Katherine Daignault* and Olli Saarela

Doubly Robust Estimator for Indirectly Standardized Mortality Ratios

https://doi.org/10.1515/em-2016-0016

Abstract: Routinely collected administrative and clinical data are increasingly being utilized for comparing quality of care outcomes between hospitals. This problem can be considered in a causal inference framework, as such comparisons have to be adjusted for hospital-specific patient case-mix, which can be done using either an outcome or assignment model. It is often of interest to compare the performance of hospitals against the average level of care in the health care system, using indirectly standardized mortality ratios, calculated as a ratio of observed to expected quality outcome. A doubly robust estimator makes use of both outcome and assignment models in the case-mix adjustment, requiring only one of these to be correctly specified for valid inferences. Doubly robust estimators have been proposed for direct standardization in the quality comparison context, and for standardized risk differences and ratios in the exposed population, but as far as we know, not for indirect standardization. We present the causal estimand in indirect standardization in terms of potential outcome variables, propose a doubly robust estimator for this, and study its properties. We also consider the use of a modified assignment model in the presence of small hospitals.

Keywords: quality indicators, causal inference, indirect standardization, direct standardization, doubly robust estimation, provider profiling

1 Introduction

Institutional comparisons have become popular in recent years as a means of assessing the care levels of hospitals for the purpose of resource allocation and policy decisions. With the increasing availability of large administrative databases comprising patient data from multiple hospitals, there is a need for reliable statistical methods for such comparisons that address issues common to these data formats.

In this paper, we consider statistical methods for institutional comparisons for binary outcomes resulting in proportion-type quality indicators, such as the proportion of patients treated with a particular procedure, or proportion experiencing complications from a treatment procedure. In particular, we focus on comparisons made using the standardized mortality ratio (SMR), a ratio of observed to expected outcomes. As patients can not be randomized to hospitals for treatment, adjustment for case-mix must be made since for instance high volume hospitals may also receive more complex cases (Shahian and Normand, 2008). Such adjustment can be made through standardization where the choice between direct or indirect methods depends on the particular comparison of interest. Adopting methods from causal inference, direct standardization can be seen as comparing the potential expected outcomes had all patients in the standard population experienced the care level of a given hospital. Such a comparison would be of particular interest for determining how each hospital might care for the average population. However, policy makers might be interested in how best to allocate resources across the hospital system. In this case, indirect standardization could be more appropriate, as it contrasts the observed outcomes for patients treated in a specific hospital to their potential expected outcomes had these patients experienced the care level of some reference system. In particular, comparing to an average nationwide care level is relevant when the data available capture all hospitals from across

***Corresponding author: Katherine Daignault,** Dalla Lana School of Public Health, University of Toronto, Toronto, Ontario, Canada, E-mail: katherine.daignault@mail.utoronto.ca
http://orcid.org/0000-0003-3665-9184
Olli Saarela, Dalla Lana School of Public Health, University of Toronto, Toronto, Ontario, Canada, E-mail: olli.saarela@utoronto.ca

the country, and thus standardization is relative to a nationwide average standard level of care. An example would be assessing the quality of surgical care for rectal cancer using the positive margin proportion as the quality indicator (Massarweh et al., 2014) and data from the National Cancer Data Base (NCDB, Raval et al., 2009), which captures hospitals across the United States.

Regardless of the institutional comparison of interest, statistical adjustment is required to attempt to ensure that any differences in indicators are due solely to differences in actual institutional performance. One such method is to calculate the propensity score for each hospital (Shahian and Normand, 2008). By determining the probability of being treated at each institution based on patient characteristics (i.e. the propensity score), it is then possible to compare the observed outcomes using the propensity score through simple matching or stratification, or through weighting in regression modelling. Risk adjustment can also be implemented using outcome models (Spiegelhalter, 2005), which directly give the expected outcome conditional on patient characteristics. By summing over the patients in each hospital, the expected outcome needed for the SMR is obtained. A comprehensive summary of the evolution of standardization methods can be found in Keiding and Clayton (2014).

A common issue that arises in any modelling scenario is that of model misspecification, which can be due to a number of reasons, including unmeasured confounders, omission of observed confounders in the model, and misspecification of the functional form of relationships. When making institutional comparisons in an effort to identify under/over-performing hospital practices, model misspecification can have a potentially serious effect on the classification of these hospitals as outliers. An attempt to overcome or at least alleviate such issues is to use doubly robust (DR) methods that incorporate both the propensity score and the outcome model into a single estimator (Bang and Robins, 2005; Funk et al., 2011). We propose a DR estimator for the SMR under indirect standardization, where the causal quantity being estimated is specified through the expected potential outcome had the patients treated in a given hospital experienced a system-wide average level of care. In the context of institutional quality comparisons, a DR estimator has been proposed for direct standardization (Varewyck et al., 2014). In addition, Shinozaki and Matsuyama (2015) propose a DR estimator for standardized risk differences and ratios in the exposed population. While the intended use of their estimator was not for the purpose of institutional comparisons, it may be adopted in order to make pairwise comparisons between hospitals, namely to estimate the expected potential outcome had patients treated in hospital A been treated in hospital B. The causal comparison being made in Shinozaki and Matsuyama (2015) differs from the proposed estimator to follow as we attempt to compare each hospital to an average level of care in a healthcare system instead of to a given reference hospital's level of care.

The paper proceeds as follows. In Section 2 we review the ideas of direct and indirect standardization, and specify the causal estimand in indirect standardization using potential outcomes notation. The proposed DR estimator for the SMR under indirect standardization is developed and shown to be consistent. Simulation study results demonstrating the doubly robust property of the proposed estimator are presented in Section 3. A discussion follows in Section 4.

2 Proposed estimator

2.1 Notation and assumptions

For the extent of the paper, $Y \in \{0, 1\}$ is the observed binary quality outcome variable, $Z \in \{1, \ldots, m\}$ is the hospital in which the patient was actually treated, and $X \equiv (X_1, \ldots, X_p)$ is a vector of patient-level characteristics relevant to case-mix adjustment, capturing for example demographic information, medical history, and disease progression. The triples $W \equiv (Y, Z, X)$ are assumed independent and identically distributed across the patients. As is the convention in the causal inference literature, we denote by Y_z the potential outcome that would have been observed had the patient been treated in hospital z. Throughout, we make the following standard causal assumptions. We assume that X is sufficient to control for any confounding (conditional exchangeability) so that $(Y_1, \ldots, Y_m) \perp\!\!\!\perp Z \mid X$. In addition, we assume consistency, under which the

observed outcome is determined by $Y = \sum_{z=1}^{m} \mathbf{1}_{\{Z=z\}} Y_z$ (Hernan and Robins, 2006), where $\mathbf{1}_{\{Z=z\}}$ is the indicator function which takes on values $\{0, 1\}$ depending on if the condition is false or true respectively. Finally, we assume positivity, under which all patients have a non-zero probability of being treated at any hospital, i.e. $P(Z = z \mid X) > 0$ for all $z \in \{1, \ldots, m\}$ and X combinations.

2.2 Direct versus indirect standardization

In this section, we briefly review the two common standardization procedures used in epidemiology, and discuss their causal interpretation in the quality comparison context. The main difference between the two standardization methods is that direct standardization provides the expected outcome if the standard population were to experience the event rate observed in the index population, whereas indirect standardization provides the expected outcome if the index population had experienced the event rate from the standard population. Table 1 provides an illustration of the different elements from each population used to compute the expected outcome. Each standardization method computes the expected outcome by considering a covariate stratum-specific (e.g. age, gender) event rate applied to a stratum-specific population size.

Table 1: Difference between the standardization methods; The asterisk refers to the standard population, k indicates the covariate strata, $\hat{\pi}_k$ is the estimated event rate, and E is the expected outcome.

Method	Standard population	Index population	Expected outcome
Direct	n_k^*	$\hat{\pi}_k$	$E = \frac{1}{\sum_k n_k^*} \sum_k n_k^* \hat{\pi}_k$
Indirect	$\hat{\pi}_k^*$	n_k	$E = \frac{1}{\sum_k n_k} \sum_k n_k \hat{\pi}_k^*$

Direct standardization, as seen in the first row of Table 1, takes the event rate of the index population and applies it to the standard population in each strata, and then averages over all covariate strata in the standard population, resulting in $E = \left(\sum_k n_k^*\right)^{-1} \sum_k n_k^* \hat{\pi}_k$. Direct standardization assumes that the stratum membership in the standard population is known and the stratum-specific rates must be estimated from the index population. In the present context, the index population are patients treated in a given hospital, while the standard population may be another hospital, or in the case of nationwide comparisons, the entire patient population across all hospitals. The case-mix adjustment required for the quality comparisons usually involves a large number of covariate strata, and therefore the stratum-specific event rates are in practice found using regression modelling techniques. The causal estimand under direct standardization where the standard population is all hospitals nationwide, can be written as $E[Y_z]$ as in Varewyck et al. (2014), which under the causal assumptions of Section 2.1 can be expressed as $E[Y_z] = \sum_x E[Y \mid Z = z, x]P(X = x)$ (e.g. Hernan and Robins, 2006). This can be interpreted as the expected outcome had all patients experienced the care level of hospital z.

In contrast, indirect standardization takes the event rate in each covariate stratum in the standard population, applies it to the stratum membership in the index population, and then averages over the membership of all strata in the index population, as $E = \left(\sum_k n_k\right)^{-1} \sum_k n_k \hat{\pi}_k^*$ (see row 2 of Table 1). More conceptually, indirect standardization can be thought of as determining the expected outcome if patients in the index population were to experience the same event rate as the standard population. In this case, the stratum-specific event rates are obtained from the standard population, while the stratum membership is determined from the study index population. Because the stratum-specific event rates do not need to be estimated from the index population, indirect standardization is still applicable when the index population is small, provided that the standard population is large. Finally, the expected event counts are contrasted to the observed ones through the standardized mortality ratio $SMR = O/E$.

The specific causal comparison being made in indirect standardization is determined by the choice of standard population. Suppose for instance that the comparison of interest is how patients from hospital z (index population) would fare if they were instead treated at hospital z' (standard population). The causal

estimand for indirect standardization must feature a conditional expectation of the form $E[\cdot \mid Z = z]$. When comparing two hospitals, this becomes simply $E[Y_{z'} \mid Z = z]$, with SMR $= E[Y_z \mid Z = z]/E[Y_{z'} \mid Z = z]$. The latter corresponds to the exposure effect among the exposed risk ratio discussed by Shinozaki and Matsuyama (2015). However, to express the causal estimand in comparison to the nationwide average care level, instead of fixing the subscript of the potential outcome, we need to consider this as a random variable. This corresponds to a hypothetical intervention of randomly assigning a patient actually treated in hospital z to be treated in one of the hospitals under comparison. To express this, as a notational device, we define a new random variable A to denote the unobserved potential hospital assignment, where $A \in \{1, \ldots, m\}$. Further, we let $(Y_1, \ldots, Y_m) \perp\!\!\!\perp A \mid (Z, X)$ and $A \perp\!\!\!\perp Z \mid X$ so that we have the causal relationships presented in Figure 1. We can now generally express the causal estimand in indirect standardization through the conditional expectation $E[Y_A \mid Z = z]$. Several interesting special cases may be obtained by choosing the hypothetical assignment probabilities $P(A \mid X)$. The usual exposure effect among the exposed comparison would be obtained by taking $P(A = z' \mid X) = 1$. Comparison to the care level of an average provider would be obtained by taking $P(A = a \mid X) = 1/m$ for all $a \in \{1, \ldots, m\}$ (see Section 2.4 and Varewyck et al., 2014). However, herein we are specifically interested in comparisons to nationwide care level. In Section 2.4 we show that the corresponding causal estimand is specified by choosing the hypothetical assignment probabilites as equal to the actual ones, effectively weighting the hospitals in the average by their patient volumes.

Regardless of the standardization method used, it is necessary to estimate the respective expected potential outcomes. When the number of covariate strata becomes too large, it is common to fit an outcome model to the data and to estimate the particular expected number from the fitted values. However, such estimates are subject to the possibility of bias from model misspecification due to any number of factors. One attempt to protect against misspecification of the outcome model used in estimation is to instead use a doubly robust estimator.

2.3 Doubly robust estimation in direct standardization

Doubly robust (DR) estimation attempts to eliminate bias due to misspecification of a single model by utilizing two separate models in the estimation process (Funk et al., 2011). By doing so, it incorporates as much information about the causal pathway between the outcome, the exposure and the covariates/confounders as possible (see Figure 1). DR estimators combine the use of an outcome model, $m(X, z, \phi) \equiv E[Y \mid Z = z, X; \phi]$, and a propensity/assignment model, $e(X, z, \gamma) \equiv P(Z = z \mid X; \gamma)$, parametrized with respect to ϕ and γ respectively, into an estimator such that, as long as one of the models is correctly specified, the results should be unbiased or at least consistent (Bang and Robins, 2005).

In general, DR estimators of a mean effect are composed of three terms that contain either the outcome model, the propensity/assignment model, or both such that the terms that contain the misspecified model will cancel thereby resulting in estimation using only the correct model. Therefore, making use of fitted outcome and assignment probabilities $m(X, z, \hat{\phi})$ and $e(X, z, \hat{\gamma})$, where estimated model parameters were

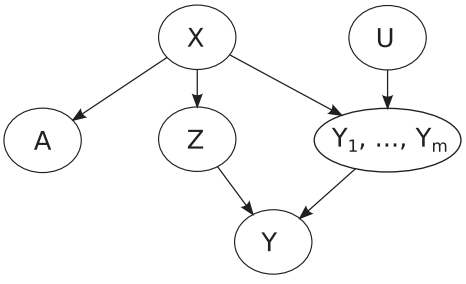

Figure 1: The postulated causal mechanism (U is a non-confounder latent variable representing the correlation between potential outcomes for an individual).

denoted by $\hat{\phi}$ and $\hat{\gamma}$ respectively, a DR estimator (Robins et al., 2007) for a marginal mean μ_z under direct standardization can be written as

$$\hat{\mu}_z^{\mathrm{DR}} = n^{-1}\sum_{i=1}^{n} m(x_i, z, \hat{\phi}) + n^{-1}\sum_{i=1}^{n} \frac{\mathbf{1}_{\{Z_i=z\}}}{e(x_i, z, \hat{\gamma})}\left[Y_i - m(x_i, z, \hat{\phi})\right] \tag{1}$$

$$= n^{-1}\sum_{i=1}^{n} \frac{\mathbf{1}_{\{Z_i=z\}}}{e(x_i, z, \hat{\gamma})} Y_i + n^{-1}\sum_{i=1}^{n}\left[1 - \frac{\mathbf{1}_{\{Z_i=z\}}}{e(x_i, z, \hat{\gamma})}\right] m(x_i, z, \hat{\phi}). \tag{2}$$

Here the outcome model estimator, with hospital effect, has been augmented by weighting by the inverse of the assignment probability. An estimator of this form is doubly robust if either the outcome or the assignment/propensity model is correctly specified. This is evident as the second term in eq. (1) will in large samples have mean zero if the outcome model is correctly specified, leaving the first term to provide the estimate, while the second term of eq. (2) will also in large samples have mean zero if the propensity/assignment model is correctly specified and thus leaving only the propensity model to provide the estimate. Varewyck et al. (2014) used such an estimator for estimating the potential full population risk, $E[Y_z]$, had all patients received the care level of hospital z, a directly standardized quantity.

An issue that arises in direct standardization is the need to specify hospital effects in the outcome model. In the case of a large nationwide database of hospitals, some of them small in volume, this requires the estimation of a large number of parameters which might not be feasible without smoothing/shrinkage. Although such smoothing could be employed through mixed effect models, this might be a questionable approach if the purpose of the institutional comparison is to identify outliers. Varewyck et al. (2014) discuss possible ways to reduce shrinkage, such as clustered mixed effect models on hospitals or Firth corrected fixed effects logistic regression as the outcome model.

Additionally, direct standardization raises the question as to whether modelling hospital-patient interactions would also be needed, especially in the case of hospitals that specialize in particular patient subgroups, such as children or the elderly (Varewyck et al., 2016). Including interaction terms would further contribute to the large number of parameters that require estimation in the outcome model. Varewyck et al. (2016) have shown that the omission of hospital-patient interactions in the models used for standardization can contribute bias towards the estimated excess risks.

To sum up, in the case of nationwide comparisons involving hundreds or thousands of hospitals, many of these small volume, direct standardization may be more problematic due to the large number of hospital effects and patient-hospital interaction effects that would need to be estimated to ensure unbiased estimation. However, if comparison to an average level of care as the reference is of interest, indirect standardization avoids the issue of modelling hospital effects as well as hospital-patient interactions, which we will demonstrate in Section 2.4. Nevertheless, indirect standardization still requires specification of an outcome model. We thus propose a doubly robust estimator for the standardized mortality ratio under indirect standardization. In order to do this, we must first express the SMR as a causal estimand.

2.4 Causal estimand under indirect standardization

As per the discussion in Section 2.2, we define the causal estimand for hospital z in indirect standardization as

$$\mathrm{SMR} = \frac{E[Y_z \mid Z = z]}{E[Y_A \mid Z = z]}, \tag{3}$$

where the observed response in the numerator results simply from considering the potential outcomes of the patients of hospital z had they been treated in hospital z (i.e. the consistency assumption). In contrast, the expected response in the denominator depends on the specified target assignment regime, $P(A \mid X)$. As mentioned in Section 2.2, notable special cases may be obtained by choosing the assignment probabilities. First, consider the target assignment regime that gives patients an equal probability of being treated at each

hospital, $P(A = a \mid X) = m^{-1}$. We then show in Appendix A that, for this choice of assignment regime, eq. (3) is equivalent to

$$\text{SMR} = \frac{E[Y_z \mid Z = z]}{m^{-1} \sum_{a=1}^{m} E[Y_a \mid Z = z]}, \tag{4}$$

which is the causal estimand briefly considered by Varewyck et al. (2014). This takes an equally weighted average across all hospitals in the denominator and thus corresponds to using the care level of an average provider as the reference in indirect standardization. In contrast, we want to use the national average level of care as the reference, and choose as the target assignment regime $P(A = a \mid X) = P(Z = a \mid X)$. Then the denominator of eq. (3) involves an average across all hospitals but weighted by their actual volume. For the causal estimand for hospital z under this special case, we introduce the shorthand notation $\theta_z \equiv \text{SMR}$.

Now, utilizing the causal assumptions listed in Section 2.1, and the additional conditional independence properties $(Y_1, \ldots, Y_m) \perp\!\!\!\perp A \mid (Z, X)$ and $A \perp\!\!\!\perp Z \mid X$, it can be shown (Appendix A) that eq. (3) can be expressed in terms of observable quantities as

$$\begin{aligned} \theta_z &= \frac{\sum_x P(Y = 1 \mid X = x, Z = z) P(X = x \mid Z = z)}{\sum_x \sum_a P(Y = 1 \mid X = x, Z = a) P(A = a \mid X = x) P(X = x \mid Z = z)} \\ &= \frac{E[Y \mid Z = z]}{E\left\{E[Y \mid X] \mid Z = z\right\}}, \end{aligned} \tag{5}$$

where the denominator corresponds to using an outcome model without hospital effects, and averaging the predictions from such a model over the patients of hospital z. This is similar to the indirect standardization approach considered by e.g. Faris et al. (2003) and Tang et al. (2015), and is the appropriate modelling approach when the reference is chosen as the average level of care in the health care system. While it depends on the context whether this is the relevant comparison, we will now demonstrate how to obtain a simple doubly robust estimator for the causal SMR under such standardization. We show in Appendix A that similar manipulations of the causal estimand (that resulted in the equivalence of eqs. (3) and (5)) further result in two other equivalent expressions in terms of observable quantities, that is,

$$\theta_z = \frac{E\left[\mathbf{1}_{\{Z=z\}} Y\right]}{E\left[P(Z = z \mid X) Y\right]} \tag{6}$$

and

$$\theta_z = \frac{E\left[\mathbf{1}_{\{Z=z\}} Y\right]}{E\left\{E[Y \mid X] P(Z = z \mid X)\right\}}. \tag{7}$$

Thus, under the causal assumptions, the causal estimand θ_z can be written in terms of observable quantities in three equivalent forms that could be estimated using either an outcome model (eq. (5)), an assignment model (eq. (6)), or a combination of both models (eq. (7)). Therefore, we may utilize all three forms in a doubly robust estimator.

2.5 Proposed doubly robust estimator

As eqs. (5–7) are all equivalent and contain either one or both of an outcome model and propensity/assignment model, we may now write the causal SMR as

$$\theta_z = \frac{E\left[\mathbf{1}_{\{Z=z\}} Y\right]}{E\left[P(Z = z \mid X) Y\right]} + \frac{E[Y \mid Z = z]}{E\left\{E[Y \mid X] \mid Z = z\right\}} - \frac{E\left[\mathbf{1}_{\{Z=z\}} Y\right]}{E\left\{E[Y \mid X] P(Z = z \mid X)\right\}}. \tag{8}$$

This motivates the proposed DR estimator for the SMR of hospital z under indirect standardization

$$\hat{\theta}_z \equiv \frac{\sum_{i=1}^n \mathbf{1}_{\{Z_i=z\}} Y_i}{\sum_{i=1}^n e(x_i, z, \hat{\gamma}) Y_i} + \frac{\sum_{i=1}^n \mathbf{1}_{\{Z_i=z\}} Y_i}{\sum_{i=1}^n \mathbf{1}_{\{Z_i=z\}} m(x_i, \hat{\phi})} - \frac{\sum_{i=1}^n \mathbf{1}_{\{Z_i=z\}} Y_i}{\sum_{i=1}^n m(x_i, \hat{\phi}) e(x_i, z, \hat{\gamma})}. \tag{9}$$

Here $m(x_i, \phi) \equiv E[Y_i \mid X_i = x_i, \phi]$ is an outcome model parametrized in terms of ϕ. In the case of a binary outcome variable, this would be a logistic regression model of the form

$$m(x_i, \phi) \equiv \text{expit}\{\phi_0 + \phi_1' x_i\}, \tag{10}$$

where $\phi \equiv (\phi_0, \phi_1)$. The corresponding parameter estimates are denoted by $\hat{\phi}$. Further, $e(x_i, z, \gamma) \equiv P(Z_i = z \mid X_i = x_i, \gamma)$ is a multinomial logistic assignment probability model parametrized in terms of γ, given by

$$e(x_i, z, \gamma) \equiv \frac{\exp(\gamma_{0z} + \gamma_{1z}' x_i)}{1 + \sum_{a=2}^m \exp(\gamma_{0a} + \gamma_{1a}' x_i)}, \quad z = 2, \ldots, m \tag{11}$$

and $e(x_i, 1, \gamma) = 1 - \sum_{z=2}^m e(x_i, z, \gamma)$, with $\gamma \equiv (\gamma_{02}, \ldots, \gamma_{0m}, \gamma_{12}, \ldots, \gamma_{1m})$ denoting the collection of all the parameters, and the corresponding parameter estimates denoted by $\hat{\gamma}$.

The estimator in eq. (9) is applied in turn to each hospital $z = 1, \ldots, m$, with the parameters ϕ and γ estimated through fitting the regression models in eqs. (10) and (11) to the pooled patient population. We note that the outcome model $m(x_i, \phi)$ no longer contains a term for the hospital effect, (as opposed to $m(x_i, z, \phi)$ in eq. (1)), and thus we are estimating fewer parameters compared to the outcome model used for direct standardization. However, the hospital assignment model requires estimation of hospital-level regression parameters, and thus it is worth considering the case where the observational database may contain information on small hospitals, in which very few patients are being treated. In such situations, there may not be sufficient data to estimate the model parameters for all hospitals in the multinomial assignment model. However, we still want to include all the hospitals in the standardization since the reference is the national average level of care. We therefore also propose a modification of the multinomial assignment model of eq. (11) that only specifies covariate effects for the hospitals that are large enough. Suppose that, out of m hospitals, the first l hospitals are 'small' and the rest are 'large'. We may then pool the small hospitals together as the reference category and specify the multinomial assignment model as

$$e(x_i, z, \gamma) = \begin{cases} \frac{\exp(\gamma_{0z})}{1 + \sum_{a=2}^l \exp(\gamma_{0a}) + \sum_{a=l+1}^m \exp(\gamma_{0a} + \gamma_{1a} x_i)} & \text{for } z = 2, \ldots, l \\ \\ \frac{\exp(\gamma_{0z} + \gamma_{1z} x_i)}{1 + \sum_{a=2}^l \exp(\gamma_{0a}) + \sum_{a=l+1}^m \exp(\gamma_{0a} + \gamma_{1a} x_i)} & \text{for } z = l+1, \ldots, m \end{cases} \tag{12}$$

and $e(x_i, 1, \gamma) = 1 - \sum_{z=2}^m e(x_i, z, \gamma)$. While the assignment model in eq. (12) is obviously misspecified for the small hospitals $z = 2, \ldots, l$, it will still help in estimation of the SMRs for the large hospitals, and thus is an improvement over using only the outcome model, as the estimation of the θ_zs for $z = l+1, \ldots, m$ will be doubly robust.

While the DR estimator in eq. (9) is composed of three terms as required by the form of eqs. (1) and (2), as well as using two models for estimation, it is worth noting that the assignment model is not actually being utilized through inverse probability weighting. Nevertheless, our estimator does in fact have the doubly robust property, namely that $\hat{\theta}_z$ converges in probability to θ_z as long as either the outcome or assignment model is correctly specified, the proof of which can be found in Appendix B. We will demonstrate this property in a simulation study in the following section.

3 Simulation

We now present the results of a simulation study that illustrates the doubly robust property of the proposed estimator $\hat{\theta}_z$ (eq. (9)). To this end, we purposefully kept the number of hospitals in the simulation small. We simulated 1000 datasets according to the causal pathway in Figure 1. Each dataset consists of $n = 1000$ patients that are assigned into $m = 5$ hospitals. Each patient has $p = 2$ measured covariates which are associated with both the hospital assignment and the quality outcome: $X_{1i} \sim N(0, 1)$, a standard normal variable, which, to demonstrate model misspecification, we transform into $V_{1i} = \frac{|X_{1i}|}{\sqrt{1-\frac{2}{\pi}}}$ such that more extreme values of X_{1i} represent increased risk, and $X_{2i} \sim$ Bernoulli(0.5). We also generate another standard normal random variable U_i to represent the similarity among the potential outcomes for each patient (see Figure 1). The binary potential outcomes are generated as

$$Y_{zi} \sim \text{Bernoulli}(\text{expit}(\alpha_{0z} + \alpha_1 V_{1i} + \alpha_2 X_{2i} + \alpha_3 U_i))$$

independently for $z = 1, \ldots, 5$, where the coefficients were chosen as $(\alpha_1, \alpha_2, \alpha_3) = (0.5, 1.5, 1.0)$ for all simulations and $\alpha_0 = (\alpha_{01}, \ldots, \alpha_{05})$, which dictates the level of the true quality of care of each hospital, are chosen according to two different scenarios. In the first, there is no difference in the quality of care between the hospitals, and thus we set $\alpha_0 = (0, \ldots, 0)$, corresponding to SMR=1 for each hospital. In the second, we set $\alpha_0 = (0, -1, 0, 1, 0)$ such that 3 hospitals have SMR near 1 while one hospital has SMR larger than 1 and one has SMR smaller than 1.

The observed hospital assignment for each patient is generated as $Z_i \sim \text{Multinomial}(p_{1i}, \ldots, p_{5i})$, where $p_{zi} = \text{expit}(\beta_{0z} + \beta_{1z} V_{1i} + \beta_{2z} X_{2i})$ for $z = 2, \ldots, 5$ and $p_{1i} = 1 - \sum_{z=2}^{5} p_{zi}$. Here $\beta_0 = (\beta_{02}, \ldots, \beta_{05})$ dictates the volumes of the hospitals, and $\beta_1 = (\beta_{12}, \ldots, \beta_{15})$ and $\beta_2 = (\beta_{22}, \ldots, \beta_{25})$ dictate how the hospital assignment probabilities depend on the patient-level characteristics. We let $\beta_1 = (0, 0, 0.5, 1)$ and $\beta_2 = (-1, -0.5, 0.5, 1)$ for all simulations, while for the hospital volume we consider $\beta_0 = (-1, -0.5, 0.5, 1)$, which results in three small volume and two large volume hospitals. This choice of β_0 results in hospitals 1-3 having average sizes of 57.6, 17.0 and 30.7, while hospitals 4 and 5 have average sizes of 181.8 and 712.9 respectively, for all simulations.

Finally, the observed outcome Y is given by the potential outcome corresponding to the hospital assignment of each patient, as required by the consistency assumption. Although the true SMRs are not directly specified by the parameters in the data generating mechanism, we estimated the true SMRs from the simulated potential outcomes using the definition given by eq. (3) and the true assignment probabilities, averaged over the 1000 simulation rounds.

For each dataset, under these specifications, we compute the SMR for each hospital using the estimators based on both the outcome model (eq. (5)) and the assignment model (eq. (6)) alone and the proposed doubly robust estimator (eq. (9)) when all models are correctly specified. We then misspecifiy each model in turn and then simultaneously, and compare the performance of all three estimators under each scenario. The type of misspecification considered here is that of misspecifying the functional form of a covariate, i.e. using the original untransformed variable X_{1i} in place of V_{1i}.

Figure 2 presents the sampling distribution of the three SMR estimators (based on eqs. (5), (6) and (9)) as well as the true value for this ratio under the scenario where there is no difference between the quality of care. We see that when there is no misspecification of the models (top left panel), the sampling distributions of the three estimators are nearly identical within each hospital. When either the outcome model (top right) or the assignment model (bottom left) alone is misspecified, we see that the doubly robust estimator in both cases produces results similar to the estimator featuring only the correctly specified model, demonstrating the double robustness property, while the estimators featuring only the misspecified model produce biased results. When both models are misspecified, all three estimators are biased, but the doubly robust estimator does not introduce additional bias compared to the other two estimators. In each of the four misspecification

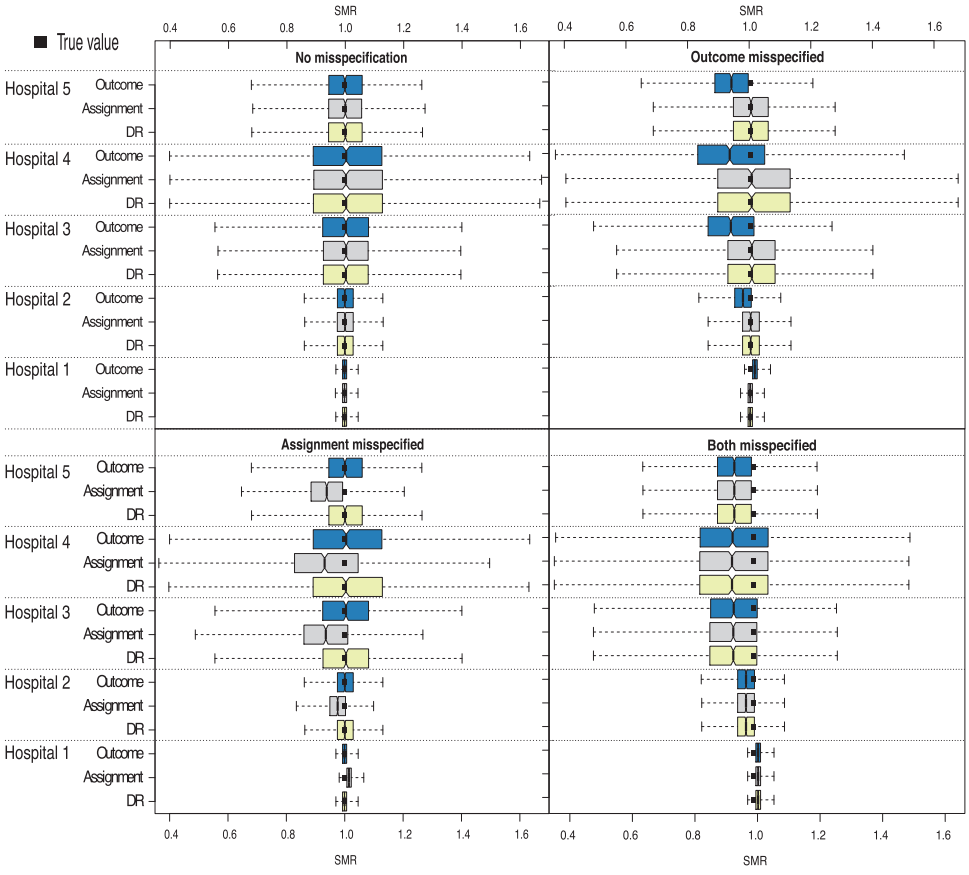

Figure 2: Sampling distributions of observed-to-expected ratios based on outcome model (eq. (5)) only, assignment model (eq. (6)) only and doubly robust (eq. (9)) estimators when true SMR = 1.0 for all hospitals.

scenarios considered in Figure 2, due to their small volume, the sampling distributions of hospitals 1-3 exhibit more variability than those of hospitals 4 and 5.

Figure 3 presents the sampling distributions under the scenario where the SMR is allowed to vary across hospitals. When all models are correctly specified, the three estimators produce a similar sampling distribution of SMRs for each hospital. Once again, when one of the models is misspecified, we see that the doubly robust estimator and the estimator featuring only the correctly specified model produce nearly identical results while the estimator featuring only the misspecified model produces biased results. As expected, when both models are misspecified, all three estimators produce biased estimates.

Although we did not simulate a scenario where there are hospitals so small that the full multinomial assignment model cannot be fit, it is of interest to consider the effect of pooling hospitals in the assignment model, as discussed in Section 2.5. This requires fitting a multinomial logistic model of the form in eq. (12) where only the intercept terms are estimated for hospitals 1, 2 and 3, and both intercept terms and regression coefficients are estimated for hospitals 4 and 5. The results of this scenario are presented in Figure 4. Relative to Figure 3, there is a small difference in the bias of the assignment model based estimator when the models are correctly specified, yet we do not see much difference when it is misspecified. The doubly robust estimator, as expected, consistently estimates the true SMR when the outcome model is correctly specified, despite the presence of the added misspecification to the assignment model. When the outcome model is misspecified and the doubly robust estimator relies on the assignment model for estimation, there is additional bias introduced by the use of the modified multinomial model for hospitals 1, 2, and 3. However, the double robustness property still applies to hospitals 4 and 5, as discussed in Section 2.5.

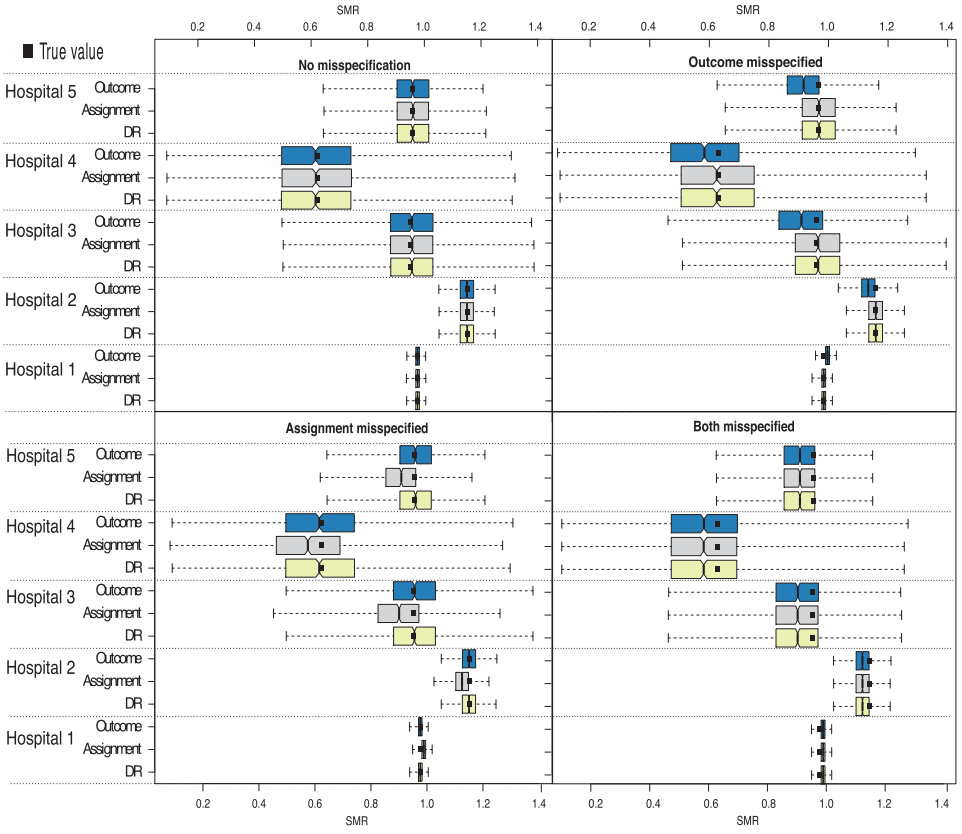

Figure 3: Sampling distributions of observed-to-expected ratios based on outcome model (eq. (5)) only, assignment model (eq. (6)) only and doubly robust (eq. (9)) estimators when true level of care varies across hospitals.

4 Discussion

The doubly robust estimator for the SMR under indirect standardization that we have proposed in eq. (9) has been shown (in Appendix B) to be an asymptotically consistent estimator when either the outcome or the assignment model is correctly specified and we have also demonstrated this property through simulations. The simulation results demonstrated that our proposed estimator is robust to model misspecification of one but not both of the models used for estimation, but performs no worse than the outcome model estimator or the assignment model estimator when both models are misspecified. Some authors have discussed scenarios where doubly robust estimators have the potential to increase bias (Kang and Schafer, 2007). We did not encounter these in our simulations, although as a caution it should be noted that the results of the simulation study apply only to the types of misspecification that we considered.

When small hospitals are present, a modified multinomial assignment model, pooling some of the hospitals, can be used to avoid problems in estimating covariate effects. While the modified assignment model is inherently a type of misspecified model, we see that the bias introduced by its use only concerns the small hospitals, with the proposed estimator still demonstrating the double robustness property for the large hospitals.

The simulation results also demonstrated that there is little difference in the variance of the sampling distributions of the estimated SMRs, regardless of the estimator being used. To explain this, we note that the numerators of the three forms for the causal estimand (eqs. (5), (6) and (7)) can all be estimated through the same quantity, $\sum_{i=1}^{n} \mathbf{1}_{\{Z_i=z\}} Y_i$ which has a binomial sampling variance. Therefore, if the variance resulting from the estimation of the denominator terms is small, the variance of all three estimators will be similar. Secondly, though it is known that inverse probability weighted estimators can be highly variable

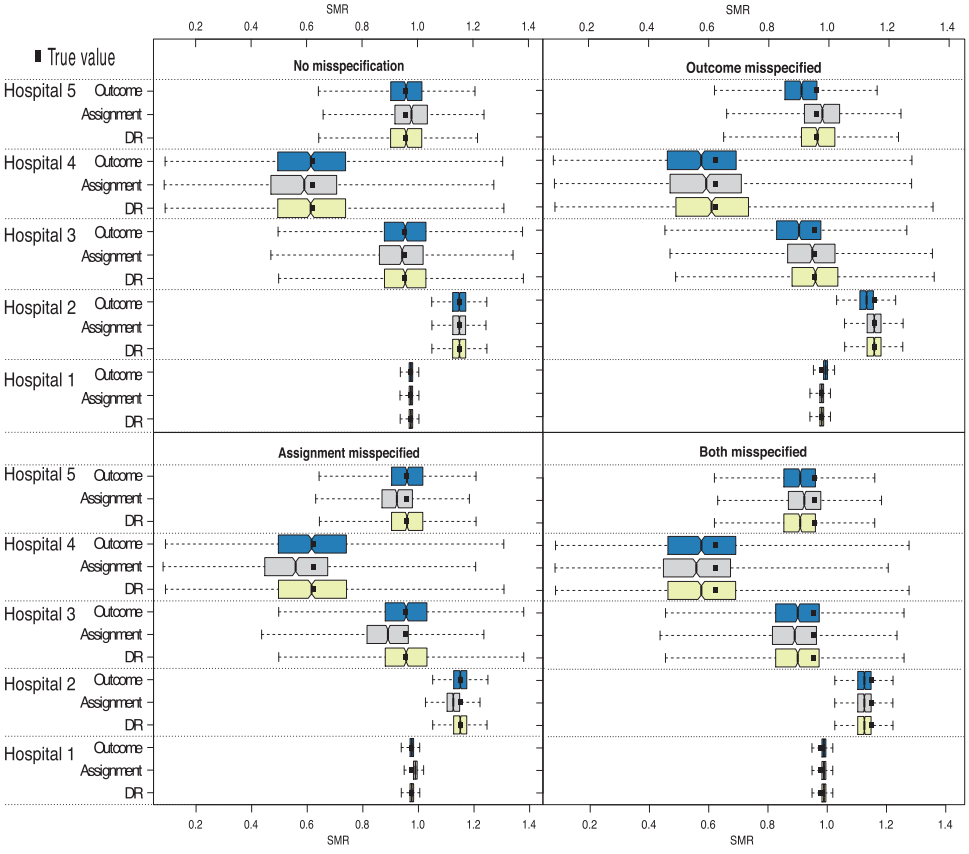

Figure 4: Sampling distributions of observed-to-expected ratios based on outcome model (eq. (5)) only, modified assignment model (eq. (12)) only and doubly robust (eq. (9)) estimators when true level of care varies across hospitals.

even when a correctly specified propensity model is used (Kang and Schafer, 2007), our estimator is not actually employing the assignment model in inverse probability weighting and therefore is not subject to this added variability.

An important consideration to be made is the estimation of the variance of the proposed doubly robust estimator. It is a common practice in indirect standardization to assume that the estimated expected number of events in the SMR contributes no variability to the overall estimate of the SMR and thus confidence intervals are built solely on the variability of the observed counts. Faris et al. (2003) have shown that, by ignoring the modelling error in the estimated expected counts, bias is introduced into the confidence intervals and can result in misclassification of hospitals as outliers. In the case where the expected counts are estimated using a logistic model of the binary outcome on a single risk score that incorporates the information from many patient characteristics deemed relevant, Tang et al. (2015) have proposed an asymptotic distribution for the SMR from which confidence intervals can be obtained. They also show that these asymptotic confidence intervals perform similarly to intervals obtained through a bootstrap procedure. We thus suggest that confidence intervals for the proposed doubly robust estimator be computed via bootstrap, and the derivation of an explicit form for the variance of our estimator is left as future work.

The present methodological work has several possible extensions. One is to consider doubly robust estimators for composite scores based on multiple quality indicators. This is motivated by the fact that policy makers would likely base their decisions for the allocation of funding and resources on multiple dimensions of quality of care. Further, the proposed framework can be generalized to also incorporate within hospital comparisons over time, in addition to between-hospital comparisons in the same time period. Here again we

need to be able to remove the confounding due to changes in the patient population over time, possibly in a doubly robust way.

Appendix A: Proofs for equations (4–7)

Throughout, we make use of the notation and assumptions introduced in Sections 2.1 and 2.2. First, under a general target assignment regime $P(A \mid X)$ we can write

$$\begin{aligned}
\text{SMR} &= \frac{E_{X|Z=z}\left\{E[Y_z \mid X, Z=z]\right\}}{E_{A,X|Z=z}\left\{E[Y_A \mid A, X, Z=z]\right\}} \\
&= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(X=x \mid Z=z)}{\sum_x \sum_a P(Y_a=1 \mid A=a, Z=z, X=x)P(A=a, X=x \mid Z=z)} \\
&= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(X=x, Z=z)}{\sum_x \sum_a P(Y_a=1 \mid Z=z, X=x)P(A=a, X=x, Z=z)} \\
&= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(Z=z \mid X=x)P(X=x)}{\sum_x \sum_a P(Y_a=1 \mid Z=a, X=x)P(A=a, Z=z \mid X=x)P(X=x)} \\
&= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(Z=z \mid X=x)P(X=x)}{\sum_x \sum_a P(Y_a=1 \mid Z=a, X=x)P(A=a \mid X=x)P(Z=z \mid X=x)P(X=x)}.
\end{aligned} \tag{13}$$

In the above, the third equality followed from the conditional independence property $(Y_1, \ldots, Y_m) \perp\!\!\!\perp A \mid (Z, X)$ and the fifth equality from the conditional independence property $A \perp\!\!\!\perp Z \mid X$, both of which are taken to be true by the definition of A. To show the equivalence between eqs. (3) and (4), under the target assignment regime where $P(A = z \mid X) = m^{-1}$ we can further write this as

$$\begin{aligned}
\text{SMR} &= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(X=x \mid Z=z)}{\sum_x \sum_a P(Y_a=1 \mid Z=a, X=x)P(A=a \mid X=x)P(X=x \mid Z=z)} \\
&= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(X=x \mid Z=z)}{m^{-1}\sum_x \sum_a P(Y_a=1 \mid Z=a, X=x)P(X=x \mid Z=z)} \\
&= \frac{P(Y_z=1 \mid Z=z)}{m^{-1}\sum_a P(Y_a=1 \mid Z=z)} \\
&= \frac{E[Y_z \mid Z=z]}{m^{-1}\sum_a E[Y_a|Z=z]}.
\end{aligned} \tag{14}$$

On the other hand, under the hypothetical assignment regime under which $P(Z = z \mid X) = P(A = z \mid X)$, starting from the general form of eq. (13) above, we can express the causal parameter θ_z in the form

$$\begin{aligned}
\theta_z &= \frac{\sum_x P(Y_z=1 \mid X=x, Z=z)P(X=x \mid Z=z)}{\sum_x \sum_a P(Y=1 \mid Z=a, X=x)P(A=a \mid X=x)P(X=x \mid Z=z)} \\
&= \frac{\sum_x P(Y=1 \mid X=x, Z=z)P(X=x \mid Z=z)}{\sum_x \sum_a P(Y=1 \mid Z=a, X=x)P(Z=a \mid X=x)P(X=x \mid Z=z)} \\
&= \frac{P(Y=1 \mid Z=z)}{\sum_x P(Y=1 \mid X=x)P(X=x \mid Z=z)} \\
&= \frac{E[Y \mid Z=z]}{\sum_x E[Y \mid X=x]P(X=x \mid Z=z)} \\
&= \frac{E[Y \mid Z=z]}{E\left\{E[Y \mid X] \mid Z=z\right\}},
\end{aligned} \tag{15}$$

which proves equality to eq. (5). We note that eq. (15) contains a term $E[Y \mid X]$ which could be estimated by fitting an outcome model. An alternative form may be obtained as

$$\begin{aligned}\theta_z &= \frac{\sum_x P(Y_z = 1 \mid X = x, Z = z)P(Z = z \mid X = x)P(X = x)}{\sum_x \sum_a P(Y_a = 1 \mid Z = a, X = x)P(A = a \mid X = x)P(Z = z \mid X = x)P(X = x)} \\ &= \frac{\sum_x P(Y = 1 \mid X = x, Z = z)P(Z = z \mid X = x)P(X = x)}{\sum_x P(Z = z \mid X = x) \sum_a P(Y = 1 \mid Z = a, X = x)P(A = a \mid X = x)P(X = x)} \\ &= \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{\sum_x \sum_{y=0}^{1} P(Z = z \mid X = x)yP(Y = y \mid X = x)P(X = x)} \\ &= \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E[P(Z = z \mid X)Y]}, \end{aligned} \tag{16}$$

which proves equality to eq. (6). Expression (16) only involves a term $P(Z = z \mid X)$ which could be estimated by fitting a multinomial assignment model. Finally we can derive one more expression for θ_z, beginning from eq. (13), as

$$\begin{aligned}\theta_z &= \frac{\sum_x P(Y_z = 1 \mid X = x, Z = z)P(Z = z \mid X = x)P(X = x)}{\sum_x \sum_a P(Y_a = 1 \mid Z = a, X = x)P(A = a \mid X = x)P(Z = z \mid X = x)P(X = x)} \\ &= \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{\sum_x P(Z = z \mid X = x)P(Y = 1 \mid X = x)P(X = x)} \\ &= \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{P(Z = z \mid X)P(Y = 1 \mid X)\right\}} \\ &= \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{E[Y \mid X]P(Z = z \mid X)\right\}}, \end{aligned} \tag{17}$$

which proves equality to eq. (7). Expression (17) is the final term in the proposed doubly-robust estimator. The denominator combines two terms that could be estimated by an outcome model and a multinomial assignment model, and serves as the cancellation term to achieve double robustness.

Appendix B: Consistency of the Proposed Estimator

In this appendix we show that eq. (9) is a consistent estimator. This will be done asymptotically using the Law of Large Numbers combined with Slutsky's theorem. We show here that the estimator is consistent when all models are correctly specified, as well as when each model in turn is misspecified.

B.1 A note on correctly specified models

We assume that the triples $W_i \equiv (X_i, Y_i, Z_i)$ and $W_j \equiv (X_j, Y_j, Z_j)$ are independent and identically distributed for $i \neq j$. Further, we assume that for $n \to \infty$ we have $\hat{\phi} \to \phi_0$, namely, the estimator of the outcome model parameters converges to some unknown constant in probability. We say that the relationship between an outcome variable and the covariates is correctly specified when

$$E[Y_i \mid X_i, \phi_0] = E[Y_i \mid X_i].$$

Since the parameter ϕ_0 is unknown and must be estimated by $\hat{\phi}$, by the continuous mapping theorem, $m(x_i, \hat{\phi}) \equiv E[Y_i \mid X_i, \hat{\phi}] \to E[Y_i \mid X_i, \phi_0]$. Therefore, if the model is correctly specified, we have

$$m(x_i, \hat{\phi}) \to E[Y_i \mid X_i, \phi_0] = E[Y_i \mid X_i] \text{ as } n \to \infty.$$

The above is the case when we are considering an outcome model. For the case of a correctly specified assignment model with estimated parameters $\hat{\gamma}$, we have that

$$\hat{\gamma} \to \gamma_0 \Rightarrow e(x_i, z, \hat{\gamma}) \equiv P(Z_i \mid X_i, \hat{\gamma}) \to P(Z_i \mid X_i, \gamma_0)$$

and thus if the assignment model is correctly specified, we have

$$e(x_i, z, \hat{\gamma}) \to P(Z_i \mid X_i, \gamma_0) = P(Z_i \mid X_i)$$

B.2 Consistency for correctly specified models

Our proof of general consistency will show that the numerators and denominators of each term in the summation converge in probability to an expectation that is equivalent to the numerator and denominators of the quantity of interest. Then, by Slutsky's theorem, we see that, since all three summation terms are equivalent, the summation itself converges in probability to the estimand, and therefore our estimator is consistent.

Numerators: Let $g(W_i) = \mathbf{1}_{\{Z_i=z\}}Y_i$, namely the numerator of all three terms in eq. (9). Then by Law of Large Numbers (LLN) we have

$$\frac{1}{n}\sum_{i=1}^{n} g(W_i) \xrightarrow{P} E[g(W)] = E[\mathbf{1}_{\{Z=z\}}Y]$$

where we can write

$$\begin{aligned} E[\mathbf{1}_{\{Z=z\}}Y] &= \sum_x P(Y=1 \mid Z=z, X=x)P(Z=z, X=x) \\ &= \sum_x P(Y_z=1 \mid Z=z, X=x)P(Z=z, X=x) \\ &= P(Y_z=1 \mid Z=z)P(Z=z) \\ &= P(Z=z)E[Y_z \mid Z=z] \end{aligned}$$

Thus we have that

$$\frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_{\{Z=z\}}Y_i \xrightarrow{P} P(Z=z)E[Y_z \mid Z=z].$$

Denominators: Here we will show, for each term in the estimator, that each denominator converges in probability to $P(Z=z)E[Y_A \mid Z=z]$ and thus, by Slutsky's theorem, the ratio converges in probability to the causal estimand.

1. Let $g(W_i; \hat{\gamma}) = e(x_i, z, \hat{\gamma})Y_i$, the denominator of the first term in eq. (9). Under the assumption that the assignment model is correctly specified,

$$\frac{1}{n}\sum_{i=1}^{n} g(W_i; \hat{\gamma}) \xrightarrow{P} E[g(W; \gamma_0)] = E[P(Z=z \mid X)Y]$$

where, under the causal assumptions made in Section 2, we may write

$$
\begin{aligned}
&E[P(Z = z \mid X)Y] \\
&\quad = \sum_x P(Z = z \mid X = x)P(Y = 1 \mid X = x) \\
&\quad = \sum_x \sum_a P(Y = 1 \mid Z = a, X = x)P(Z = a \mid X = x)P(X = x)P(Z = z \mid X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = z, A = a, X = x)P(A = a \mid X = x)P(X = x)P(Z = z \mid X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = z, A = a, X = x)P(A = a, Z = z \mid X = x)P(X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = z, A = a, X = x)P(A = a, X = x \mid Z = z)P(Z = z) \\
&\quad = \sum_x P(Y_A = 1 \mid Z = z, X = x)P(X = x \mid Z = z)P(Z = z) \\
&\quad = P(Z = z)E[Y_A \mid Z = z]
\end{aligned}
$$

so we have that

$$
\frac{1}{n}\sum_{i=1}^{n} e(x_i, z, \hat{\gamma})Y_i \xrightarrow{P} P(Z = z)E[Y_A \mid Z = z]
$$

and therefore, by Slutsky's theorem, we have that the first term of the estimator in eq. (9) of the main paper converges to the causal estimand,

$$
\frac{n^{-1}\sum_{i=1}^{n} \mathbf{1}_{\{Z=z\}}Y_i}{n^{-1}\sum_{i=1}^{n} e(x_i, z, \hat{\gamma})Y_i} \xrightarrow{P} \frac{P(Z = z)E[Y_z \mid Z = z]}{P(Z = z)E[Y_A \mid Z = z]} = \frac{E[Y_z \mid Z = z]}{E[Y_A \mid Z = z]}
$$

2. Let $g(W_i; \hat{\phi}) = \mathbf{1}_{\{Z_i=z\}}m(x_i, \hat{\phi})$, the denominator of the middle term of eq. (9) in the main paper. Under the assumption that the outcome model is correctly specified,

$$
\frac{1}{n}\sum_{i=1}^{n} g(W_i; \hat{\phi}) \xrightarrow{P} E[g(W; \phi_0)] = E\left\{\mathbf{1}_{\{Z=z\}}E[Y \mid X]\right\}
$$

where we may write

$$
\begin{aligned}
&E\left\{\mathbf{1}_{\{Z=z\}}E[Y \mid X]\right\} \\
&\quad = \sum_x E[Y \mid X = x]P(X = x, Z = z) \\
&\quad = \sum_x P(Y = 1 \mid X = x)P(X = x, Z = z) \\
&\quad = \sum_x \sum_a P(Y = 1 \mid Z = a, X = x)P(Z = a \mid X = x)P(X = x, Z = z) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = a, X = x)P(A = a \mid X = x)P(X = x, Z = z) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = a, X = x)P(A = a \mid X = x)P(Z = z \mid X = x)P(X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = a, X = x)P(A = a, Z = z \mid X = x)P(X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = a, X = x)P(A = a, Z = z, X = x)
\end{aligned}
$$

$$
\begin{aligned}
&= \sum_x \sum_a P(Y_a = 1 \mid Z = z, X = x) P(A = a \mid Z = z, X = x) P(X = x, Z = z) \\
&= \sum_x P(Y_A = 1 \mid Z = z, X = x) P(X = x, Z = z) \\
&= \sum_x P(Y_A = 1 \mid Z = z, X = x) P(X = x \mid Z = z) P(Z = z) \\
&= P(Z = z) E[Y_A \mid Z = z].
\end{aligned}
$$

Therefore we have that the outcome model-based term of eq. (9) converges to the causal estimand,

$$
\frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\{Z=z\}} m(x_i, \hat{\phi}) \xrightarrow{P} P(Z = z) E[Y_A \mid Z = z],
$$

so using Slutsky, we have

$$
\frac{n^{-1} \sum_{i=1}^{n} \mathbf{1}_{\{Z=z\}} Y_i}{n^{-1} \sum_{i=1}^{n} \mathbf{1}_{\{Z=z\}} m(x_i, \hat{\phi})} \xrightarrow{P} \frac{P(Z = z) E[Y_z \mid Z = z]}{P(Z = z) E[Y_A \mid Z = z]} = \frac{E[Y_z \mid Z = z]}{E[Y_A \mid Z = z]}
$$

and so the term that uses only the outcome model is a consistent estimator for our quantity of interest.

3. Let $g(W_i; \hat{\phi}, \hat{\gamma}) = m(x_i, \hat{\phi}) e(x_i, z, \hat{\gamma})$, the last term in eq. (9) of the main paper. Under the assumption that both models are correctly specified, then by the LLN

$$
\frac{1}{n} \sum_{i=1}^{n} g(W_i; \hat{\phi}, \hat{\gamma}) \xrightarrow{P} E[g(W; \phi_0, \gamma_0)] = E\left\{E[Y \mid X] P(Z = z \mid X)\right\}
$$

where we may write

$$
\begin{aligned}
&E\left\{E[Y \mid X] P(Z = z \mid X)\right\} \\
&\quad = \sum_x P(Z = z \mid X) P(Y = 1 \mid X = x) P(X = x) \\
&\quad = \sum_x \sum_a P(Z = z \mid X = x) P(Y = 1 \mid Z = a, X = x) P(Z = a \mid X = x) P(X = x) \\
&\quad = \sum_x \sum_a P(Z = z \mid X = x) P(Y_a = 1 \mid Z = a, A = a, X = x) P(A = a \mid X = x) P(X = x) \\
&\quad = \sum_x \sum_a P(Z = z \mid X = x) P(Y_a = 1 \mid Z = z, A = a, X = x) P(A = a \mid X = x) P(X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = z, A = a, X = x) P(A = a, Z = z \mid X = x) P(X = x) \\
&\quad = \sum_x \sum_a P(Y_a = 1 \mid Z = z, A = a, X = x) P(A = a, X = x \mid Z = z) P(Z = z) \\
&\quad = \sum_x P(Y_A = 1 \mid Z = z, X = x) P(X = x \mid Z = z) P(Z = z) \\
&\quad = P(Z = z) P(Y_A = 1 \mid Z = z) \\
&\quad = P(Z = z) E[Y_A \mid Z = z],
\end{aligned}
$$

so we have that

$$
\frac{1}{n} \sum_{i=1}^{n} m(x_i, \hat{\phi}) e(x_i, z, \hat{\gamma}) \xrightarrow{P} P(Z = z) E[Y_A \mid Z = z]
$$

and therefore, by Slutsky

$$\frac{n^{-1}\sum_{i=1}^{n}\mathbf{1}_{\{Z=z\}}Y_i}{n^{-1}\sum_{i=1}^{n}m(x_i,\hat{\phi})e(x_i,z,\hat{\gamma})} \xrightarrow{P} \frac{P(Z=z)E[Y_z \mid Z=z]}{P(Z=z)E[Y_A \mid Z=z]} = \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]}.$$

Finally, since each term in the summation is a consistent estimator of the causal estimand θ_z, we can use Slutsky again to show that the entire estimator is a consistent estimator for θ_z:

$$\hat{\theta}_z \xrightarrow{P} \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]} - \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]} + \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]} = \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]} = \theta_z.$$

Thus, estimator given by eq. (9) is asymptotically consistent, when the models are correctly specified.

B.3 Consistency under misspecified assignment model

Now we can check the consistency of the estimator when each of the models in turn are misspecified in order to show the double robust property. We begin by assuming that the assignment model is misspecified, but the outcome model remains correct. The misspecified assignment model, denoted by asterisk, is assumed to converge towards a constant different from the true assignment probability as $e^*(x_i,z,\hat{\gamma}) \xrightarrow{P} P^*(Z_i = z \mid X_i,\gamma_0) \neq P(Z_i \mid X_i)$. The second term in eq. (9) (estimator for eq. (5)) will consistently estimate the causal quantity of interest as the outcome model is correctly specified and the assignment model is not present in this term. Now consider the denominator of the first term of the sum, $\sum_{i=1}^{n} e^*(x_i,z,\hat{\gamma})Y_i$, when the assignment model is misspecified. We have by the law of large numbers

$$\frac{1}{n}\sum_{i=1}^{n} e^*(x_i,z,\hat{\gamma})Y_i \xrightarrow{P} E\left[P^*(Z=z \mid X,\gamma_0)Y\right].$$

For the third term in the summation, we have by LLN that, for the denominator,

$$\frac{1}{n}\sum_{i=1}^{n} m(x_i,\hat{\phi})e^*(x_i,z,\hat{\gamma}) \xrightarrow{P} E\left\{E[Y \mid X]P^*(Z=z \mid X,\gamma_0)\right\},$$

where we can write

$$\begin{aligned} E\left\{P^*(Z=z \mid X,\gamma_0)E[Y \mid X]\right\} &= \sum_x P^*(Z=z \mid X=x,\gamma_0)P(Y=1 \mid X=x)P(X=x) \\ &= \sum_x \sum_{y=0}^{1} P^*(Z=z \mid X=x,\gamma_0)yP(Y=y \mid X=x)P(X=x) \\ &= E\left[P^*(Z=z \mid X,\gamma_0)Y\right], \end{aligned}$$

which is equivalent to the asymptotic denominator of the third term above. Thus, using information from the previous section, we have

$$\hat{\theta}_z \xrightarrow{P} \frac{P(Z=z)E[Y_z \mid Z=z]}{E[P^*(Z=z \mid X,\gamma_0)Y]} - \frac{P(Z=z)E[Y_z \mid Z=z]}{E[P^*(Z=z \mid X,\gamma_0)Y]} + \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]} = \theta_z.$$

Therefore, when the assignment model is misspecified but the outcome model is correct, we have that the doubly robust estimator remains a consistent estimator.

B.4 Consistency under misspecified outcome model

As we have shown earlier, the numerators converge as follows:

$$\frac{1}{n}\sum_{i=1}^{n}\mathbf{1}_{\{Z=z\}}Y_i \xrightarrow{P} E[\mathbf{1}_{\{Z=z\}}Y],$$

where it is possible to write the asymptotic numerator alternatively as

$$\begin{aligned}
E[\mathbf{1}_{\{Z=z\}}Y] &= \sum_x P(Y=1 \mid Z=z, X=x)P(Z=z \mid X=x)P(X=x) \\
&= \sum_x P(Y=1 \mid Z=z, X=x)P(X=x \mid Z=z)P(Z=z) \\
&= P(Z=z)P(Y=1 \mid Z=z) \\
&= P(Z=z)E[Y \mid Z=z].
\end{aligned}$$

Further, the misspecified outcome model, denoted by asterisk, is assumed to converge to a constant different from the true expected outcome, as $m^*(x_i, \hat{\phi}) \xrightarrow{P} E^*[Y_i \mid X_i, \phi_0] \neq E[Y_i \mid X_i]$. We also have by the LLN that

$$\frac{1}{n}\sum_{i=1}^{n}\mathbf{1}_{\{Z_i=z\}}m^*(x_i, \hat{\phi}) \xrightarrow{P} E\left\{\mathbf{1}_{\{Z=z\}}E^*[Y \mid X, \phi_0]\right\}$$

and

$$\frac{1}{n}\sum_{i=1}^{n}m^*(x_i, \hat{\phi})e(x_i, z, \hat{\gamma}) \xrightarrow{P} E\left\{E^*[Y \mid X, \phi_0]P(Z=z \mid X)\right\}.$$

The first of these may be expressed as

$$\begin{aligned}
E\left\{\mathbf{1}_{\{Z=z\}}m^*(x_i, \phi_0)\right\} &= \sum_x P(Z=z \mid X=x)P(X=x)E^*[Y \mid X=x, \phi_0] \\
&= \sum_x P(X=x \mid Z=z)P(Z=z)E^*[Y \mid X=x, \phi_0] \\
&= P(Z=z)E\left\{E^*[Y \mid X, \phi_0] \mid Z=z\right\}.
\end{aligned}$$

We note that under the misspecified outcome model, for the middle term of eq. (9) we have the convergence

$$\frac{\sum_{i=1}^{n}\mathbf{1}_{\{Z=z\}}Y_i}{\sum_{i=1}^{n}\mathbf{1}_{\{Z=z\}}m^*(x_i, \hat{\phi})} \xrightarrow{P} \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{\mathbf{1}_{\{Z=z\}}E^*[Y \mid X, \phi_0]\right\}}, \tag{18}$$

and for the third term of eq. (9) the convergence

$$\frac{\sum_{i=1}^{n}\mathbf{1}_{\{Z=z\}}Y_i}{\sum_{i=1}^{n}m^*(x_i, \hat{\phi})e(x_i, z, \hat{\gamma})} \xrightarrow{P} \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{E^*[Y \mid X, \phi_0]P(Z=z \mid X)\right\}}, \tag{19}$$

Here for the right hand side of eq. (18) we get

$$\begin{aligned}
\frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{\mathbf{1}_{\{Z=z\}}E^*[Y \mid X, \phi_0]\right\}} &= \frac{E[Y \mid Z=z]}{E\left\{E^*[Y \mid X, \phi_0] \mid Z=z\right\}} \\
&= \frac{P(Y=1 \mid Z=1)}{\sum_x E^*[Y \mid X=x, \phi_0]P(X=x \mid Z=z)} \\
&= \frac{\sum_x P(Y=1 \mid X=x, Z=z)P(X=x \mid Z=z)}{\sum_x E^*[Y \mid X=x, \phi_0]P(X=x \mid Z=z)} \\
&= \frac{\sum_x P(Y=1 \mid Z=z, X=x)P(Z=z \mid X=x)P(X=x)}{\sum_x E^*[Y \mid X=x, \phi_0]P(Z=z \mid X=x)P(X=x)} \\
&= \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{E^*[Y \mid X, \phi_0]P(Z=z \mid X)\right\}}.
\end{aligned}$$

Therefore, we have that

$$\hat{\theta}_z \xrightarrow{P} \frac{E[Y_z \mid Z=z]}{E[Y_A \mid Z=z]} - \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{E^*[Y \mid X, \phi_0]P(Z=z \mid X)\right\}} + \frac{E[\mathbf{1}_{\{Z=z\}}Y]}{E\left\{E^*[Y \mid X, \phi_0]P(Z=z \mid X)\right\}} = \theta_z,$$

and thus, when the assignment model is correctly specified, the DR estimator remains an asymptotically consistent estimator, and thus we have the doubly robust property that we require.

References

Bang, H., and Robins, J. M. (2005). Doubly robust estimation in missing data and causal inference models. Biometrics, 61:962–972.

Faris, P. D., Ghali, W. A., and Brant, R. (2003). Bias in estimates of confidence intervals for health outcome report cards. Journal of Clinical Epidemiology, 56:553–558.

Funk, M. J., Westreich, D., Wiesen, C., Sturmer, T., Brookhart, A. M., and Davidian, M. (2011). Doubly robust estimation of causal effects. American Journal of Epidemiology, 173:761–767.

Hernan, M. A., and Robins, J. M. (2006). Estimating causal effects from epidemiological data. Journal of Epidemiology and Community Health, 60:578–586.

Kang, J. D., and Schafer, J. L. (2007). Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. Statistical Sciences, 22:523–539.

Keiding, N., and Clayton, D. (2014). Standardization and control for confounding in observational studies: A historical perspective. Statistical Sciences, 29:529–558.

Massarweh, N. N., Hu, C. Y., You, Y. N., Bednarski, B. K., Rodriguez-Bigas, M. A., Skibber, J. M., Cantor, S. B., Cormier, J. N., Feig, B. W., and Chang, G. J. (2014). Risk-adjusted pathologic margin positivity rate as a quality indicator in rectal cancer surgery. Journal of Clinical Oncology, 32:2967–2974.

Raval, M. V., Bilimoria, K. Y., Stewart, A. K., Bentrem, D. J., and Ko, C. Y. (2009). Using the NCDB for cancer care improvement: an introduction to available quality assessment tools. Journal of Surgical Oncology, 99:488–490.

Robins, J., Sued, M., Lei-Gomez, Q., and Rotnitzky, A. (2007). Comment: Performance of doubly-robust estimators when "inverse probability" weights are highly variable. Statistical Sciences, 22:544–559.

Shahian, D. M., and Normand, S.-L. T. (2008). Comparison of “risk-adjusted” hospital outcomes. Circulation: Journal of the American Heart Association, 117:1955–1963.

Shinozaki, T., and Matsuyama, Y. (2015). Doubly robust estimation of standardized risk difference and ratio in the exposed population. Epidemiology, 26:873–877.

Spiegelhalter, D. (2005). Funnel plots for comparing institutional performance. Statistics in Medicine, 24:1185–1202.

Tang, X., Gan, F. F., and Zhang, L. (2015). Standardized mortality ratio for an estimated number of deaths. Journal of Applied Statistics, 42:1348–1366.

Varewyck, M., Goetghebeur, E., Eriksson, M., and Vansteelandt, S. (2014). On shrinkage and model extrapolation in the evaluation of clinical center performance. Biostatistics, 15:651–664.

Varewyck, M., Vansteelandt, S., Eriksson, M., and Goetghebeur, E. (2016). On the practice of ignoring center-patient interactions in evaluating hospital performance. Statistics in Medicine, 35:227–238.