

Clarke, Paul; Windmeijer, Frank

Working Paper

Identification of causal effects on binary outcomes using structural mean models

cemmap working paper, No. CWP02/10

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Clarke, Paul; Windmeijer, Frank (2010) : Identification of causal effects on binary outcomes using structural mean models, cemmap working paper, No. CWP02/10, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.1920/wp.cem.2010.0210>

This Version is available at:

<https://hdl.handle.net/10419/64758>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Identification of causal effects on binary outcomes using structural mean models

Paul Clarke
Frank Windmeijer

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP02/10

Identification of Causal Effects on Binary Outcomes Using Structural Mean Models

Paul Clarke* and Frank Windmeijer**
CMPO, University of Bristol, UK

*email: Paul.Clarke@bristol.ac.uk

**email: F.Windmeijer@bristol.ac.uk

January 2010

Abstract

Structural mean models (SMMs) were originally formulated to estimate causal effects among those selecting treatment in randomised controlled trials affected by non-ignorable non-compliance. It has already been established that SMM estimators identify these causal effects in randomised placebo-controlled trials where no-one assigned to the control group can receive the treatment. However, SMMs are starting to be used for randomised controlled trials without placebo-controls, and for instrumental variable analysis of observational studies; for example, Mendelian randomisation studies, and studies where physicians select patients' treatments. In such scenarios, identification depends on the assumption of no effect modification, namely, the causal effect is equal for the subgroups defined by the instrument. We consider the nature of this assumption by showing how it depends crucially on the underlying causal model generating the data, which in applications is almost always unknown. If its no effect modification assumption does not hold then an SMM estimator does not estimate its associated causal effect. However, if treatment selection is monotonic we highlight that additive and multiplicative SMMs do identify local (or complier) causal effects, but that the double-logistic SMM estimator does not without further assumptions. We clarify the proper interpretation of inferences from SMM estimators using a data example and simulation study.

Key Words: Structural Mean Models; Identification; Local Average Treatment Effects; Complier Average Causal Effects.

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 treat-

ment 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 the ‘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. More precisely, we consider the conditions under which each SMM estimator identifies its target causal parameter, and the consequences if these conditions do not hold. Until recently, SMMs have mainly been applied to randomised placebo-controlled trials for which the identification issue is fairly straightforward. However, SMMs can be applied to other types of randomised controlled trial, and more generally to the causal analysis of observational studies using instrumental variables. For these more general designs, the usual identification assumption for a SMM estimator is

‘no effect modification’ (NEM) by randomisation (Hernán and Robins, 2006). We highlight how identification depends crucially on the unknown data generating process from which the data arise. While data generating processes satisfying the various NEM assumptions do exist, not all data generating processes satisfy NEM, including simple and widely known mechanisms for binary data like the bivariate probit model.

Another question we address in this paper is: what causal parameter is being identified if the NEM assumption does not hold? As such, we highlight previous results showing that additive and multiplicative SMMs identify ‘local’, or ‘complier’, causal effects under the alternative assumption that patients’ treatment selection is monotonic (e.g., Angrist et al., 1996). Local effects are special cases of principal strata and thus widely used in biostatistics (Frangakis and Rubin, 2002). However, we also highlight that the double-logistic SMM does not identify the local odds ratio under monotonic selection, but that the local odds ratio can be identified using an alternative estimator.

The remainder of this paper is organised as follows. In Section 2, we review the potential outcomes causal framework within which SMMs are specified, and the three important SMMs considered in this paper. In Section 3, we consider the identification of each SMM’s causal effect for randomised placebo-controlled trial designs, before going on in Section 4 to consider identification for more general designs. An alternative identification strategy based on monotonic treatment selection is considered in Section 5. Finally, in Section 6 we consider a data example and present some numerical results to illustrate the potential impact on results if NEM does not hold, before making our concluding remarks in Section 7.

2 Structural Mean Models

2.1 Potential Outcomes

Before introducing SMMs, 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; $X \in \{0, 1\}$ 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 $\{(y_i, x_i, z_i) : i = 1, \dots, n\}$ 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. Similarly, let $X(z)$ be the potential treatment that would be obtained if treatment assignment was set to z by external intervention.

Five important conditions for identification of causal effects 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; 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); and the 'independence assumption' implying that Z is independent of the potential treatments and outcomes $\{X(0), X(1), Y(0), Y(1)\}$.

More generally, Z can be any instrumental variable (IV) satisfying the assumptions we have just introduced. The scope of SMMs is thus broader than randomised controlled trials and encompasses observational studies too. However, the practical difficulties associated with choosing an IV are well known. In particular, if Z is a randomisation indicator then the independence assumption can be taken for granted, but for observational studies it must be justified and this cannot be done on empirical grounds alone. To maintain focus, all of these assumptions will be taken to hold throughout this paper, and so we assume that a valid IV Z is available to the analyst.

2.2 The Additive and Multiplicative SMMs

For the simple scenario just described, the additive SMM is

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

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

SMM estimators work by exploiting the conditional mean independence (CMI), or randomisation, assumption

$$E \{Y(0) | Z = 1\} = E \{Y(0) | Z = 0\}, \quad (1)$$

which follows from the identification conditions in Section 2.1. Under the additive SMM, (1) can be rewritten as

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

from which an estimating equation can be constructed.

The saturated multiplicative SMM for the same scenario is defined as

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

The parameters of the multiplicative SMM are

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

and

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

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

$$E[Y \exp\{-(\theta_0 + \theta_1 Z) X\} | Z = 1] = E\{Y \exp(-X\theta_0) | Z = 0\}. \quad (3)$$

It is clear that neither set of SMM parameters is identified by its corresponding moment condition because both constitute systems with two unknowns and one equation. Therefore, further assumptions are required to identify the SMM parameters. Hernán and Robins (2006) highlight the role of the no effect modification (NEM) by Z assumption. Each SMM has its own distinct NEM assumption: for the additive SMM, it corresponds to constraining $\psi_1 = 0$, and for the multiplicative SMM it corresponds to $\theta_1 = 0$. Under

NEM, there is only one unknown in (2) and (3) and the usual target parameters are identified, namely, for the additive SMM

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

the average treatment (or causal) effect among the treated, and for the multiplicative SMM

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

the risk ratio among the treated.

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

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

and

$$\widehat{\exp(\theta_0)} = 1 - \frac{E(Y|Z = 1) - E(Y|Z = 0)}{E \{(1 - X)Y|Z = 1\} - E \{(1 - X)Y|Z = 0\}}, \quad (5)$$

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 these NEM assumptions for both SMM estimators is thus apparent. It is the validity of this assumption in the binary case that we will consider in more detail in Section 4.

2.3 The Double-Logistic SMM

Robins et al. (1999) proposed the logistic SMM

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

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

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

and we can write the logistic SMM as

$$\text{logit}\{E(Y|X, Z)\} - \text{logit}\{E(Y(0)|X, Z)\} = \xi_0 X,$$

where $\text{logit}(a) = \log\{a/(1-a)\}$.

The logistic SMM is considered separately here because Robins (1999) showed that no G-estimator for ξ_0 can be constructed. 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’

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

which is indexed by parameter vector η . The double-logistic estimator is based on specifying $m_\eta(X, Z)$ to be logistic. A drawback to this approach in more general settings, as acknowledged by Vansteelandt and Goetghebeur (2003), is that both the SMM and the association model cannot both be logistic, so the double-logistic SMM is ‘uncongenial’ in the sense described by Meng (1994). However, it has been shown that this is not a problem for the saturated logistic SMMs considered here (Babanezhad et al., 2009). The double-logistic SMM estimator is then the solution to the moment condition

$$E[\text{expit}\{\eta_{00} + \eta_{01} + (\eta_{10} + \eta_{11} - \xi_0)X\} | Z=1] = E[\text{expit}\{\eta_{00} + (\eta_{10} - \xi_0)X\} | Z=0], \quad (6)$$

where an estimate of $(\eta_{00}, \eta_{10}, \eta_{01}, \eta_{11})$ is obtained at the first stage by fitting the saturated logistic association model $m_\eta(X, Z) = \text{expit}(\eta_{00} + Z\eta_{01} + X\eta_{10} + ZX\eta_{11})$, and $\text{expit}(a) = \exp(a)/\{1 + \exp(a)\}$.

3 SMMs for Randomised Placebo-controlled Designs

There is a wide scope for applications of SMMs to randomised placebo-controlled trial designs such as those considered by Greenland (2000), Nagelkerke et al. (2000) and Vansteelandt and Goetghebeur (2003). For these designs, neither compliers nor non-compliers randomised to control can receive the treatment because non-compliers ($Z = 0, X = 1$) receive only the placebo, equating to the condition $\Pr(X = 0|Z = 0) = 1$. Cuzick et al. (2007) refer to this as a ‘no contamination’ restriction; it is a special case of the identifying assumptions for binary outcome SMMs described by Robins and Rotnitzky (2004). To analyse placebo-control designs, an additional assumption of no placebo effect is also needed that we herein take to hold.

Under the no contamination restriction, the SMM parameters ψ_0 , θ_0 and ξ_0 are not defined because all three are conditioned on the measure-zero event $\{X = 1, Z = 0\}$. 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

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

for the additive, multiplicative and logistic SMMs, respectively.

Under the no contamination restriction, $E\{Y(0)|Z = 0\}$ is always non-parametrically identified because it equals $E(Y|Z = 0)$. Further, expanding $E\{Y(0)|Z = 1\}$ gives

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

which can be used in conjunction with CMI (1) to show that

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

and the key counterfactual parameters are identifiable from the observed data. Hence,

the estimators of the additive and multiplicative SMM parameters under the no contamination restriction are, respectively,

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

and

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

Finally, the double-logistic SMM estimator is

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

4 No Effect Modification

The role of the no effect modification (NEM) assumption becomes crucial for more general designs. Each SMM has its own distinct NEM assumption, which acts to constrain the causal effects among the treated to be equal for those randomised to treatment and those randomised to control. To take just one example, recall that the additive NEM assumption constrains $\psi_1 = 0$ in the additive SMM, and thus

$$\psi_0 = E\{Y(1) - Y(0)|X=1, Z=0\} = E\{Y(1) - Y(0)|X=1, Z=1\} = \psi_0 + \psi_1.$$

The NEM assumptions for the multiplicative and logistic SMMs can be similarly expressed.

To investigate the validity of NEM for binary outcomes, we make a link between structural models and potential outcomes by following Hernán and Robins (2006, Appendix 1) and, less directly, Pearl (2000). Suppose that the analyst is faced with data from a randomised controlled trial for which the no contamination restriction does not hold. In any application, the observed data and all the counterfactual potential outcomes and potential treatments are realisations from an unknown ‘non-parametric structural equation model’ that satisfies CMI and the constraints set out in Section 2.1. Note that

‘non-parametric’ here does not imply that the true data generating process cannot be parametric, but only that no constraints are placed on its unknown form.

Hence, the potential outcome can be written

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

where indicator function $I(a) = 1$ if a is true and 0 otherwise, and $f_Y^*(x, U)$ is a function that depends on the fixed value of treatment and the latent random variable/vector U . It is usual to interpret U as the combined effect of all unobserved ‘confounding’ variables on the outcome, although it also involves the contributions from other variables which are independent of the exposure selection mechanism; to ensure independence of Z and $Y(x)$ it follows that U must be independent of Z . The potential treatment is similarly defined as $X(z) = I \{f_X^*(z, V) > 0\}$, where V is another latent random variable/vector representing unobservable factors influencing treatment choice; it also follows that V and Z must be independent. If U and V are independent then non-compliance is ignorable, otherwise it is non-ignorable. For fixed x , it is U that determines whether the potential outcome is zero or one for a particular patient, with V playing the same role for $X(z)$. This set-up straightforwardly extends to continuous potential treatments by dropping the indicator function and specifying $X(z) = f_X^*(z, V)$. We discuss a simple model with continuous treatments in the Appendix.

This class of structural models is extremely general because $f_Y^*(x, u)$ can be any function generating, for example, non-linear or heterogeneous treatment effects. However, it does not include models where the joint support of U and V depends on the observed variables. For example, this excludes models like $Y(x) = f_Y^*(x) + U$ and $Y(x) = f_Y^*(x)U$ for which $Y = f_Y^*(X) + U$ and $Y = f_Y^*(X)U$, respectively, because the support of U clearly depends on X in both cases to ensure the outcome lies in $\{0, 1\}$. Thus, structurally implausible data generating models where the support of U depends on a variable to which it is antecedent are excluded from consideration.

Crucially, all of the SMM parameters are functions of $E\{Y(x)|X, Z\}$ and so can be written in terms of the underlying structural model using

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

An advantage of defining the class of models in this way is that all its members satisfy the CMI assumption, which can be shown by expanding (1) and using the identity $\Pr\{f_X^*(z, V) > 0\} = E(X|Z = z)$. We can therefore focus on each of the NEM assumptions. For a specific example, consider the family of simple parametric structural models with (U, V) a bivariate continuous random vector related to the potential outcomes by

$$Y(x) = I(\alpha + \beta x - U > 0), \quad X(z) = I(\gamma + \delta z - V > 0), \quad (8)$$

where $E(U) = E(V) = 0$ and (U, V) has distribution function $F_\rho(u, v)$, with ‘correlation’ parameter ρ indexing all non-zero moments involving products of U and V . In this case,

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

where $G(v)$ is the marginal distribution function of V . Clearly, if non-compliance is ignorable then $\rho = 0$ and all three NEM assumptions automatically hold. However, if $\rho \neq 0$ then none of the NEM assumptions will necessarily hold. For example, if F is the distribution function for a zero-mean, unit-variance bivariate normal distribution then (8) is known as the bivariate probit model; for the additive SMM,

$$\begin{aligned} \psi_0 + \psi_1 &= \frac{F_\rho(\alpha + \beta, \gamma + \delta)}{G(\gamma + \delta)} - \frac{F_\rho(\alpha, \gamma + \delta)}{G(\gamma + \delta)} \\ &\neq \frac{F_\rho(\alpha + \beta, \gamma)}{G(\gamma)} - \frac{F_\rho(\alpha, \gamma)}{G(\gamma)} = \psi_0 \end{aligned}$$

almost everywhere; for the multiplicative SMM

$$\begin{aligned}\exp(\theta_0 + \theta_1) &= \frac{F_\rho(\alpha + \beta, \gamma + \delta) / G(\gamma + \delta)}{F_\rho(\alpha, \gamma + \delta) / G(\gamma + \delta)} \\ &\neq \frac{F_\rho(\alpha + \beta, \gamma) / G(\gamma)}{F_\rho(\alpha, \gamma) / G(\gamma)} = \exp(\theta_0)\end{aligned}$$

almost everywhere; and for the logistic SMM

$$\begin{aligned}\exp(\xi_0 + \xi_1) &= \frac{F_\rho(\alpha + \beta, \gamma + \delta) / \{G(\gamma + \delta) - F_\rho(\alpha + \beta, \gamma + \delta)\}}{F_\rho(\alpha, \gamma + \delta) / \{G(\gamma + \delta) - F_\rho(\alpha, \gamma + \delta)\}} \\ &\neq \frac{F_\rho(\alpha + \beta, \gamma) / \{G(\gamma) - F_\rho(\alpha + \beta, \gamma)\}}{F_\rho(\alpha, \gamma) / \{G(\gamma) - F_\rho(\alpha, \gamma)\}} = \exp(\xi_0)\end{aligned}$$

almost everywhere. As such, the additive, multiplicative and double-logistic SMMs do not estimate ψ_0 , $\exp(\theta_0)$ and $\exp(\xi_0)$, respectively, if the data come from a causal model which is closely approximated by the bivariate probit structural model. However, as we now discuss, this does not mean that families of distributions for which a NEM assumption holds cannot be found.

The class of structural models is far broader than that defined by (8): it can be extended to allow for non-linear effects of treatment and IV on the latent scale, and the latent variables can be multivariate with semi-continuous or even discrete distributions. However, we contend that each NEM assumption is highly restrictive. We illustrate this point by focussing on the logistic SMM for which Babanezhad et al. (2009) show how data can be generated to satisfy both the logistic SMM and its NEM assumption without specifying the underlying structural model; their approach is based on the SMM parameterisation developed by Robins and Rotnitzky (2004). Data are generated as follows: first, generate X as Bernoulli with success probability $E(X|Z) = \text{expit}(\gamma + \delta Z)$; second, generate the treatment-free outcomes as Bernoulli with success probability

$$E\{Y(0) | X, Z\} = \text{expit}(\beta_0 + \beta_1 X + \beta_2 Z);$$

and finally, generate observed outcomes using success probability

$$E(Y | X, Z) = \text{expit}\{\beta_0 + (\beta_1 + \xi_0) X + \beta_2 Z\},$$

where ξ_0 is the target parameter of a logistic SMM model satisfying the NEM assumption $\xi_1 = 0$. Finally, to identify ξ_0 , the β parameters must be constrained so that

$$\begin{aligned} & \text{expit}(\beta_0 + \beta_1 + \beta_2) E(X|Z = 1) + \text{expit}(\beta_0 + \beta_2) E\{(1 - X)|Z = 1\} \\ & = \text{expit}(\beta_0 + \beta_1) E(X|Z = 0) + \text{expit}(\beta_0) E\{(1 - X)|Z = 0\} \end{aligned}$$

to ensure that the CMI assumption is satisfied.

An example of a structural model satisfying this data generating process can be written as

$$Y(x) = I(\alpha + \xi_0 x + U > 0), X(z) = I(\gamma + \delta z + V > 0),$$

where

$$U = (\beta_0 - \alpha) + \beta_1 X + \beta_2 Z + W,$$

and V and W both have standard logistic marginal distributions, are mutually independent and independent of Z . This model does not fit into the structural set-up defined above because U and V are associated only indirectly through X and Z , and U is clearly not independent of Z , but it does show that the family of models satisfying the logistic NEM is very restrictive: the integral of $Y = I\{f_Y^*(X, U) > 0\}$ with respect to the conditional distribution of U given X and Z (and an appropriate measure) must be logistic, whereas the class of structural models we consider places no such restriction on this conditional distribution. The existence of other families is straightforward to show: for example, the family of structural models satisfying NEM for the probit SMM, where the conditional distribution must be normal rather than logistic, does not satisfy logistic NEM (Goetghebeur and Vansteelandt, 2005). Furthermore, these models do not automatically satisfy the other NEM assumptions and so ψ_0 and $\exp(\theta_0)$ may not be estimated by the additive and multiplicative SMMs, even if $\exp(\xi_0)$ is estimated by the double-logistic SMM. We discuss this point further in the Appendix.

5 Monotonic Selection

We have argued that the families of models for binary outcome and treatment satisfying the additive, multiplicative or logistic NEM assumptions are very restrictive, and so the problem facing the practitioner analysing data using any SMM is ascertaining whether the true data generating process is closely approximated by a model that satisfies its NEM assumption. Establishing this with any certainty depends on application-specific background knowledge, and is extremely difficult - if not impossible - to do. In many applications, it may be a reasonable working assumption, but more generally the practitioner may not be prepared to make it. As such, we now discuss an alternative assumption to NEM, namely, monotonic selection of treatment by patients, under which local, or complier, causal effects can be estimated. 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

$$X(1) \geq X(0) \tag{9}$$

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 all patients exist in two universes, one in which they are randomised to control, and another in which they are randomised to treatment. So the ‘no defiers’ assumption corresponds to saying that, while patients can disobey their treatment assignments in one or other of these universes, they cannot disobey their assignments in both. For example, the simple structural model described in Section 4 is monotonic because

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

if $\delta > 0$.

While 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),

$$\text{LATE} = E \{Y(1) - Y(0) | X(1) > X(0)\}, \quad (10)$$

which is also called the ‘complier average causal effect’ (CACE). Note that no contamination 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),

$$\text{LRR} = \frac{E \{Y(1) | X(1) > X(0)\}}{E \{Y(0) | X(1) > X(0)\}}, \quad (11)$$

see also Greenland (2000) and Hernán and Robins (2006).

The local odds ratio (LOR) is defined as

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

Our numerical examples below illustrate that the double-logistic estimator based on (6) is biased for the LOR under monotonic selection. Clarke and Windmeijer (2009, Appendix 3) show that the double-logistic estimator is not consistent for the LOR under monotonic selection unless $E \{Y(1) | X(1) = X(0) = 1\} = E \{Y(1) | X(1) > X(0)\}$. However, a consistent estimator for the LOR is available. Abadie (2003) proposes an estimator and van der Laan et al. (2007) note how this estimator can be derived based on the relative risk

estimator (5): first calculate $\widehat{\exp(\theta_0)}$ as per usual, then recode the outcome variable as $Y^* = 1 - Y$ and calculate $\widehat{\exp(\theta_0^*)}$ replacing Y by Y^* in (5), then the ratio $\widehat{\exp(\theta_0)}/\widehat{\exp(\theta_0^*)}$ is consistent for the LOR by symmetry of the relative risk. We refer to this below as the ‘LOR estimator’.

6 Numerical Examples

6.1 Data Example

To summarise and make the implications of these results concrete, consider the following example of an observational study to which instrumental variables have been used to obtain causal inferences. A study of patients attending clinical practice was carried out to assess if the ‘Cox-2’ inhibitor treatment performed better than the standard, non-selective non-steroidal anti-inflammatory (NSAID) treatment in preventing the unwanted side-effect of gastrointestinal bleeding after sixty days’ follow-up (Brookhart et al., 2006). The analysis here is based on a subset of 37842 patients who took part in the original study, of which 26407 were allocated Cox-2 and 11435 were allocated NSAID by their physicians (Brookhart and Schneeweiss, 2007; Babanezhad et al., 2009).

In our set-up, Y is 1 if the patient experiences gastrointestinal bleeding within 60 days of being treated and 0 otherwise; and X is 1 if the patient receives Cox-2, and 0 otherwise. The IV Z for each patient is taken to be the treatment allocated by the prescribing physician to the preceding patient. Brookhart et al. (2006) originally proposed the use of physician preference for this study. We take Z to be a valid IV and refer the reader to Hernán and Robins (2006) for a detailed discussion of how well physician preference satisfies the conditions set out in Section 2.1.

We fit the additive, multiplicative and logistic SMMs to these data, along with the naive logistic regression of Y on X and estimate the local odds ratio as described above; two-tailed 95-percent percentile confidence intervals are also calculated based on 100 non-

parametric bootstrap samples. The results are displayed in Table 1. The naive odds ratio based on the logistic model is 1.032, indicating a negative effect of Cox-2 over NSAIDs in the trial; the confidence interval (CI) is (0.80, 1.37) and includes 1, which indicates that there is insufficient evidence to reject the hypothesis that the treatments are the same.

[TABLE 1 ABOUT HERE]

The naive odds ratio cannot be interpreted as a causal effect but only as a measure of association because we hypothesise that physicians allocate Cox-2 treatment based on unobserved factors which could be associated with the risk of gastrointestinal bleeding. Hence, we use SMMs in order to estimate causal effects among those treated with Cox-2 inhibitors. To recap, for a specific SMM we know that its associated NEM assumption is required for identification of the causal effect, but in Section 4 we showed that it does not always hold. However, both the additive and multiplicative SMMs do identify local causal effects if physicians' treatment selection is monotonic. In this example, monotonicity corresponds to the assumption that no physicians who prescribe Cox-2 for patients after prescribing NSAID for their previous patients ($X(0) = 1$) would have prescribed NSAID for the same patients had they instead (counterfactually) prescribed Cox-2 to their previous patients ($X(1) = 0$). As such, unless we know that the additive NEM or monotonicity assumption holds, we cannot know if the estimate based on the additive SMM ($\widehat{\psi}_0 = -0.0092$) can be interpreted as the average treatment effect among the treated, or as the local treatment effect. However, the effect itself is clearly indicating less risk of gastrointestinal bleeding as the confidence interval excludes 0.

The same scenario holds for the multiplicative SMM, but here $\exp(\widehat{\theta}_0) = -0.176$ and so is out of the valid range for a risk ratio. Out of range estimates are not uncommon for moment-based estimators like these. If the multiplicative NEM assumption holds then this could be because of sampling variability: although the sample size is large, the gastrointestinal bleeding event is rare (fewer than 250 patients have events) and is

sensitive to sampling variability. Alternatively, if the multiplicative NEM assumption has failed then the negative risk ratio may indicate a failure of the monotonicity assumption. The estimate again indicates a positive effect of Cox-2 inhibitors because the CI does not include one

The double-logistic SMM estimate (denoted DL SMM in Table 1) - using the saturated association model described in Section 2.3 - is $\widehat{\exp(\xi_0)} = 0.029$ (CI: 0.01, 0.73), which again indicates a positive effect of Cox-2 inhibitors. Inferences can be made about $\exp(\xi_0)$ only if the logistic NEM assumption holds. If one is not prepared to believe that the logistic NEM assumption is even approximately correct, an alternative is to assume monotonicity and use the local odds ratio estimator from Section 5. Here it is estimated to be -0.174 , which is very close to the estimate for the multiplicative SMM and so again out of range. (Note that we would expect these estimates to be close because the gastrointestinal bleeding event is rare and so the any odds ratio approximates the risk ratio closely.) In this example, the out of range estimate again raises some doubt as to whether treatment selection is monotonic; a more likely explanation, perhaps, is that the logistic NEM assumption approximately holds and the double-logistic SMM estimate can be interpreted as evidence of a substantial positive effect of Cox-2 inhibitors among the patients to which it was allocated. The inherent problem is that these questions cannot be answered on the basis of the available data. Thus this should be interpreted as a sensitivity analysis in which we find some degree of robustness because a positive effect of Cox-2 is inferred using all of the causal estimators; in addition, we know that previously conducted randomised controlled trials have also found positive effects of Cox-2 (Brookhart et al., 2006).

To demonstrate further the important role of the different NEM assumptions, we now conduct two analyses where we know the true structural model generating the data. We first look at a scenario where the data come from a bivariate probit model in which

the true treatment selection mechanism is monotonic and the unobserved confounders normally distributed. In this scenario, we can analytically calculate and compare the key causal parameters. In the second setting, we replicate the design of Didelez et al. (2010) in which treatment and outcome data are generated using a logistic model, and assess the key causal parameters and the estimators in a Monte Carlo study. In this example, we also consider the setting where the true selection mechanism is not monotonic. The aim of both studies is to show the impact of misinterpreting SMM estimates.

6.2 Bivariate Probit

The first illustration is based on structural model (8) from Section 3, namely,

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

where here 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 $\Pr(Z = 1) = 0.5$. 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: $\psi_0 + \psi_1 = E\{Y(1) - Y(0) | X = 1, Z = 1\}$ by ATEX1Z1, $\psi_0 = E\{Y(1) - Y(0) | X = 1, Z = 0\}$ 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 ($RRX1Z1 = \exp(\theta_0 + \theta_1)$, $ORX1Z0 = \exp(\xi_0)$, LRR, etc.); for the odds ratio, there is an additional parameter, denoted by DL, corresponding to the estimand of the double-logistic SMM (6).

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. 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 SMM's corresponding NEM assumption 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, DL, tracks the odds ratio OR quite closely here, but not LOR or ORX1: at $\rho = 0.5$, DL 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 DL 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 no contamination 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 DL are in this case identical to ORX1Z1 and ORX1.

[FIGURES 1-3 ABOUT HERE]

6.3 Mixed Logistic

Didelez et al. (2010) considers a more complex model for generating non-ignorable non-compliance using a logistic structural model. In our notation, it is written

$$X(z) = I(\alpha_1 + z\alpha_2 + H\alpha_3 + zH\alpha_4 + V > 0), \quad (12)$$

$$Y(x) = I(\beta_1 + x\beta_2 + H\beta_3 + xH\beta_4 + U > 0), \quad (13)$$

where U and V are independent logistically distributed random variables, and H is unobserved. An equivalent expression to (12) is $E\{X(z)|H = h\} = \text{expit}(\alpha_1 + z\alpha_2 + h\alpha_3 + zh\alpha_4)$ and for (13) is $E\{Y(x)|H = h\} = \text{expit}(\beta_1 + x\beta_2 + h\beta_3 + xh\beta_4)$. Both models contain interaction terms allowing the effect of latent H 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 $H \sim N(0, 1)$ and $P(Z = 1) = 0.5$. Table 2 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). We further present estimation results for the consistent estimator of the LOR described in Section 5 and for the double-logistic SMM, again denoted DL. To minimise the impact of finite sample bias and maintain our focus on consistency, we generated samples of size 500,000. The population parameters ATE, ATEX1 etc. are as defined above, but are here calculated using the generated data samples and approximated by the averages over the 1000 replications. The column denoted *stdev* contains the Monte Carlo standard deviations of the estimates.

[TABLE 2 ABOUT HERE]

The results for $\alpha_4 = \beta_4 = 0$ are given in column 1 and are similar to the results found in the first example above. When we introduce an extra source of treatment heterogeneity by setting $\beta_4 = 1$ (column 2), we see again that the additive, multiplicative and LOR SMM 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. The DL SMM estimator is close to the OR in this case.

When the monotonicity assumption is violated by further setting $\alpha_4 = 1$ (column 3), we see that the 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. The mean of the DL SMM estimator is here also much higher than any of the causal treatment effect parameters.

7 Conclusions

We have highlighted that causal effects on binary outcomes in studies with non-ignorable non-compliance are not always identified by SMMs, and that additional assumptions about the causal process generating the observed data are required. For example, the double-logistic estimator will be valid for its target parameter, the odds ratio $\exp(\xi_0)$, if the true data generating process is well approximated by that described in Babanezhad et al. (2009); a similar situation exists for the additive and multiplicative SMM models. Our examples show that failure of these assumptions can lead to misleading inferences.

Causal parameters can be identified by all three SMMs if the design satisfies the no contamination restriction that the control group has no access to treatment (e.g., randomised placebo-controlled trials). While applications of logistic SMMs have mainly been to designs satisfying the no contamination restriction, not all randomised controlled trials satisfy it. SMMs can also be applied to observational studies without a randomisation indicator but where Z is chosen to satisfy the assumptions of an instrumental variable (e.g., Angrist et al., 1996). For applications such as genetic instruments used within the ‘Mendelian randomisation’ context (e.g., Didelez and Sheehan, 2007), the no contamination restriction is likely to be implausible, as it is for our data example where the instrumental variable for physician treatment selection is the previous prescribing behaviour by the physician. At present, the form of the structural model generating the potential outcomes in applications such as these is completely unknown (although scientific advances may eventually shed some light on its form) and so a naive interpretation of SMM estimates as causal effects must always be qualified.

An alternative assumption is to assume that the mechanism by which patients select treatment is monotonic. Under monotonicity, the additive and multiplicative SMM estimators are valid for local causal effects, but these can be quite different from treatment effects for the treated. Caution is therefore required when interpreting SMM estimates for

binary outcomes if patients in the control group can receive treatment, with the issues of monotonicity and the interpretation of local/complier average effects paramount. When the NEM assumption fails we find that the double-logistic estimator does not estimate the local odds ratio under monotonicity, but an alternative estimator is available that does.

If the practitioner is agnostic about these as reasonable working assumptions then we would recommend he/she performs a sensitivity analysis. In addition to the SMMs discussed here, various causal estimators for binary outcomes based on instrumental variable estimators have been proposed in the literature; see reviews by Babanezhad et al. (2009), Clarke and Windmeijer (2009) and Didelez et al. (2010). Each estimator makes alternative identifying assumptions, and assessing robustness to these assumptions should be regarded as essential. Bounds on causal effects can be calculated for the simple all-binary-variable case without any assumptions further to those set out in Section 2.1 (Balke and Pearl, 1997); for more general problems, a method for calculating bounds for causal effects for a flexible class of structural models has recently been developed by Chesher (2010).

Finally, we note that recent work by van der Laan et al. (2007) has extended the estimating equation approaches developed by Vansteelandt and Goetghebeur (2003) and Robins and Rotnitzky (2004) to incorporate assumptions about two models (both possibly misspecified) for the conditional means of Y and $Y(0)$ given X and Z . The resulting estimating equations are consistent under fairly weak conditions, but inferences must be interpreted carefully; the implications of this approach for practice have yet to be investigated.

Appendix: Structural Mean Models and the Control Function Approach

We begin with a simple structural model where treatment $X(z)$ is a continuous mea-

sure. The model specification is

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

where U and V follow a zero-mean bivariate normal distribution. We can then write

$$U = \zeta V + W,$$

with $\zeta = \text{Cov}(U, V) / \text{Var}(V)$ and $W \sim N(0, \sigma_w^2)$, leading to a probit SMM (e.g., Goetghebeur and Vansteelandt, 2005)

$$\Phi^{-1}\{E(Y|X, Z)\} - \Phi^{-1}\{E(Y(0)|X, Z)\} = \beta^* X,$$

where Φ^{-1} is the inverse distribution function of the standard normal distribution, and

$$E\{Y(0)|X, Z\} = \Phi(\pi_0 + \pi_1 X + \pi_2 Z),$$

with $\pi_0 = (\alpha + \zeta\gamma) / \sigma_w$, $\pi_1 = \zeta / \sigma_w$, $\pi_2 = -\zeta\delta / \sigma_w$, $\beta^* = \beta / \sigma_w$. Clearly, this model satisfies the NEM and the CMI assumptions, and thus the probit SMM estimator is a consistent estimator for β^* .

Rivers and Vuong (1988) study the ML and control function estimator for the same probit model. The control function estimator is a two-stage estimator in which stage one involves obtaining ordinary least squares estimates of $\hat{\gamma}$ and $\hat{\delta}$, and stage two involves fitting a probit regression including X and the fitted residual, or control function, $\hat{V} = X - \hat{\gamma} - \hat{\delta}Z$. The estimated coefficient on X is then a consistent estimator of β^* . In Babanezhad et al. (2009) this estimator is referred to as the ‘Two-stage IV-estimator I’.

Keeping the distribution of V normal, we can obtain a model that satisfies the assumptions of the logistic SMM by defining the distribution of U to be such that again $U = \rho V + W$, but now W is logistically distributed, independent of Z and V .

We now further explore the relationship between the control function approach and the logistic SMM for binary outcomes and treatments. The structural model in Section

4, satisfying the logistic SMM assumptions, was given by

$$Y(x) = I(\alpha + \xi_0 x + U > 0), \quad X(z) = I(\gamma + \delta z + V > 0),$$

where

$$U = (\beta_0 - \alpha) + \beta_1 X + \beta_2 Z + W,$$

and V and W are marginally standard logistic, mutually independent and independent of Z , with the β parameters restricted to satisfy the CMI assumption. It is not directly clear what this specification implies for the correlation between U and V , but if we specify the residual $R = X - E(X|Z) = X - p_0 - (p_1 - p_0)Z$, with $p_0 = E\{X(0)\}$ and $p_1 = E\{X(1)\}$, we see that

$$U = (\beta_0 - \alpha + \beta_1 p_0) + \beta_1 R + \{\beta_2 + \beta_1(p_1 - p_0)\}Z + W.$$

If we then further restrict $E(U|Z) = 0$, we get that $\beta_2 = -\beta_1(p_1 - p_0)$, $\beta_0 = \alpha - \beta_1 p_0$ and thus

$$U = \beta_1 R + W. \tag{14}$$

It is interesting to note that the CMI assumption can hold when $E(U|Z) \neq 0$ and vice versa, although in both cases Z is not independent of U . If (14) holds, then the simple logit estimator, regressing Y on X and the estimated residual \hat{R} will result in a consistent estimator of ξ_0 . If the parameters are such that both (14) and the CMI assumptions hold, then both this simple logit estimator and the logistic SMM estimator are consistent.

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.

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] Babanezhad, M., Vansteelandt, S. and Goetghebeur, E. (2009), On the performance of instrumental variable estimators of the causal odds ratio, submitted manuscript.
- [5] Balke, A. and Pearl, J. (1997), Bounds on treatment effects from studies with imperfect compliance, *Journal of the American Statistical Association*, 92, 1171-1176.
- [6] Brookhart, M.A. and Schneeweiss, S. (2007), Preference-based instrumental variable methods for the estimation of treatment effects: assessing validity and interpreting results, *The International Journal of Biostatistics* 3, article 14.
- [7] Brookhart, M.A., Wang, P.S., Solomon, D.H. and Schneeweiss, S. (2006), Evaluating short-term drug effects using a physician-specific prescribing preference as an instrumental variable, *Epidemiology* 17, 268-275.
- [8] Chesher, A. (2010), Instrumental variable models for discrete data, *Econometrica* (in press).

- [9] 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.
- [10] Cuzick, J., Sasieni, P., Myles, J. and Tyrer, J. (2007), Estimating the effect of treatment in a proportional hazards model in the presence of non-compliance and contamination, *Journal of the Royal Statistical Society, Series B* 69, 565-588.
- [11] Didelez, V. and Sheehan, N. (2007), Mendelian randomization as an instrumental variable approach to causal inference, *Statistical Methods in Medical Research* 16, 309-330.
- [12] Didelez, V., Meng, S. and Sheehan, N. (2010), On the bias of IV estimators for Mendelian Randomisation, *Statistical Science* (in press).
- [13] 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.
- [14] Frangakis, C.E. and Rubin, D.B. (2002), Principal stratification in causal inference, *Biometrics*, 58, 21-29.
- [15] 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.
- [16] 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.
- [17] Greenland, S. (2000), An introduction to instrumental variables for epidemiologists, *International Journal of Epidemiology* 29, 722-729.

- [18] Hernán, M.A. and Robins, J.M. (2006), Instruments for causal inference: an epidemiologist's dream?, *Epidemiology* 17, 360-372.
- [19] Imbens, G.W. and Angrist, J. (1994), Identification and estimation of local average treatment effects, *Econometrica* 62, 467-476.
- [20] Meng, X.L. (1994), Multiple-imputation inferences with uncongenial sources of input (with discussion), *Statistical Science* 9, 538-573.
- [21] Moodie, E.E.M., Platt, R.W. and Kramer, M.S. (2009), Estimating response-maximized decision rules with applications to breast-feeding, *Journal of the American Statistical Association* 104, 155-165.
- [22] Nagelkerke, N., Fidler, V., Bernsen, R. and Bordgorff, M. (2000), Estimating treatment effects in randomized clinical trials in the presence of non-compliance, *Statistics in Medicine* 19, 1849-1864.
- [23] Pearl, J. (2000), *Causality: Models, Reasoning and Inference*, Cambridge: Cambridge University Press.
- [24] Rivers, D. and Vuong, Q.H. (1988), Limited information estimators and exogeneity tests for simultaneous probit models, *Journal of Econometrics* 39, 347-366.
- [25] 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.
- [26] 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.

- [27] 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.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] Ten Have, T.R., Elliott, M.R., Joffe, M., Zanutto, E. and Datto, C. (2004), Causal models for randomized physician encouragement trials in treating primary care depression, *Journal of the American Statistical Association* 99, 16-25.
- [32] 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.
- [33] Vansteelandt, S. and Goetghebeur, E. (2003), Causal inference with generalized structural mean models, *Journal of the Royal Statistical Society, Series B* 65, 817-835.

- [34] Witteman, J.C.M., D'Agostino, R.B., Stijnen, 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.

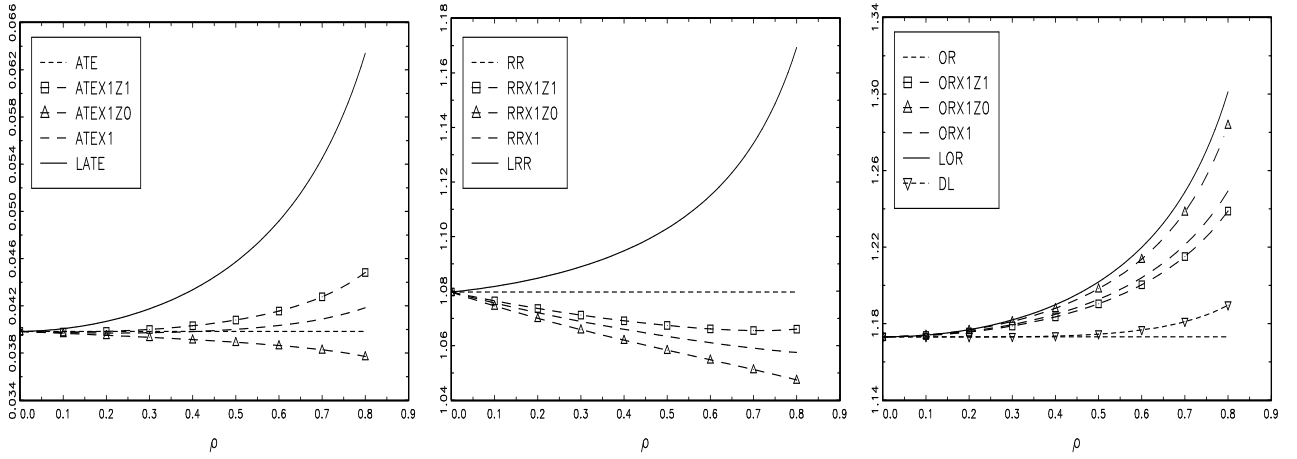


Figure 1. $\gamma = 0, \delta = 0.500$

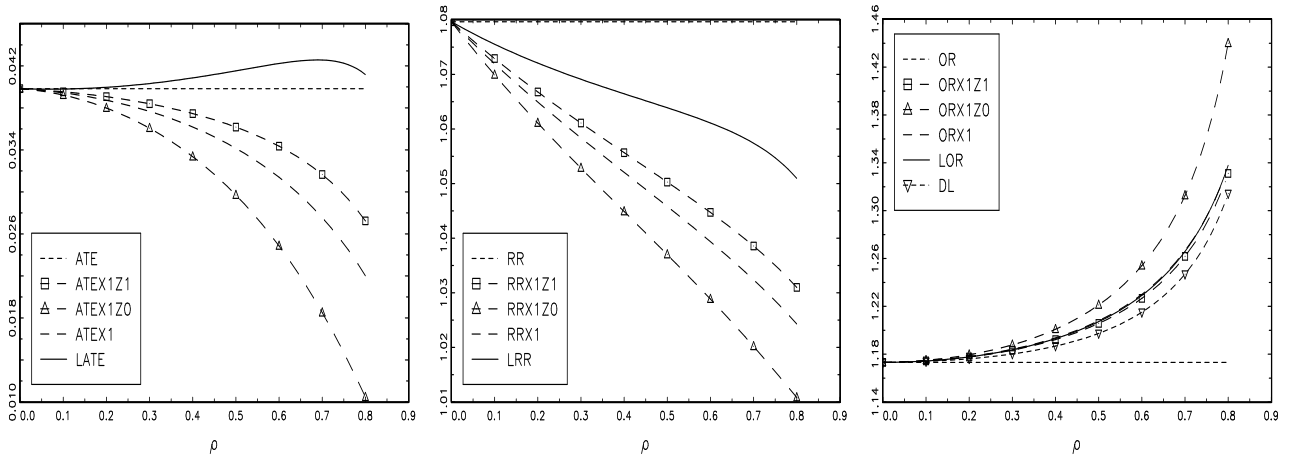


Figure 2. $\gamma = -1, \delta = 0.615$

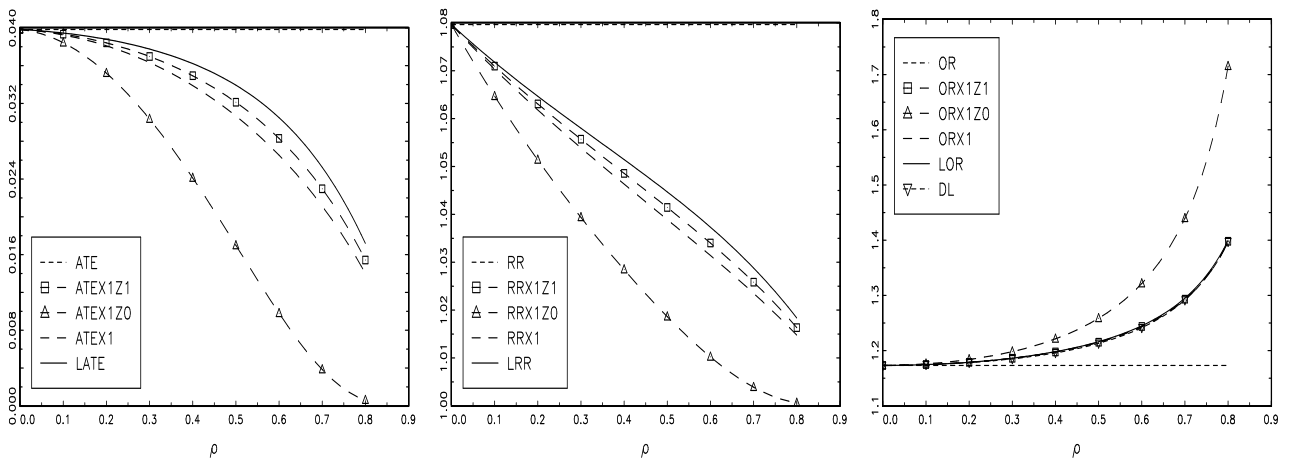


Figure 3. $\gamma = -2, \delta = 1.208$

Table 1. Estimation results

Estimator	Estimate	95% CI
Logistic OR	1.032	[0.80, 1.37]
Add. SMM	-0.0092	[-0.017, -0.002]
Mult. SMM	-0.176	[-1.56, 0.81]
LOR Estimator	-0.174	[-1.56, 0.70]
DL SMM	0.029	[0.01, 0.73]

Notes: 95% CI calculated from 100 bootstrap samples

Table 2. 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	<i>stdev</i>	mean	<i>stdev</i>	mean	<i>stdev</i>
ATE	0.0453		0.0344		0.0344	
ATEX1Z1	0.0455		0.0624		0.0708	
ATEX1Z0	0.0437		0.0658		0.0658	
ATEX1	0.0447		0.0640		0.0684	
LATE	0.0574		0.0402		0.0923	
Add. SMM	0.0573	<i>0.0179</i>	0.0402	<i>0.0181</i>	0.1143	<i>0.0229</i>
RR	1.0907		1.0688		1.0688	
RRX1Z1	1.0682		1.0935		1.1006	
RRX1Z0	1.0627		1.0944		1.0944	
RRX1	1.0656		1.0939		1.0977	
LRR	1.1219		1.0854		1.1376	
Mult. SMM	1.1224	<i>0.0402</i>	1.0861	<i>0.0400</i>	1.1515	<i>0.0332</i>
OR	1.1994		1.1478		1.1478	
ORX1Z1	1.2376		1.3464		1.4454	
ORX1Z0	1.2421		1.3981		1.3981	
ORX1	1.2394		1.3688		1.4225	
LOR	1.2586		1.1748		1.5800	
LOR estimator	1.2614	<i>0.0903</i>	1.1778	<i>0.0852</i>	2.2142	<i>0.3659</i>
DL SMM	1.2199	<i>0.0772</i>	1.1420	<i>0.0685</i>	2.0390	<i>0.2610</i>

Notes: 1000 Monte Carlo replications; sample size 500,000. Population parameters ATE etc. are calculated within the samples and approximated by the averages of the 1000 replications