THE CENTRE FOR MARKET AND PUBLIC ORGANISATION

Identification of Causal Effects on Binary Outcomes Using Structural Mean Models*

Paul Clarke and Frank Windmeijer

January 2010

Working Paper No. 09/217
*Please Note: this is a substantially revised version of the original June 2009 paper

Published as Cemmap Working Paper Series number 02/10 (2010)
Published in Biostatistics Advance Access (June 2010)

Centre for Market and Public Organisation
Bristol Institute of Public Affairs
University of Bristol
2 Priory Road
Bristol BS8 1TX
http://www.bristol.ac.uk/cmpo/

Tel: (0117) 33 10799
Fax: (0117) 33 10705
E-mail: cmpo-office@bristol.ac.uk

The Centre for Market and Public Organisation (CMPO) is a leading research centre, combining expertise in economics, geography and law. Our objective is to study the intersection between the public and private sectors of the economy, and in particular to understand the right way to organise and deliver public services. The Centre aims to develop research, contribute to the public debate and inform policy-making.

CMPO, now an ESRC Research Centre was established in 1998 with two large grants from The Leverhulme Trust. In 2004 we were awarded ESRC Research Centre status, and CMPO now combines core funding from both the ESRC and the Trust.

ISSN 1473-625X

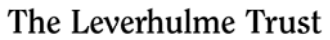

CMPO Working Paper Series No. 09/217

Identification of Causal Effects on Binary Outcomes Using Structural Mean Models

Paul Clarke[1]
and
Frank Windmeijer[2]

[1]*CMPO, University of Bristol, UK*
[2]*Department of Economics and CMPO, University of Bristol, UK*

June 2009

Abstract
Structural mean models (SMMs) are used to estimate causal effects among those selecting treatment in randomised controlled trials affected by non-ignorable non-compliance. These causal effects can be identified by assuming that there is no effect modification, namely, that the causal effect is equal for the treated subgroups randomised to treatment and to control. By analysing simple structural models for binary outcomes, we argue that the no effect modification assumption does not hold in general, and so SMMs do not estimate causal effects for the treated. An exception is for designs in which those randomised to control can be completely excluded from receiving the treatment. However, when there is non-compliance in the control arm, local (or complier) causal effects can be identified provided that the further assumption of monotonic selection into treatment holds. We demonstrate these issues using numerical examples.

Keywords: structural mean models, identification, local average treatment effects, complier average treatment effects.

JEL Classification: C13, C14

Electronic version: www.bristol.ac.uk/cmpo/publications/papers/2008/wp217.pdf

Acknowledgements
This work was funded by UK Economic & Social Research Council grant RES-060-23-0011 and UK Medical Research Council grant G0601625. The authors would like to thank Vanessa Didelez, Roger Harbord, Tom Palmer, Nuala Sheehan and Graham Dunn for their comments on an earlier draft.

Address for correspondence
CMPO, Bristol Institute for Public Affairs
University of Bristol
2 Priory Road
Bristol
BS8 1TX
f.windmeijer@bristol.ac.uk
www.bristol.ac.uk/cmpo/

1 Introduction

Robins (1989, 1994) introduced the class of semi-parametric structural mean models (SMMs) and their associated 'G-estimators' for the estimation of causal effects of treatment regimes on outcomes from randomised controlled trials affected by non-compliance. Non-compliance comes about when participants choose treatments other than those to which they were randomised. Of most interest are SMM estimators that allow for the effects of non-ignorable non-compliance, that is, where participants choose their treatments in a manner associated with their study outcomes, even after baseline (and possibly time-varying) covariates have been adjusted for. SMMs for non-ignorable non-compliance are widely used in biomedical research: see, for example, Goetghebeur and Lapp (1997), Witteman et al. (1998), Fischer-Lapp and Goetghebeur (1999), Ten Have et al. (2004), Tanaka et al. (2008), and Moodie et al. (2009).

The parameters of SMMs correspond to meaningful functions of expected potential outcomes for the population of participants exposed to the treatment. For example, additive SMMs are specified in terms of average treatment (or causal) effects, and multiplicative SMMs in terms of causal risk ratios. Vansteelandt and Goetghebeur (2003) developed the generalised SMM and we consider its important special case, the logistic SMM and its associated 'double-logistic' estimator for causal odds ratios. Hernán and Robins (2006) review additive and multiplicative SMMs and consider the relationship between these and econometric instrumental variable estimators; Goetghebeur and Vansteelandt (2005) review all of the SMMs considered here.

In this paper, we consider the estimation of causal effects using SMMs from studies in which the outcome is binary. Superficially, SMMs are applicable no matter what the outcome's measurement scale, but we will show that the binary case poses problems for SMM estimators. The usual identification assumption is 'no effect modification by randomisation', but we argue that it does not generally hold for binary outcomes. In

fact, the usual target parameters are identified only with assumptions like the 'treatment exclusion restriction' wherein no-one randomised to the control group can receive the treatment. For more general designs, provided that patients' treatment selection is monotonic (e.g., Angrist et al., 1996), additive and multiplicative SMMs identify 'local' (or 'complier') causal effects, but the double-logistic SMM does not. If researchers fail to recognise the difficulties that arise with binary outcomes then misleading interpretations of SMM estimates could result. We present some numerical results for simple examples to demonstrate this.

2 Structural Mean Models

2.1 Potential Outcomes

Before introducing the models, we first set out the potential outcomes notation to be used throughout. To simplify notation and highlight concepts, we consider only the simplest set-up: a randomised controlled trial in which patients are randomised to a fixed treatment dose or to the control group, which they comply with or not according to some non-ignorable mechanism; the binary study outcome is measured after some fixed follow-up period. The focus on this simple set-up is done without loss of generality and our findings apply equally to situations including pre-randomisation covariates, variable treatment dose, and treatment regimes involving repeated doses with time-varying covariates recorded.

Following Hernán and Robins (2006), let Y, X and Z denote random variables representing the following observed quantities: Z is the randomisation assignment indicator, with $Z = 1$ denoting treatment and $Z = 0$ control; $X \in \{1, 0\}$ is the corresponding indicator for the actual treatment chosen by the patient, where $X \neq Z$ is possible due to non-compliance; and $Y \in \{0, 1\}$ is the binary study outcome. It is assumed throughout

that the observed data $\{(y_i, x_i, z_i) : i = 1, ..., n\}$ constitute an *i.i.d.* sample from the target population.

The potential outcomes can now be defined in the usual way. Let $Y(x, z)$ be the potential outcome that would be obtained if the treatment assignment was set to z and the treatment received to x by external intervention, rather than by randomising and letting the patient choose himself/herself. Similarly, let $X(z)$ be the potential treatment that would be obtained if treatment assignment was set to z by external intervention.

Four necessary (but not sufficient) conditions for identification can now be stated as follows: the 'stable unit treatment value' assumption that each patient's potential outcomes are mutually independent of those of any other patient; the existence of 'causal effects' of Z on X and on Y; the 'consistency assumption' $Y = Y(X, Z)$ and $X = X(Z)$, linking the observed and potential outcomes; and the 'exclusion restriction' $Y(x, z) = Y(x)$ constraining the effect of treatment assignment to affect the study outcome *only* through its effect on treatment choice (e.g., Angrist et al., 1996). All of these are taken to hold throughout this paper.

2.2 The Additive and Multiplicative SMMs

For the simple scenario just described, the additive SMM is

$$E\left(Y|X, Z\right) - E\left\{Y\left(0\right)|X, Z\right\} = \left(\psi_0 + \psi_1 Z\right) X,$$

where $Y\left(0\right)$ is the treatment-free potential outcome. While this model is saturated, or non-parametric, in more general scenarios the right hand side will be a parametric function incorporating the effect of pre-randomisation covariates or variable treatment dose, which is why SMMs are usually referred to as semi-parametric models. The parameters of the additive model are $\psi_0 = E\left\{Y(1) - Y(0)|X = 1, Z = 0\right\}$ and $\psi_0 + \psi_1 = E\left\{Y(1) - Y(0)|X = 1, Z = 1\right\}$, that is, the average causal effect among those who chose treatment but were assigned the control, and the average causal effect among

those who were assigned to and chose treatment, respectively.

SMM estimators are based on the conditional mean independence, or randomisation assumption

$$E\left\{Y\left(0\right)|Z=1\right\}=E\left\{Y\left(0\right)|Z=0\right\}, \tag{1}$$

which holds provided that randomisation is unrelated to the untreated potential outcome. Under the additive SMM, (1) can be rewritten as

$$E\left\{Y-\left(\psi_0+\psi_1\right)X|Z=1\right\}=E\left(Y-\psi_0X|Z=0\right), \tag{2}$$

from which an estimating equation can be constructed.

The saturated multiplicative SMM for the same scenario is defined as

$$\frac{E\left(Y|X,Z\right)}{E\left\{Y\left(0\right)|X,Z\right\}}=\exp\left\{\left(\theta_0+\theta_1Z\right)X\right\}.$$

The parameters of the multiplicative SMM are

$$\exp\left(\theta_0\right)=\frac{E\left\{Y(1)|X=1,Z=0\right\}}{E\left\{Y(0)|X=1,Z=0\right\}}$$

and

$$\exp\left(\theta_0+\theta_1\right)=\frac{E\left\{Y(1)|X=1,Z=1\right\}}{E\left\{Y(0)|X=1,Z=1\right\}},$$

that is, causal risk ratios among the same two subgroups as before. Under the multiplicative SMM, the conditional mean independence assumption (1) leads to the moment condition

$$E\left[Y\exp\left\{-\left(\theta_0+\theta_1Z\right)X\right\}|Z=1\right]=E\left\{Y\exp\left(-X\theta_0\right)|Z=0\right\}. \tag{3}$$

It is clear that neither set of SMM parameters is identified by its corresponding moment condition because both constitute systems with two unknowns in one equation. Therefore, further assumptions are required to identify the SMM parameters. Hernán and Robins (2006) highlight the role of the 'no effect modification by Z' assumption: for

the additive SMM, it corresponds to constraining $\psi_1 = 0$, and for the multiplicative SMM it corresponds to $\theta_1 = 0$. Under no effect modification, there is only one unknown in both (2) and (3) and identification is achieved. The target parameter for the additive SMM is then

$$\psi_0 = E\left\{Y(1) - Y(0)|X = 1\right\},$$

that is, the average causal effect among the treated, and for the multiplicative SMM

$$\exp(\theta_0) = \frac{E\left\{Y(1)|X = 1\right\}}{E\left\{Y(0)|X = 1\right\}},$$

i.e. the risk ratio among the treated.

The estimators of the additive and multiplicative SMM target parameters can be written as

$$\widehat{\psi}_0 = \frac{E(Y|Z = 1) - E(Y|Z = 0)}{E(X|Z = 1) - E(X|Z = 0)}, \tag{4}$$

and

$$\exp(\widehat{\theta}_0) = 1 - \frac{E(Y|Z = 1) - E(Y|Z = 0)}{E\left\{(1 - X)Y|Z = 1\right\} - E\left\{(1 - X)Y|Z = 0\right\}}, \tag{5}$$

respectively (e.g., Angrist, 2001; Hernán and Robins, 2006). The additive SMM estimator has the same form as the classical instrumental variable estimator (Angrist et al., 1996); the numerator in both expressions is called the 'intention to treat' estimator. More generally, the estimating equations under additive and multiplicative SMMs based on (1) can be solved by G-estimation (Robins, 1994).

The crucial nature of the no effect modification assumption in identifying both SMM estimators is thus apparent. It is the validity of this assumption in the binary case that we will consider in detail below in Section 3.

2.3 The Double-Logistic SMM

Robins et al. (1999) proposed the logistic SMM

$$\left\{\frac{E\left(Y|X, Z\right)}{E\left(1 - Y|X, Z\right)}\right\} / \left[\frac{E\left\{Y(0)|X, Z\right\}}{E\left\{1 - Y(0)|X, Z\right\}}\right] = \exp\left\{\left(\xi_0 + \xi_1 Z\right) X\right\},$$

parameterised in terms of the causal odds ratios among the two treated groups. Under the no effect modification assumption $\xi_1 = 0$, its target parameter is

$$\exp(\xi_0) = \frac{E\{Y(1)|X=1,Z=z\}/E\{1-Y(1)|X=1,Z=z\}}{E\{Y(0)|X=1,Z=z\}/E\{1-Y(0)|X=1,Z=z\}}.$$

The logistic SMM is considered separately here because Robins (1999) showed that no consistent G-estimator for ξ_0 can be constructed without further assumptions. Vansteelandt and Goetghebeur (2003) developed the double-logistic estimator by exploiting the result that ξ_0 can potentially be identified if the researcher specifies a parametric 'association model'

$$E(Y|X,Z) = m_\beta(X,Z),$$

which is indexed by parameter vector β. The double-logistic estimator is based on specifying $m_\beta(X,Z)$ to be logistic. A drawback to this approach, as acknowledged by Vansteelandt and Goetghebeur (2003), is that both the SMM and the association model cannot generally be logistic (i.e., one but not the other can be logistic), so the double-logistic SMM is generally 'uncongenial' in the sense described by Meng (1994). Robins and Rotnitzky (2004) propose an alternative estimator that avoids this problem, but it is still affected by the identification problem that concerns us here. Hence, we shall continue to focus on the double-logistic estimator because of its relative simplicity, and so proceed by assuming the effect of uncongeniality is negligible.

3 No Effect Modification and Binary Outcomes

The role of the no effect modification (NEM) assumption is to constrain the causal effects among the treated to be equal for those randomised to treatment and those randomised to control. For example, recall that the NEM assumption for the additive SMM constrains $\psi_1 = 0$ and thus

$$E\{Y(1)-Y(0)|X=1,Z=0\} = E\{Y(1)-Y(0)|X=1,Z=1\}.$$

To investigate the validity of this assumption for binary outcomes, it is necessary to consider the hypothetical structural model that generated the observed data. Pearl (2000) gives a full account of the link between structural models and potential outcomes, but we follow the overview given by Hernán and Robins (2006, Appendix 1).

Informally, we posit that there must exist a non-parametric structural equation model generating the observed data that satisfies the assumptions set out in Section 2.1. Note that by 'non-parametric' we are not implying that the true data generating process cannot be parametric, merely that no constraints are placed on its unknown form. Under these assumptions, the 'one-step ahead counterfactual' can be written

$$Y(x) = I\left\{f_Y^*(x, U) > 0\right\},$$

where $I(w)$ is the indicator function taking the value 1 if w is true and 0 otherwise, $f_Y^*(x, U)$ is a function that depends on the fixed value of treatment and on a latent random variable, or vector, U representing all unobserved factors influencing the outcome. It is usual to interpret U as representing the impact of unobserved 'confounding' variables. Similarly, the potential outcome for chosen treatment is $X(z) = I\left\{f_X^*(z, V)\right\}$, where V is another latent random variable, or vector, representing unobservable factors influencing treatment choice. As the outcome is binary the indicator function is unavoidable, but within it $f_Y^*(x, u)$ can be any suitable function that generates, say, non-linear or heterogeneous treatment effects. For fixed x, it is U that determines whether the potential outcome is zero or one for a particular patient, and similarly for $X(z)$.

To understand the implications of this, note that

$$E\left\{Y(x)|X = 1, Z = z\right\} = \Pr\left\{f_Y^*\left(x, U\right) > 0|f_X^*\left(z, V\right) > 0\right\}.$$

It is clear that, in general, the NEM assumption cannot hold under such models. For a specific example, consider the simple parametric structural model with (U, V) a bivariate

continuous random vector related to the potential outcomes by

$$Y(x) = I(\alpha + \beta x - U > 0), X(z) = I(\gamma + \delta z - V > 0), \tag{6}$$

where $E(U) = E(V) = 0$ and (U, V) has distribution function $F(u, v; \rho)$, with 'correlation' parameter ρ indexing all non-zero moments involving products of U and V. This model specification is of course very simple because it does not allow for heterogeneity in the effect of treatment on the latent scale. Even in this case, however, it can be shown that

$$\begin{aligned}
E\{Y(1)|X=1,Z=1\} &= \Pr(U < \alpha+\beta | V < \gamma+\delta) = F(\alpha+\beta, \gamma+\delta; \rho)/F_V(\gamma+\delta) \\
E\{Y(1)|X=1,Z=0\} &= \Pr(U < \alpha+\beta | V < \gamma) = F(\alpha+\beta, \gamma; \rho)/F_V(\gamma) \\
E\{Y(0)|X=1,Z=1\} &= \Pr(U < \alpha | V < \gamma+\delta) = F(\alpha, \gamma+\delta; \rho)/F_V(\gamma+\delta) \\
E\{Y(0)|X=1,Z=0\} &= \Pr(U < \alpha | V < \gamma) = F(\alpha, \gamma; \rho)/F_V(\gamma),
\end{aligned}$$

where $F_V(v)$ is the marginal distribution function of V, and so for the additive SMM

$$\psi_0 + \psi_1 = \frac{F(\alpha+\beta, \gamma+\delta; \rho)}{F_V(\gamma+\delta)} - \frac{F(\alpha, \gamma+\delta; \rho)}{F_V(\gamma+\delta)} \neq \frac{F(\alpha+\beta, \gamma; \rho)}{F_V(\gamma)} - \frac{F(\alpha, \gamma; \rho)}{F_V(\gamma)} = \psi_0$$

almost everywhere. Hence, we can see that the NEM assumption does not hold for the additive SMM model, and so its target parameter, the average treatment effect among the treated, is not identified. In other words, the additive SMM is not estimating the average treatment effect among the treated. Writing out the appropriate risk and odds ratios for the multiplicative SMM and logistic SMM models in a similar fashion leads us to the same conclusion regarding NEM for both.

It now remains to establish what can be identified using additive, multiplicative and double-logistic SMM estimators for binary outcomes.

4 Identification By Treatment Exclusion

In the previous section, we established that the target parameters of all three SMMs are non-identified, even in our simple scenario, because the NEM assumption does not hold

for binary outcomes. For certain designs, however, it is possible to identify the SMM parameters without requiring that NEM holds.

Vansteelandt and Goetghebeur (2003) originally proposed the double-logistic estimator in the context of 'placebo-control' randomised controlled trials. For this design, neither compliers nor non-compliers randomised to control can receive the treatment because non-compliers receive only the placebo, equating to the condition $\Pr(X = 0|Z = 0) = 1$. More generally, as we will now outline, this treatment exclusion restriction is crucial and a special case of the identifying assumptions for binary outcome SMMs described by Robins and Rotnitzky (2004). In placebo-control designs, an additional assumption of no placebo effect is also needed that we herein take to hold.

Under treatment exclusion, the SMM parameters ψ_0, θ_0 and ξ_0 are not defined because all three are conditioned on the event $\{X = 1, Z = 0\}$, which has measure zero. Conversely, $\{X = 1\} = \{X = 1, Z = 1\}$ and so $\psi_0 + \psi_1 = \psi = E\{Y(1) - Y(0)|X = 1\}$, $\exp(\theta_0 + \theta_1) = \exp(\theta) = E\{Y(1)|X = 1\}/E\{Y(0)|X = 1\}$, and

$$\exp(\xi_0 + \xi_1) = \exp(\xi) = \frac{E\{Y(1)|X = 1\}}{E\{1 - Y(1)|X = 1\}} \Big/ \frac{E\{Y(0)|X = 1\}}{E\{1 - Y(0)|X = 1\}},$$

for the additive, multiplicative and logistic SMMs, respectively.

Under treatment exclusion, $E\{Y(0)|Z = 0\}$ is always non-parametrically identified because

$$E\left(Y|Z = 0\right) = E\{Y(0)|Z = 0\}.$$

Now consider $E\{Y(0)|Z = 1\}$ and expand it to give

$$\begin{aligned} E\left\{Y\left(0\right)|Z = 1\right\} &= E\left\{Y\left(0\right)|X = 1, Z = 1\right\} E\left(X|Z = 1\right) \\ &\quad + E\left\{Y\left(0\right)|X = 0, Z = 1\right\} E(1 - X|Z = 1). \end{aligned}$$

Using the conditional mean independence assumption (1), it then follows that we can estimate the counterfactual values directly from the data:

$$E\left\{Y\left(0\right)|X = 1, Z = 1\right\} = \frac{E\left(Y|Z = 0\right) - E\left\{\left(1 - X\right)Y|Z = 1\right\}}{E\left(X|Z = 1\right)}.$$

Hence, the estimators of the additive and multiplicative SMM parameters under treatment exclusion are, respectively,

$$\hat{\psi} = \frac{E(Y|Z=1) - E(Y|Z=0)}{E(X|Z=1)},$$

and

$$\exp(\hat{\theta}) = \frac{E(XY|Z=1)}{E(Y|Z=0) - E\{(1-X)Y|Z=1\}}. \tag{7}$$

The double-logistic SMM estimator is

$$\exp(\hat{\xi}) = \frac{E(Y|X=1,Z=1)}{E(1-Y|X=1,Z=1)} \Big/ \frac{E(Y|Z=0) - E\{(1-X)Y|Z=1\}}{E(X|Z=1) - E(Y|Z=0) + E\{(1-X)Y|Z=1\}},$$

which corresponds to the estimator proposed by Vansteelandt and Goetghebeur (2003) for our simple scenario.

5 Identification Under Monotonic Selection

The treatment exclusion restriction may be unrealistic for randomised controlled trials without a placebo-control. In such circumstances, one is forced to focus on alternative causal parameters because the target SMM parameters are non-identified.

Imbens and Angrist (1994) and Angrist et al. (1996) highlight the importance of monotonicity in problems affected by non-ignorable non-compliance. Patient treatment selection is monotonic if

$$X(1) \geq X(0) \tag{8}$$

for all patients for some coding of X, Z.

In this set-up, monotonic selection corresponds to the assumption that no patient will be a defier, such that $X(0) = 1, X(1) = 0$, with probability one. For this definition to make sense, we must assume that for each patient there are two universes in which he/she is randomised to control and randomised to treatment, and so the 'no defiers' assumption corresponds to saying that while patients can disobey their treatment assignments in one

or other of these universes, patients cannot disobey their assignments in both. As an example, the simple structural model described in Section 3 is monotonic because

$$X(1) = I\left(\gamma + \delta - V > 0\right) \geq I\left(\gamma - V > 0\right) = X(0),$$

if $\delta > 0$.

While we have shown that the NEM assumption does not generally hold, additive and multiplicative SMM estimators (4) and (5) do identify local (or complier) effects under monotonic selection. Compliers are those people who comply with their treatment assignments in both hypothetical universes, such that they satisfy $X(0) = 0, X(1) = 1$, which we write as $X(1) > X(0)$. Specifically, consider estimator (4) based on the additive SMM. As noted previously, it has the same form as the classical instrumental variable estimator and so from the results of Imbens and Angrist (1994) it follows that it is consistent for the local average treatment effect (LATE),

$$\text{LATE} = E\left\{Y\left(1\right) - Y\left(0\right) | X\left(1\right) > X\left(0\right)\right\}, \tag{9}$$

which is also called the 'complier average causal effect' (CACE). Note that treatment exclusion can be seen as an extreme special case of monotonic selection in which $X(1) \geq X(0) = 0$ and the complier and treated groups are equivalent.

Similarly, Angrist (2001) showed that estimator (5) based on the multiplicative SMM under NEM is consistent for the local relative risk (LRR),

$$\text{LRR} = \frac{E\left\{Y\left(1\right) | X\left(1\right) > X\left(0\right)\right\}}{E\left\{Y\left(0\right) | X\left(1\right) > X\left(0\right)\right\}}. \tag{10}$$

See also Hernán and Robins (2006) and an alternative derivation without using SMMs by Greenland (2000).

The equivalent double-logistic SMM estimator is the solution to the moment condition

$$E\left[\text{expit}\left\{\beta_{01} + (\beta_{11} - \xi_0)X\right\} | Z = 1\right] = E\left[\text{expit}\left\{\beta_{00} + (\beta_{10} - \xi_0)X\right\} | Z = 0\right], \tag{11}$$

where $\text{expit}(a) = \exp(a) / \{1 + \exp(a)\}$. Estimates for $(\beta_{00}, \beta_{10}, \beta_{01}, \beta_{11})$ are obtained at the first stage by fitting logistic association model $m_\beta(X, Z) = \text{expit}(\beta_{00} + Z\beta_{01} + X\beta_{10} + XZ\beta_{11})$. Moment condition (11) can be solved iteratively. However, as shown by the numerical illustrations below, even assuming monotonicity it is not generally consistent for the local odds ratio (LOR), defined as

$$\text{LOR} = \frac{E\{Y(1)|X(1) > X(0)\}}{E\{1 - Y(1)|X(1) > X(0)\}} \Big/ \frac{E\{Y(0)|X(1) > X(0)\}}{E\{1 - Y(0)|X(1) > X(0)\}}.$$

Clarke and Windmeijer (2009, Appendix 3) show that the exception to this rule is if $E\{Y(1)|X(1) = X(0) = 1\} = E\{Y(1)|X(1) > X(0)\}$, in which case $\exp(\widehat{\xi}_0)$ is consistent for the LOR. However, the LOR can be estimated without this assumption. Abadie (2003) proposes an estimator for the LOR, but van der Laan et al. (2007) note how the same estimator can be derived based on the relative risk estimator (5): first calculate $\exp(\widehat{\theta}_0)$ as per usual, then recode the outcome variable as $Y^* = 1 - Y$ and calculate $\exp\left(\widehat{\theta}_0^*\right)$ replacing Y by Y^* in (5), then the ratio $\exp(\widehat{\theta}_0)/\exp\left(\widehat{\theta}_0^*\right)$ is consistent for the LOR by symmetry of the relative risk.

6 Numerical Examples

To recap, the starting point of our analysis is that SMM parameters are identified only by making assumptions additional to those set out in Section 2.1. The NEM assumption is often used for identification, but in Section 3 we showed that it does not hold for simple binary structural models. An alternative to NEM is the treatment exclusion restriction, appropriate for randomised controlled trials such as those with placebo-control designs, but implausible for other scenarios. In Section 5, however, we showed that the additive and multiplicative SMM estimators are consistent for local, or complier, effects under monotonic selection, but that the double-logistic SMM is not consistent without further assumptions.

In this section, we conduct two numerical studies to investigate the implications of these findings for scenarios where the treatment exclusion restriction does not hold. We first look at a scenario in which the true selection mechanism is monotonic and the unobserved confounders normally distributed. In this setting we can analytically calculate and compare the key causal parameters. In the second scenario, we replicate the design of Didelez et al. (2008) in which treatment and outcome data are generated using a logistic model and calculate the key causal parameters and SMM estimates in a Monte Carlo study. In the latter example, we also consider the case where the true selection mechanism is not monotonic. Our aim in both studies is to show the impact of misinterpreting the estimand of a SMM estimator as a causal effect among the treated.

6.1 Example 1

The first illustration is based on the structural model (6) defined in Section 3:

$$Y(x) = I\left(\alpha + \beta x - U > 0\right), X(z) = I\left(\gamma + \delta z - V > 0\right).$$

We set (U, V) to have the bivariate normal distribution

$$\begin{pmatrix} U \\ V \end{pmatrix} \sim N\left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right\},$$

and set $\Pr(Z = 1) = 0.5$. Note that ρ indexes the strength of non-ignorability in the selection mechanism determining compliance, with $\rho = 0$ corresponding to ignorable compliance. For each set of parameter values $(\alpha, \beta, \gamma, \delta, \rho)$, we can calculate the corresponding values of the key causal parameters. We fix the parameters in the outcome model to $\alpha = 0$, $\beta = 0.1$ and look at how the causal parameters vary as a function of (γ, δ, ρ).

Figure 1 displays the values of average treatment effects, relative risks and odds ratios as a function of ρ for $\gamma = 0$ and $\delta = 0.5$. In the first panel, ATE denotes the average treatment effect $E\{Y(1) - Y(0)\}$, and the parameters of the additive SMM are denoted as

follows: $E\{Y(1)-Y(0)\,|X=1,Z=1\}$ by ATEX1Z1, $E\{Y(1)-Y(0)\,|X=1,Z=0\}$ by ATEX1Z0, and the average treatment effect among the treated $E\{Y(1)-Y(0)\,|X=1\}$ by ATEX1. LATE denotes $E\{Y(1)-Y(0)\,|X(1)>X(0)\}$. The parameters are similarly defined in the second and third panel for the relative risk and odds ratio respectively (RRX1Z1, ORX1Z0, etc.). For the odds ratio, there is an additional parameter denoted by 'VG' corresponding to the estimand of the double-logistic SMM (11).

For $\alpha = 0$ and $\beta = 0.1$, the marginal expectations are $E\{Y(1)\} = \Phi(0.1) = 0.5398$ and $E\{Y(0)\} = \Phi(0) = 0.5$, and hence ATE $= 0.5398 - 0.5 = 0.0398$, RR $= 1.0796$ and OR $= 1.1730$. Likewise, as $\gamma = 0$ and $\delta = 0.5$ then $E\{X(1)\} = \Phi(0.5) = 0.6915$ and $E\{X(0)\} = \Phi(0) = 0.5$, indicating a large degree of non-compliance in the control arm, as $E\{X(0)\} = \Pr(X=1|Z=0)$. The proportion of compliers in the population is $\Pr\{X(1)>X(0)\} = E\{X(1)-X(0)\} = 0.1915$.

Figure 1 shows the differences between the local parameters that are identified by the SMM estimands, LATE and LRR, and their respective parameters in the treated group, ATEX1 and RRX1. Clearly, the differences are increasing functions of ρ. We take ATEX1 and RRX1 as the comparison here, as these are the parameters estimated if the no effect modification by Z (NEM) holds. The differences are quite substantial for large ρ: for example, if $\rho = 0.5$ the LATE equals 0.0457 and the ATEX1 is equal to 0.0400, a difference of 14%. In terms of risk ratios, the LRR minus 1 equals 0.0634 and the RRX1 minus 1 equals 0.1030, a 62% difference. The magnitude by which NEM is violated is indicated by the difference between ATEX1Z1 and ATEX1Z0 for the additive SMM, and between RRX1Z1 and RRX1Z0 for the multiplicative SMM. Both are relatively small indicating a minor failure of NEM, but the local parameters take quite different values. For the odds ratio, the LOR and ORX1 are quite close: for example, if $\rho = 0.5$ the LOR minus 1 is equal to 0.2018 and the ORX1 minus 1 equal to 0.1926, only a small difference of 4.8%. Interestingly, the estimand of the double-logistic SMM estimator, VG, tracks

the odds ratio OR quite closely here, but not LOR or ORX1: at $\rho = 0.5$, VG minus 1 is equal to 0.1745, a 10% difference from ORX1 and 15.6% difference from LOR.

Figure 2 displays the same plots for $\gamma = -1$ and $\delta = 0.615$. We now have $E\{X(0)\} = 0.159$, so there is more compliance in the control group, while the complier proportion remains 0.1915. Here we find values of LATE and ATEX1 at $\rho = 0.5$ of 0.0415 and 0.0341 respectively, a difference of 21%. For the LRR and RRX1 (minus 1) the respective values are 0.0639 and 0.0458, a difference of 40%. In contrast, the LOR and ORX1 are virtually identical in this case for all ρ, with VG now tracking both quite closely.

In Figure 3 we set $\gamma = -1$ and $\delta = 1.208$ to give $E\{X(0)\} = 0.023$. These parameter values generate data for which treatment exclusion might be expected to provide a good approximation. As expected, the local parameters LATE and LRR are very close to ATEX1Z1 and RRX1Z1 respectively, and to ATEX1 and RRX1 too. The LOR and VG are in this case identical to ORX1Z1 and ORX1.

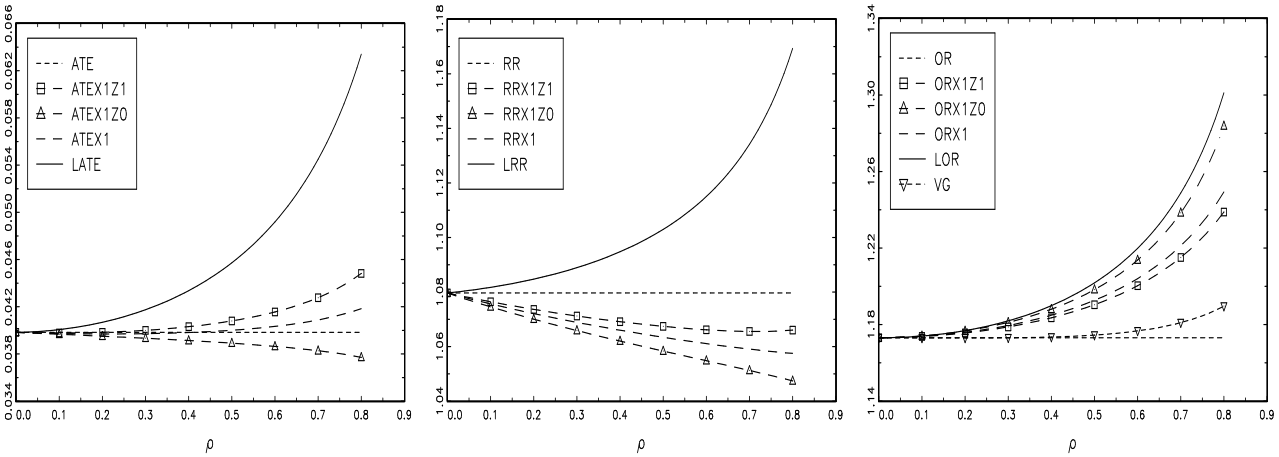

Figure 1. $\gamma = 0$, $\delta = 0.500$

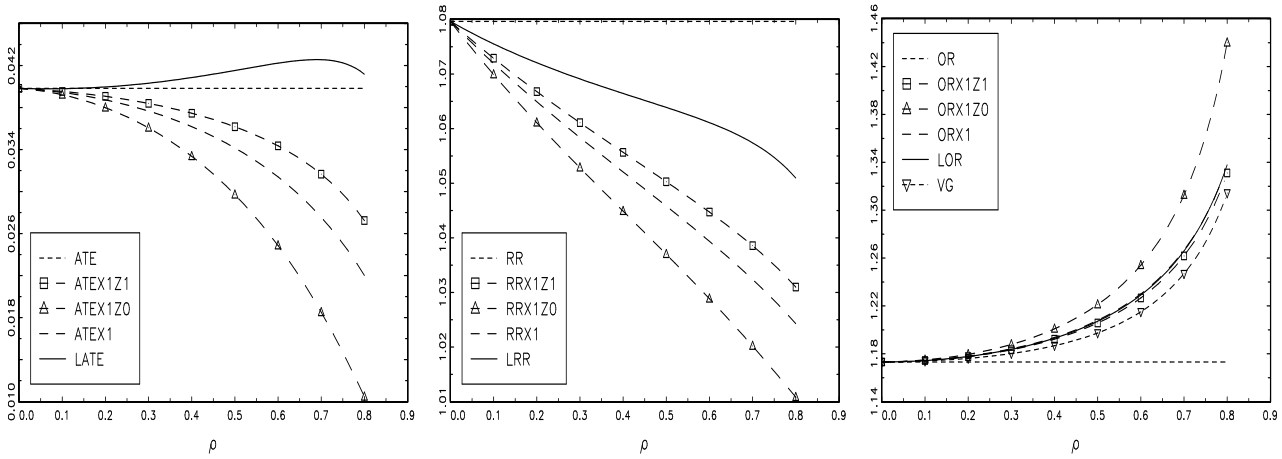

Figure 2. $\gamma = -1$, $\delta = 0.615$

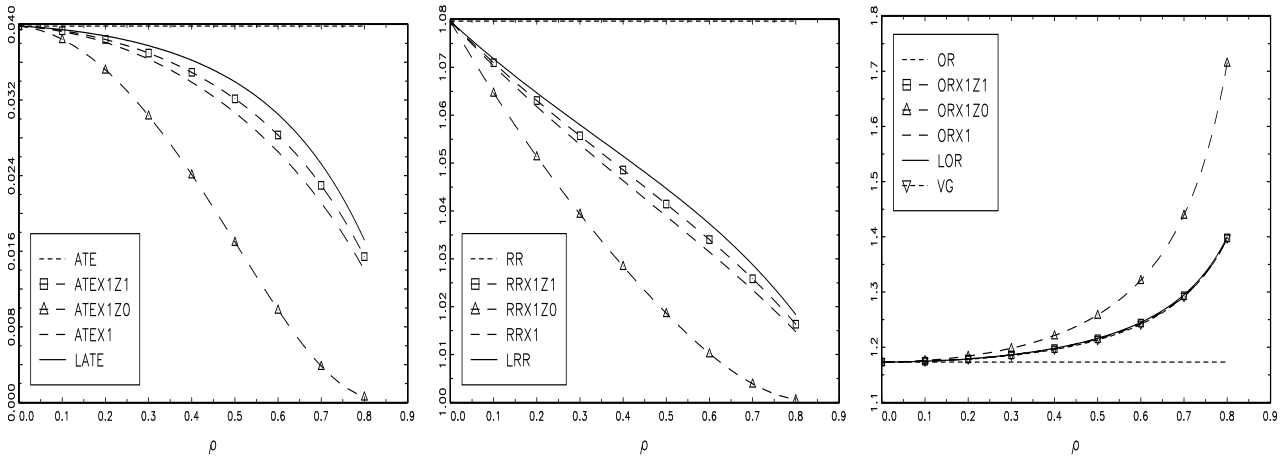

Figure 3. $\gamma = -2$, $\delta = 1.208$

6.2 Example 2

Didelez et al. (2008) considered a more complex model for generating non-ignorable non-compliance using a logistic structural model. In our notation, it is written

$$X(z) = I\{D < \text{expit}(\alpha_1 + z\alpha_2 + U\alpha_3 + zU\alpha_4)\} \quad (12)$$

$$Y(x) = I\{W < \text{expit}(\beta_1 + x\beta_2 + U\beta_3 + xU\beta_4)\} \quad (13)$$

where D and W are independent random variables uniformly distributed on $(0, 1)$. An equivalent expression for (12) is $E\{X(z)|U = u\} = \text{expit}(\alpha_1 + z\alpha_2 + u\alpha_3 + zu\alpha_4)$ and for (13) is $E\{Y(x)|U = u\} = \text{expit}(\beta_1 + x\beta_2 + u\beta_3 + xu\beta_4)$. Both models contain interaction

terms allowing the effect of latent U to vary depending on z and x, respectively. There are heterogeneous treatment effects (on the latent scale) if $\beta_4 \neq 0$ but this poses no problems as SMMs do not constrain treatment effects to be homogeneous, or indeed place any constraints on the form of treatment effect heterogeneity. More importantly, however, the monotonicity assumption $X(1) \geq X(0)$ holds only if $\alpha_4 = 0$, and monotonicity is crucial for identification of local causal effects.

We generate data according to models (12) and (13), setting the parameters $\alpha_1 = 0$, $\alpha_2 = 0.5$, $\alpha_3 = 2$, $\beta_1 = 0$, $\beta_2 = 0.3$, $\beta_3 = 2$, and specifying $U \sim N(0,1)$. Table 1 contains Monte Carlo estimates, based on 1000 replications, of the mean and standard deviation of the causal parameters defined above and the estimands of three local effect estimators. For the additive and multiplicative SMMs we use (4) and (5), and use the consistent estimator of LOR described in Section 5 rather than the double-logistic SMM. To minimise the impact of finite sample bias and maintain our focus on consistency, we generated samples of size $500,000$.

Table 1. Causal parameters and SMM estimates for logistic model

	(1)		(2)		(3)	
	$\alpha_4 = \beta_4 = 0$		$\alpha_4 = 0, \beta_4 = 1$		$\alpha_4 = 1, \beta_4 = 1$	
	mean	*stdev*	mean	*stdev*	mean	*stdev*
ATE	0.0453	*0.0003*	0.0344	*0.0003*	0.0344	*0.0004*
ATEX1Z1	0.0455	*0.0005*	0.0625	*0.0007*	0.0708	*0.0008*
ATEX1Z0	0.0437	*0.0006*	0.0658	*0.0007*	0.0658	*0.0007*
ATEX1	0.0447	*0.0004*	0.0640	*0.0005*	0.0684	*0.0005*
LATE	0.0574	*0.0012*	0.0403	*0.0013*	0.0923	*0.0016*
Add. SMM	0.0575	*0.0188*	0.0403	*0.0180*	0.1148	*0.0225*
RR	1.0907	*0.0006*	1.0689	*0.0007*	1.0689	*0.0007*
RRX1Z1	1.0682	*0.0009*	1.0936	*0.0011*	1.1007	*0.0012*
RRX1Z0	1.0627	*0.0008*	1.0944	*0.0011*	1.0944	*0.0011*
RRX1	1.0656	*0.0006*	1.0940	*0.0008*	1.0977	*0.0008*
LRR	1.1219	*0.0028*	1.0855	*0.0029*	1.1376	*0.0027*
Mult. SMM	1.1229	*0.0426*	1.0863	*0.0399*	1.1523	*0.0328*
OR	1.1994	*0.0014*	1.1479	*0.0016*	1.1478	*0.0016*
ORX1Z1	1.2377	*0.0032*	1.3467	*0.0045*	1.4457	*0.0057*
ORX1Z0	1.2420	*0.0034*	1.3984	*0.0053*	1.3982	*0.0052*
ORX1	1.2394	*0.0023*	1.3691	*0.0033*	1.4227	*0.0039*
LOR	1.2584	*0.0061*	1.1749	*0.0061*	1.5802	*0.0126*
LOR estimator	1.2629	*0.0958*	1.1781	*0.0852*	2.2247	*0.3699*

Notes: 1000 Monte Carlo replications; sample size $500,000$.

The results for $\alpha_4 = \beta_4 = 0$ are given in column 1 and are similar to the results found above. When we introduce an extra source of treatment heterogeneity by setting $\beta_4 = 1$ (column 2), we see again that all three estimators are very close to the local parameters. For the odds ratio, it can also be seen that treatment effect heterogeneity has here exacerbated the difference between the local and treated group odds ratios LOR and ORX1, being 1.175 and 1.369 respectively.

When the monotonicity assumption is violated by further setting $\alpha_4 = 1$ (column 3), we see that all three SMM estimates diverge from the local parameters, with the LOR estimator especially very poorly behaved. In this example, the divergence between the target parameters, the causal effects in the treated group, and the estimates of the local treatment effects gets more pronounced.

7 Conclusions

We have shown that for the identification of causal effects on binary outcomes with non-compliance using structural mean models, additional, non-standard assumptions need to be made. Causal effects like the average treatment effect for the treated can be identified when the control group has no access to treatment. However, we show that the common identifying assumption of no effect modification by randomisation does not hold in general for binary outcomes, meaning that treatment effects for the treated are not identified when there is also treatment non-compliance in the control group. To some extent this has already been recognised: in the context of logistic and probit SMMs, NEM is "unrealistic because subpopulations [defined by randomisation] are likely to be quite different with regard to modifiers of the effect of active treatment on the outcome of interest" (Robins and Rotnitzky, 2004, p. 778); and in the context of generalised SMMs, "a sensitivity analysis which investigates departures from this homogeneity [NEM] assumption is recommended" (Vansteelandt and Goetghebeur, 2003, p. 829). However, our study represents an investigation into the full implications of this failure for all SMMs, and through our numerical study its impact on practice.

While applications of logistic SMMs so far have been to designs satisfying the treatment exclusion requirement, not all randomised controlled trials satisfy it. Moreover, SMMs can also be applied to observational studies without a randomisation indicator but where Z is chosen to satisfy the assumptions of an econometric instrumental variable (e.g., Angrist et al., 1996). For such applications, such as genetic instruments used within the 'Mendelian randomisation' context (e.g., Didelez and Sheehan, 2007), any such assumption is likely to be implausible.

If selection into treatment is monotonic, the additive and multiplicative SMM estimands correspond to a local, or complier, causal effects that, as we show, can be quite different from treatment effects for the treated. Caution is therefore in order when inter-

preting SMM estimates when outcomes are binary, with the issue of monotonicity and interpretation of complier average affects paramount when treatment exclusion cannot be obtained in the control group. Finally, we show that the double-logistic SMM estimator is not consistent for the local odds ratio without further assumptions, but that an alternative estimator is available (Abadie, 2003; van der Laan et al., 2007)

References

[1] Abadie, A. (2003), Semiparametric instrumental variable estimation of treatment response models, Journal of Econometrics 113, 231-263.

[2] Angrist, J.D. (2001), Estimation of limited dependent variable models with dummy endogenous regressors: simple strategies for empirical practice, Journal of Business and Economic Statistics 19, 2-16.

[3] Angrist, J.D., Imbens, G.W. and Rubin, D.B. (1996), Identification of causal effects using instrumental variables, Journal of the American Statistical Association 91, 444-455.

[4] Clarke, P. and Windmeijer, F. (2009), Instrumental variable estimators for binary outcomes, CMPO Working Paper 09/209, Centre for Market and Public Organisation, University of Bristol, UK.

[5] Didelez, V. and Sheehan, N. (2007), Mendelian randomization as an instrumental variable approach to causal inference, Statistical Methods in Medical Research 16, 309-330.

[6] Didelez, V., Meng, S. and Sheehan, N. (2008), On the bias of IV estimators for Mendelian Randomisation, submitted manuscript.

[7] Fischer-Lapp, K. and Goetghebeur, E. (1999), Practical properties of some structural mean analyses of the effect of compliance in randomized trials, Controlled Clinical Trials 20, 531-546.

[8] Goetghebeur, E. and Lapp, K. (1997), The effect of treatment compliance in a placebo-controlled trial: regression with unpaired data, Applied Statistics 46, 351-364.

[9] Goetghebeur, E. and Vansteelandt, S. (2005), Structural mean models for compliance analysis in randomized clinical trials and the impact on measures of exposure, Statistical Methods in Medical Research 14, 397-415.

[10] Greenland, S. (2000), An introduction to instrumental variables for epidemiologists, International Journal of Epidemiology 29, 722-729.

[11] Hernán, M.A. and Robins, J.M. (2006), Instruments for causal inference: an epidemiologist's dream?, Epidemiology 17, 360-372.

[12] Imbens, G.W. and Angrist,. J. (1994), Identification and estimation of local average treatment effects, Econometrica 62, 467-476.

[13] Meng, X.L. (1994), Multiple-imputation inferences with uncongenial sources of input (with discussion), Statistical Science 9, 538-573.

[14] Moodie, E.E.M., Platt, R.W. and Kramer, M.S. (2009), Estimating response-maximized decision rules with applications to breastfeeding, Journal of the American Statistical Association 104, 155-165.

[15] Pearl, J. (2000), Causality: Models, Reasoning and Inference, Cambridge: Cambridge University Press.

[16] Robins, J.M. (1989), The analysis of randomised and non-randomised AIDS treatment trials using a new approach to causal inference in longitudinal studies, in: Sechrest, L., Freeman, H. and Mulley, A. (eds.), Health Service Research Methodology: A Focus on AIDS, 113-159, Washington, DC: US Public Health Service, National Center for Health Services Research.

[17] Robins, J.M. (1994), Correcting for non-compliance in randomized trials using structural nested mean models, Communications in Statistics - Theory and Methods 23, 2379-2412.

[18] Robins, J.M. (1999), Marginal structural models versus structural nested models as tools for causal inference, in: Halloran, E. and Berry, D. (eds.), Statistical Methods in Epidemiology: The Environment and Clinical Trials, 95-134, New York: Springer.

[19] Robins, J.M. and Rotnitzky, A. (2004), Estimation of treatment effects in randomised trials with non-compliance and a dichotomous outcome using structural mean models, Biometrika 91, 763-783.

[20] Robins, J.M., Rotnitzky, A. and Scharfstein, D.O. (1999), Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models, in: Halloran, E. and Berry, D. (eds.), Statistical Methods in Epidemiology: The Environment and Clinical Trials, 1-92, New York: Springer.

[21] Tanaka, Y., Matsuyama, Y. and Ohashi, Y. (2008), Estimation of the treatment effect adjusting for treatment changes using the intensity score method: application to a large primary prevention study for coronary events (MEGA study), Statistics in Medicine 27, 1718-1733.

[22] Ten Have, T.R., Elliott, M.R., Joffe, M., Zanutto, E. and Datto, C. (2004), Causal models for ranomized physician encouragement trials in treating primary care depression, Journal of the American Statistical Association 99, 16-25.

[23] van der Laan, M.J., Hubbard, A. and Jewell, N.P. (2007), Estimation of treatment effects in randomized trials with non-compliance and a dichotomous outcome, Journal of the Royal Statistical Society, Series B 69, 463-482.

[24] Vansteelandt, S. and Goetghebeur, E. (2003), Causal inference with generalized structural mean models, Journal of the Royal Statistical Society, Series B 65, 817-835.

[25] Witteman, J.C.M., D'Agostino, R.B., Stijen, T., Kannel, W.B., Cobb, J.C., de Ridder, M.A.J., Hofman, A. and Robins, J.M. (1998), G-estimation of causal effects: isolated systolic hypertension and cardiovascular death in the Framingham Heart Study, American Journal of Epidemiology 148, 390-401.