=1}^{T} s(Z_i, \delta_t) I_{L_i}(A_i)\left[\widetilde{m}_{i,t} - \Phi^{-1}\{g(Z_i, \delta_t; \beta)\}\right] = 0.$$

The estimator $\widehat{\beta}_{DR}$ can be easily computed if $\Phi^{-1}\{g(Z_i, \delta_t; \beta)\}$ is as in (18) and $s(Z_i, \delta_t)$ is as in (20) as follows. First, in the aforementioned extended dataset, we add to replicate t of subject i, new variables $X_{0,i.t} = r(A_i, L_i; \widehat{\tau})$, $X_{1,i,t} = \widehat{\lambda}_{i,t}$, $X_{2,i,t} = \widehat{\lambda}_{i,t}\delta_t$, $X_{3,i,t} = \widehat{\lambda}_{i,t} Z_i$ and $X_{4,i,t} = \widehat{\lambda}_{i,t} Z_i \delta_t$. Next, the estimator $\widehat{\gamma} = (\widehat{\gamma}_1, \ldots, \widehat{\gamma}_4)$ of step 1 is obtained as the weighted logistic regression estimator of outcome Y_i (same outcome for each replicate t of subject i) on covariates $X_{1,i,t}, \ldots, X_{4,i,t}$ with an offset $X_{0,i.t}$, no intercept, an weight $\widehat{\lambda}_{i,t}$. We then store the variable $\widetilde{m}_{i,t}$, and finally, the estimator $\widehat{\beta}_{DR}$ is the logistic regression estimator of outcome $\widetilde{m}_{i,t}$ on covariates Z_i, D_{it} and $Z_i D_{it}$ and an intercept computed from the subset of the extended dataset corresponding to replicates of units i, which satisfy $c(L_i) < A_i < d(L_i)$.

In the Supporting information, we show that under regularity conditions, the resulting estimating equation has a solution $\widehat{\beta}_{DR}$ that satisfies that $\sqrt{n}(\widehat{\beta}_{DR} - \beta^*)$ converges to a mean-zero Normal random variable when C0, C2, C3 and C4 hold for all q_δ and model (8) or (12), but not necessarily both, hold. In the Supporting information, we provide the expression for the variance of the limiting Normal distribution and a consistent estimator for it. The bootstrap can also be used as an alternative to the variance estimator given in the Supporting information.

5. Application: Operative time and post-surgical outcomes

5.1. Setting

The data that motivated this work were abstracted from clinical databases at Brigham and Women's Hospital (Boston, MA) and consist of records on $N = 707$ individuals who underwent surgical resection

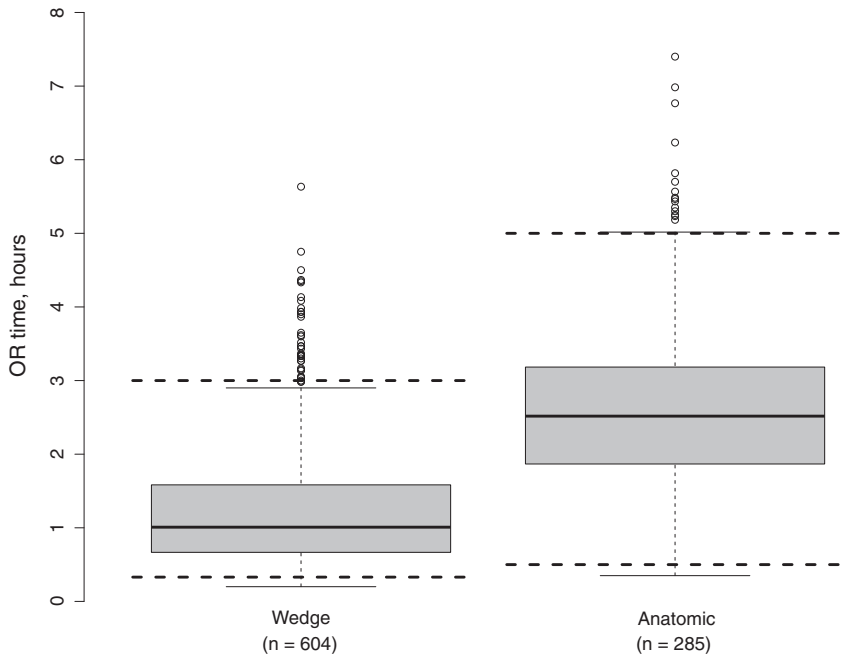

Figure 1. Boxplots of observed operating times, by resection type. Horizontal dashed lines indicate the $(c(L), d(L))$ range restrictions.

for NSCLC between January 2004 and December 2008. Additional inclusion criteria were as follows: patients were between 21 and 85 years of age at the time of surgery, had a pre-operative forced expiratory volume in 1 s (FEV_1) of less than 150% predicted and had a tumor that was less than 20 mm in diameter.

Patients who undergo lung tumor resection have a number of surgical options. The vast majority of patients receive one of three resection types: (i) lobectomy, where the lobe that contains the tumor (one of two that healthy individuals have) is removed; (ii) segmentectomy, where the segment that contains the tumor (one of five that healthy individuals have) is removed; and (ii) wedge resection, where the tumor and a minimal amount of surrounding lung tissue are removed. In this application, we combine lobectomies and segmentectomies into a single group of 'anatomic' resections. A second dimension to choice of surgery is whether a patient undergoes a thoracotomy or a minimally invasive video-assisted thoracoscopic surgery.

5.2. *Available data and analyses*

Interest in this study lies in evaluating the effect of operative time, defined as the time from initial incision to the time at which the surgeon was ready to close, on two post-surgical outcomes: (a) the occurrence of any post-surgical major complications during their hospital stay and (b) a hospital stay of at least 1 week. Confounders, identified *a priori*, included the following: surgery type (thoracotomy vs. video-assisted thoracoscopic surgery), resection type (wedge vs. anatomic), whether or not nodes were sampled, chronic obstructive pulmonary disease, body mass index and smoking status.

To estimate the causal effect of operative time on the two post-surgical outcomes, we considered an marginal structural model (MSM) of the same form as that given by (18) with Z taken to be resection type and $q_\delta(a, l)$ defined as in expression (3) with δ up to 15 min and $b(L) = 45/105$ min for wedge/anatomic resections. The upper bound of 15 min for δ was chosen in consultation with the surgeons in our consulting project so as to reflect the range of duration reductions considered ethically and practically feasible. For both outcomes, $(c(L), d(L))$ was set to (20, 180) min for wedge resection surgeries and (30, 300) min for anatomic resection surgeries (Figure 1). With these choices, we emphasize that the interpretation of the results pertain to the population defined by the constraint that the actual operative time, A, was in the interval $(c(L), d(L))$ but with the intervention (i.e., a reduction of operative time of δ minutes) only having been applied to the subpopulation with $A > b(L) + \delta$.

In the implementation of the analyses, each of the confounders was included in both the outcome regression model as well as a linear regression model for the treatment (i.e., operative time). For the

Table I. Point estimates and bootstrap-based approximate 95% confidence intervals (CI) for the causal odds ratio corresponding to a 15-min reduction in operative time for two outcomes in the analysis of the lung surgery data.

	Wedge resection		Anatomic resection	
	Est.	95% CI*	Est.	95% CI*
At least one major complication				
OR	0.77	(0.66, 0.89)	0.81	(0.68, 0.98)
IPW	0.89	(0.61, 1.29)	0.69	(0.51, 0.92)
DR	0.87	(0.66, 1.16)	0.71	(0.58, 0.88)
Hospital stay of ≥7 days				
OR	0.80	(0.72, 0.89)	0.82	(0.70, 0.96)
IPW	0.96	(0.72, 1.27)	0.71	(0.54, 0.92)
DR	0.91	(0.74, 1.12)	0.74	(0.61, 0.90)

Results are for a marginal structural model given by expression (18) with Z indicating resection type.
*Based on the bootstrap, with 10,000 bootstrap samples.

former, we also included interaction terms between each of the confounders and treatment. Further, in all models, surgeon was adjusted for via the inclusion of a surgeon-specific fixed effect. Finally, for all the IPW and doubly robust estimators, we took the conditional treatment distribution to be a Normal distribution with a constant variance.

5.3. *Results*

Table I presents point estimates and bootstrap-based approximate 95% confidence intervals for the causal odds ratio corresponding to a 15-min reduction in operative time, based on the marginal structural model given by expression (18) with Z indicating resection type. Based on the outcome regression estimator, the odds of at least one major complication decrease by 23% ($\exp\{\widehat{\beta}_{1,OR}\} = 0.77$; 95% CI: 0.66, 0.89) for every 15-min reduction in operative time among wedge resections. The point estimates based on the IPW and doubly robust estimators indicate somewhat of a decreased effect, neither of which are statistically significant. In contrast, both of these estimators indicate a relatively strong causal effect of a 15-min reduction in operative time among anatomic resection surgeries: $\exp\{\widehat{\beta}_{1,IPW} + \widehat{\beta}_{3,IPW}\} = 0.69$ (95% CI: 0.51, 0.92) and $\exp\{\widehat{\beta}_{1,DR} + \widehat{\beta}_{3,DR}\} = 0.71$ (95% CI: 0.58, 0.88), respectively. Again, the results and conclusions for the second outcome of a hospital stay of ≥7 days are similar.

Figure 2 presents estimates of $\mathrm{E}[Y_\delta | c(L) < A < d(L), Z]$ as a function of δ based on our adopted MSM; the first row for the major complication outcome and the second for the long length of stay outcome; the first column for wedge resections and the second for anatomic resections. Consistent with the point estimates in Table I, in all four panels, the estimated expected outcome rates decrease as a function of δ, with the greatest reductions among the anatomic resection surgeries. For example, given a 15-min reduction in operative times among patients who underwent an anatomic resection and whose surgery was at least 105 min long, the overall expected rate of a hospital stay of at least 7 days decreases from approximately 28.5% to 25.5%. We note that, throughout, the doubly robust estimates can be seen to be a compromise between the outcome regression and IPW estimates. Finally, while not presented in detail, we also investigated the inclusion of additional polynomial terms into the MSM; none were statistically significant indicating no evidence of non-linearity in the MSM as a function of δ.

5.4. *The treatment-variation irrelevance assumption*

As with all causal inference, the ability to learn about any particular treatment intervention relies on the interplay between the availability of data that is specific to the intervention and assumptions one makes to allow the borrowing of strength from other observed interventions. In our application, the estimated causal effect associated with a decrease in operating time of δ minutes does not correspond to any single intervention. Yet, as discussed in [27] and [14], we can interpret the target estimand Δ as quantifying the effect of an intervention that reduces the operative time by exactly δ minutes of each subject that had a operative time of A minutes by means of a random mechanism chosen among those used to attain a

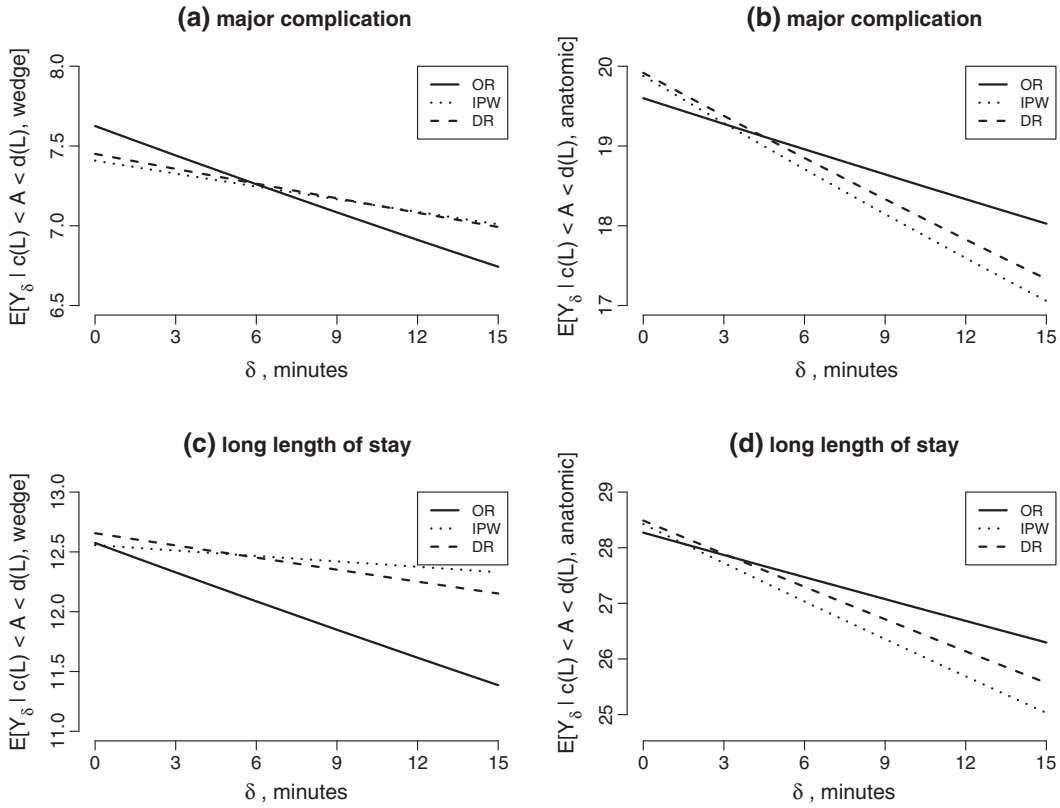

Figure 2. Estimates of $\mathrm{E}[Y_\delta | c(L) < A < d(L), Z]$ based on the adopted MSM of Section 5.2 for the two post-surgical outcomes, stratified by resection type (Z).

time of $A - \delta$ minutes in the study population. As discussed in Section 1, while such estimands cannot directly inform current clinical decision making, they are useful nonetheless as they can be used to guide the choice of promising interventions that can be evaluated in randomized studies.

6. Discussion

In this article, we propose a framework for conducting causal analysis, which relies on estimation of the outcome mean under a hypothetical world in which every subject receives a deterministic function of the dose that they actually received, and on changes in the outcome mean as this deterministic function varies. This setting is different from the standard analytic framework that contemplates hypothetical worlds where the dose received by every individual can depend at most on baseline covariates, but it is otherwise the same for all subjects.

We have argued that our framework is appealing in settings like the one that motivated this work, where the set of feasible treatments for each subject depends on attributes that are not fully captured by baseline covariates but which are reasonably captured by the treatment actually received.

Identification of the target parameters contemplated in our framework requires weaker exchangeability conditions than the ones required for identification of the target parameters of standard causal analyses. Under these identification conditions, the contrast that compares the effect of one specific modified treatment policy to no intervention coincides with a statistical parameter, which was recently studied by Diaz Munoz and van der Laan [4]. These authors show that under certain identifying conditions, this statistical parameter is equal to a causal contrast that compares the effect of a stochastic intervention to no intervention. As such, our work offers a different interpretation of the statistical parameter and estimators in that article. In addition, we have described a marginal structural model whose parameters quantify how the outcome mean changes with different modified treatment policies.

We have derived three different estimators, an IPW, an outcome regression and a doubly robust estimator, of the contrast quantifying the effect of a single intervention and of the parameters of the marginal structural mean model. Our doubly robust estimator is non-iterative, can be implemented easily

using standard software and always falls in the parameter space. In addition, we have provided analytic formulae for the variance of the limiting distribution of our estimators and for consistent estimators of it. Our variance estimators adjust for estimation of the treatment distribution, correctly accounting for the fact that the statistical parameters are functionals of the treatment distribution. In contrast, the variance estimators proposed by Munoz and van der Laan [4] do not account for this fact and can, contrary to what is claimed in that paper, be anticonservative. In the Supporting information, we provide an example which proves this.

Acknowledgements

Andrea Rotnitzky was partially funded by grant 2 R37 AI032475-16A 1 from the National Institutes of Health.

References

1. Rubin D. Estimating causal effects of treatment in randomized and nonrandomized studies. *Journal of Educational Psychology* 1974; **66**:688–701.
2. Gill R, Robins J. Causal inference for complex longitudinal data: the continuous case. *Annals of Statistics* 2001; **29**:1785–1811.
3. Hirano K, Imbens G. The propensity score with continuous treatments. In *Applied Bayesian Modeling and Causal Inference from Incomplete-data Perspectives*, Gelman A, Meng X (eds). John Wiley & Sons: Chichester, UK, 2004; 73–84.
4. Diaz Munoz I, van der Laan M. Population intervention causal effects based on stochastic interventions. *Biometrics* 2012; **68**:541–549.
5. Robins J. A new approach to causal inference in mortality studies with sustained exposure periods: applications to control of the healthy worker survivor effect. *Mathematical Modeling* 1986; **7**:1393–1512.
6. Robins J. Optimal structured nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium on Biostatistics*, Lin DY, Heagerty P (eds). Springer: New York, 2004; 189–326.
7. Murphy S. Optimal dynamic treatment regimes (with discussion). *Journal of the Royal Statistical Society, Series, B* 2003; **58**:331–366.
8. Moodie E, Richardson T, Stephens D. Demystifying optimal dynamic treatment regimes. *Biometrics* 2007; **63**:447–455.
9. Robins J, Hernan M, Seibert U. Estimation of the effects of multiple interventions. In *Comparative Quantification of Health Risks: Global and Regional Burden of Disease Attributable to Selected Major Risk Factors*, Vol. 1, Ezzati M, Lopez A, Rodgers A, Murray C (eds). World Health Organization: Geneva, Switzerland, 2004; 2191–2230.
10. Taubman S, Robins J, Mittleman M, Hernan M. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *International Journal of Epidemiology* 2009; **38**:1599–1611.
11. Shpitser I, Pearl J. Effects of treatment on the treated: identification and generalization. *Uncertainty in Artificial Intelligence*, Montreal, QC, Canada, 2009; 514–521.
12. Rubin D. Which ifs have causal answers? *Journal of the American Statistical Association* 1986; **81**:961–962.
13. Cox D. *Planning of Experiments*. John Wiley & Sons: New York, New York, 1958.
14. VanderWeele T. Concerning the consistency assumption in causal inference. *Epidemiology* 2009; **20**:880–833.
15. Rotnitzky A, Lei Q, Sued M, Robins JM. Improved double-robust estimation in missing data and causal inference models. *Biometrika* 2012; **99**:439–456.
16. Robins J, Rotnitzky A. Comment on the article 'Inference for semiparametric models: Some questions and an answer'. *Statistica Sinica* 2001; **11**:920–936.
17. Hahn J. On the role of the propensity score in efficient semi parametric estimation of average treatment effects. *Econometrica* 1998; **66**:315–331.
18. Hirano K, Imbens G, Ridder G. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 2003; **71**:1161–1189.
19. Lei Q. Improved double-robust estimation of missing data and causal inference models and efficient estimation of the average treatment effect on the treated. *Ph.D. Thesis*, Department of Biostatistics, Harvard School of Public Health, 2011.
20. van der Laan M, Robins J. *Unified Methods for Censored Longitudinal Data and Causality*. Springer: New York, 2003.
21. Robins J, Rotnitzky A, Zhao L-P. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* 1994; **89**:846–866.
22. Pierce D. The asymptotic effect of substituting estimators for parameters in certain type of statistics. *Annals of Statistics* 1982; **10**:475–478.
23. Henmi M, Eguchi S. A paradox concerning nuisance parameters and projected estimating functions. *Biometrika* 2004; **91**:929–941.
24. van der Laan M, Peterson M. Causal effect models for realistic individualized treatment and intention to treat rules. *International Journal of Biostatistics* 2007; **3**.
25. Robins J, Orellana L, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine* 2008; **27**:4678–4721.
26. Orellana L, Rotnitzky A, Robins J. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I: main content. *International Journal of Biostatistics* 2010; **2**.
27. Taubman S, Robins J, Mittleman M, Hernan M. Alternative approaches to estimating the effects of hypothetical interventions. *Proceedings of the 2008 Joint Statistical Meeting*, Denver, Colorado, 2008.