Harvard University
Harvard University Biostatistics Working Paper Series

Year 2007 *Paper* 70

Correcting Instrumental Variables Estimators for Systematic Measurement Error

Stijn Vansteelandt* Manoochehr Babanezhad†
Els Goetghebeur‡

*Ghent University, stijn.vansteelandt@ugent.be
†Ghent University and Gorgan University
‡egoetghe@hsph.harvard.edu

http://biostats.bepress.com/harvardbiostat/paper70

Correcting Instrumental Variables Estimators for Systematic Measurement Error

Stijn Vansteelandt, Manoochehr Babanezhad and Els Goetghebeur

ABSTRACT. Instrumental variables (IV) estimators are becoming increasingly popular because they allow for estimating the average causal effect of an exposure on an outcome in the presence of unmeasured confounders. Often, however, exposures are hard to measure and may carry errors which not only reflect random noise, but also contain a systematic component. In this article, we study the impact of such error-prone exposure measurements on IV estimators for the average causal effect of exposure on outcome. In addition, we propose a class of IV estimators for this effect under linear structural mean models, which correct for possibly systematic measurement error in the presence of a baseline measurement which is associated with the observed exposure and known not to modify the causal effect of interest. Simulation studies and the analysis of a small blood pressure reduction trial ($n = 105$) with treatment noncompliance confirm the adequate performance of our estimators in finite samples. Our results demonstrate that incorporating a limited amount of prior information about a weakly identified parameter (e.g., the error mean) may yield substantial improvements.

Key words: causal inference; instrumental variables; measurement error; noncompliance; prior information; structural models; weak identifiability.

Stijn Vansteelandt is Assistant Professor, Ghent University, 9000 Ghent, Belgium (email: stijn.vansteelandt@ugent.be); Manoochehr Babanezhad is Doctoral Student, Ghent University, 9000 Ghent, Belgium, and Gorgan University, Gorgan, Iran; and Els Goetghebeur is Associate Professor, Ghent University, 9000 Ghent,

Belgium, and Harvard School of Public Health, Boston MA 02115. The authors acknowledge support from IAP research network grant nr. P06/03 from the Belgian government (Belgian Science Policy). The second author would like to thank Iran's Minister of Science for financially supporting his PhD study at Ghent University.

1 INTRODUCTION

Instrumental variables (IV) methods have a long tradition in economics and econometrics, where they are used in connection with structural equation models. They have more recently entered the medical, epidemiological and biostatistical literature (for reviews, see e.g. Greenland, 2000; Martens et al., 2006). These methods succeed in estimating the average causal effect of an exposure on an outcome, even in the presence of unmeasured confounding, by using a so-called instrumental variable. This is a variable (i) which is associated with the exposure; (ii) has no direct effect on the outcome; and (iii) does not share common causes with the outcome (Hernán and Robins, 2006). Instrumental variables arise naturally in double-blind randomized controlled trials with treatment noncompliance because randomization (i.e. the instrument) is associated with received treatment (i.e. the exposure), does not affect the outcome other than through received treatment and shares no common causes with the outcome by virtue of randomization. They are hence frequently used to adjust for treatment noncompliance in randomized experiments (see e.g. Goetghebeur and Vansteelandt, 2005 for a review) and for the analysis of randomized encouragement designs (Ten Have et al., 2004). At the same time, they are becoming increasingly popular in observational settings where the conditions for an instrumental variable are nevertheless harder to justify. In genetics, for instance, the random assortment of genes transferred from parents to offspring - called 'Mendelian randomization' - resembles the use of randomization

in experiments and is therefore a natural instrumental variable for estimating the effect of genetically affected exposures on a given trait (Sheehan and Didelez, 2005). Casas et al. (2005) use this idea to assess the influence of plasma homocysteine level on the risk of stroke with homozygosity at a specific allele as an instrumental variable. In most observational studies no real or natural randomization is present, in which case the availability of an instrumental variable must be assessed on theoretical grounds. For instance, Leigh and Schembri (2004) use the observed cigarette price per region as an instrumental variable to estimate the causal effect of smoking on health, assuming that the price of cigarettes may only impact health by mediating exposure to cigarette smoke.

With the increasing popularity of IV methods, there is a growing concern as to how these methods would fare under violations of the study design, such as measurement error in the exposure. The latter concern is particularly prevalent in the context of noncompliance adjustment in clinical trials (Dunn, 1997; Goetghebeur and Vansteelandt, 2005) because simple measurements of noncompliance (e.g. the number of pills removed from the pill container) are notorious for overestimating the amount of drug actually taken (Urquhart and De Klerk, 1998).

Random measurement error on the exposure is not alarming for IV estimators in linear (structural mean) models (Robins, 1994; Goetghebeur and Lapp, 1997). Indeed, these estimators continue to be asymptotically unbiased with at most a slight loss of efficiency, when random measurement error is ignored (Goetghebeur and Vansteelandt, 2005). When measurement error is systematic, tests of the causal null hypothesis of no effect remain valid, but effect estimates may become biased. Because systematic error is often a real concern, especially in the noncompliance problem that motivated this research, our goal in this article is to investigate how IV estimators for the parameters in linear (structural

mean) models may be adjusted for systematic measurement error. Goetghebeur and Vansteelandt (2005) show how this can be done when the average size of the error is known (conditional on covariate values). This allows for sensitivity analyses to be performed, but leaves open the question of how to estimate the average size of the measurement error and subsequently correct for it. Because of identifiability problems, the latter can only be realized when extraneous information is available. One common source of information is an instrumental variable for the measurement error (Buzas and Stefanski, 1996; Carroll et al., 2004, 2006). In contrast to the previously defined instrumental variable which we used for confounder adjustment, this is a pre-exposure surrogate for the observed exposure (in the sense that it is correlated with exposure) which is known not to modify the exposure effect of interest. Our interest in such variables stems from the fact that other common sources of information on the measurement error (e.g. repeated measurements or validation samples) are typically not available in the problem setting which motivated this research.

In the next section, using ideas from linear regression models with error in the covariates (Carroll et al., 2006), we show how an instrumental variable for the measurement error may help to correct IV estimators for systematic error under linear structural mean models (Goetghebeur and Lapp, 1997; Robins, 1994). In Section 2.3, we diagnose that the error-adjusted estimator behaves poorly in small to moderate sample sizes as compared to the standard estimator which ignores measurement error. We show in Section 3 that this is due to the average magnitude of the error being weakly identified at causal effects close to zero. In Section 3, we accommodate this by incorporating weak prior information in the form of bounds on the magnitude of the average error. This leads to estimators for the causal effect of observed exposure with good performance in finite samples, as confirmed

through the analysis of a small placebo-controlled hypertension trial in Section 4 and through simulation studies in Section 5. Our results offer more general insight how to incorporate prior information about weakly identified nuisance parameters in a frequentist analysis, in favour of precision for the target parameter.

2 ADJUSTING FOR MEASUREMENT ERROR

2.1 Assumptions

We consider a study which is designed to collect data on a scalar exposure Z_i, a scalar outcome Y_i and possibly on a set of baseline (i.e. pre-exposure) covariates $\mathbf{X}_i$ for independent subjects $i = 1, ..., n$. The goal of the study is to assess the average effect of exposure Z_i on outcome Y_i, which we define as a contrast, i.e.

$$\mathrm{E}(Y_i - Y_{i0}|Z_i, \mathbf{X}_i), \tag{1}$$

of observed outcomes Y_i and potential exposure-free outcomes Y_{i0} (Rubin, 1978). The latter indicates a reference response which would have been measured for subject i if all conditions were the same as in the considered study, but no exposure were received (e.g. if the assigned experimental treatment contained no active dose). Furthermore, suppose that the exposure Z_i is imprecisely measured so that the observed exposure level W_i for subject i may differ from the actual exposure level Z_i, which is unobserved.

Due to the lack of observations on Y_{i0} and Z_i, identification of the causal effect (1) requires assumptions, which may realistically hold, but are partly untestable.

Assumption A1 (IV assumption): measurements are available for each subject i on an instrumental variable R_i, which satisfies the following assumptions:

1. within strata of baseline covariates $\mathbf{X}_i$, $\mathrm{E}(Y_{i0}|\mathbf{X}_i, R_i) = \mathrm{E}(Y_{i0}|\mathbf{X}_i)$.

2. exclusion restriction (Angrist, Imbens and Rubin, 1996): R_i has no direct effect on the outcome (only an indirect effect via the exposure is possible).

In double-blind randomized trials of an asymptomatic disease, one expects these assumptions to hold for randomization R_i since patients and physicians are unaware of the assigned treatment (Robins, 1994).

Assumption A2 (Consistency assumption): to link exposure-free outcomes to the observed data, we assume that $Y_i = Y_{i0}$ for subjects with $Z_i = 0$.

Assumption A3 (Model assumption): the causal effect (1) obeys the linear structural mean model (Robins, 1994)

$$\mathrm{E}(Y_i - Y_{i0}|Z_i, \mathbf{X}_i, R_i) = \gamma(\mathbf{X}_i, R_i; \boldsymbol{\psi}^*)Z_i \tag{2}$$

where $\gamma(\mathbf{X}_i, R_i; \boldsymbol{\psi})$ is a known function smooth in $\boldsymbol{\psi}$, satisfying $\gamma(\mathbf{X}_i, R_i; \mathbf{0}) = 0$, and where $\boldsymbol{\psi}^*$ is an unknown finite-dimensional parameter. For instance, in placebo-controlled randomized experiments with $R_i = 1$ for subjects randomized to the experimental arm and $R_i = 0$ for placebo control, we may choose

$$\mathrm{E}(Y_i - Y_{i0}|Z_i, \mathbf{X}_i, R_i) = \psi Z_i R_i \tag{3}$$

when subjects with $R_i = 0$ are not exposed to the experimental treatment. Here, ψ expresses the expected change in outcome when those exposed to $Z_i = 1$ would have their exposure set to zero. When treatment effects are potentially modified by pre-treatment covariates, one may add covariate-exposure interactions, as in

$$\mathrm{E}(Y_i - Y_{i0}|Z_i, \mathbf{X}_i, R_i) = (\psi_1 + \boldsymbol{\psi}_2'\mathbf{X}_i)Z_i R_i.$$

Here, $\boldsymbol{\psi}_2$ defines the change in the average effect of unit exposure per unit increase in $\mathbf{X}_i$. Note that we will restrict our development to models (2) which postulate the causal effect to be linear in the exposure. This is because linear structural

mean models with nonlinear exposure effects suffer from identification problems, even in the absence of measurement error (Vansteelandt and Goetghebeur, 2005).

Assumption A4 (Measurement error assumptions): Given the difficulty in obtaining information about measurement error characteristics, we will assume the availability of an instrumental variable $\mathbf{T}_i \subseteq \mathbf{X}_i$ for the measurement error for each subject i. This is surrogate for the observed exposure (in the sense that is it is conditionally associated with W_i, given $(\mathbf{S}_i, R_i)$, where $\mathbf{S}_i$ is such that $\mathbf{X}_i \equiv (\mathbf{S}_i, \mathbf{T}_i)$), which is measured prior to exposure and is such that it does not modify the causal effect of received exposure on the outcome, i.e. such that

$$\mathrm{E}(Y_i - Y_{i0}|Z_i, \mathbf{X}_i, R_i) = \mathrm{E}(Y_i - Y_{i0}|Z_i, \mathbf{S}_i, R_i) \tag{4}$$

We thus assume that $\gamma(\mathbf{X}_i, R_i; \boldsymbol{\psi}) = \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})$ in (2) does not involve $\mathbf{T}_i$. For instance, in clinical trials with a run-in period during which all patients receive placebo tablets and compliance is monitored, one possible source of such instrumental variable would be compliance during the run-in period. This is because run-in compliance is likely a good surrogate for the true exposure and because, given the actual compliance during the active study period, run-in compliance may not further relate to the treatment effect. Ten Have et al. (2007) make a similar assumption for disentangling direct from indirect causal effects. Note that $\mathbf{T}_i$ differs from and satisfies different assumptions than the instrumental variable R_i, which satisfies assumption A1.

2.2 Inference

Our goal is to estimate the parameter $\boldsymbol{\psi}^*$ indexing (2) under model $\mathcal{A}$, which is the model for the observed data $(Y_i, W_i, R_i, \mathbf{X}_i)$ defined by assumptions A1-A4

and with the conditional density

$$f(R_i|\mathbf{X}_i) \text{ known.} \tag{5}$$

It follows from Proposition 1 below that the average measurement error $\delta(\mathbf{X}_i, R_i) \equiv E(W_i - Z_i|\mathbf{X}_i, R_i)$ is all that must be known for identifying $\boldsymbol{\psi}^*$.

Proposition 1. Model $\mathcal{A}$ is the same model for the observed data as the conditional mean independence model $\mathcal{B}$ defined by (5) and

$$\begin{aligned} &E\left[Y_i - \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi}^*)\left\{W_i - \delta(\mathbf{X}_i, R_i)\right\}|\mathbf{X}_i, R_i\right] \\ &\quad = E\left[Y_i - \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi}^*)\left\{W_i - \delta(\mathbf{X}_i, R_i)\right\}|\mathbf{X}_i\right]. \end{aligned} \tag{6}$$

For convenience, we will assume that $\delta(\mathbf{X}_i, R_i) = \delta^*$ is constant, although this assumption will be straightforward to relax (see also the discussion). Our goal is thus to estimate $\boldsymbol{\psi}^*$ in model $\mathcal{A}$ when the average size δ^* of the error is unknown.

Note that the restrictions which model $\mathcal{A}$ imposes on the error distribution are very weak. First, it allows the error to be associated with both the true exposure Z_i and observed exposure W_i. As such, the error model encompasses both the classical and Berkson error model (Carroll et al., 2006). Second, by avoiding assumptions about the conditional association between W_i and Y_i, given Z_i, it allows the error to be differential (i.e. associated with outcome conditional on the exposure) (see the proof of Proposition 1 for a more formal argument). This is important because, for instance in a hypertension trial, patients may be more reluctant to 'confess' to noncompliance when their outcome (e.g. blood pressure) stayed below target (e.g. remained high). Finally, model $\mathcal{A}$ makes no assumptions on the measurement error distribution (other than restriction $\delta(\mathbf{X}_i, R_i) = \delta^*$, which is easy to relax). This is useful because the error distribution can be very complex. For instance, exposures can be very small in practice, in which case the negative errors become constrained by the fact that negative exposures (i.e. doses) are never reported.

By Proposition 1 and the fact that $\boldsymbol{\psi}^*$ is the same functional of the observed data under models $\mathcal{A}$ and $\mathcal{B}$, inference for $\boldsymbol{\psi}^*$ is the same under both models. It follows that the set of all consistent and asymptotically normal (CAN) estimators for $\boldsymbol{\psi}^*$ is the same under models $\mathcal{A}$ and $\mathcal{B}$, where the latter can be obtained as in Robins (1994) by solving the mean independence estimating equations

$$\sum_{i=1}^{n} \mathbf{d}(R_i, \mathbf{X}_i) \left[Y_i - \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})(W_i - \delta) - q(\mathbf{X}_i) \right] = 0 \tag{7}$$

jointly for $\boldsymbol{\theta} = (\boldsymbol{\psi}', \delta)'$, where $\mathbf{d}(R_i, \mathbf{X}_i) = \mathbf{g}(R_i, \mathbf{X}_i) - \mathrm{E}\left\{\mathbf{g}(R_i, \mathbf{X}_i)|\mathbf{X}_i\right\}$ and $\mathbf{g}(R_i, \mathbf{X}_i)$ and $q(\mathbf{X}_i)$ are arbitrary index functions of the dimension of $\boldsymbol{\theta}$ which can be chosen in view of efficiency. In particular, the efficient score for $\boldsymbol{\psi}^*$ under model $\mathcal{A}$ is the same as the efficient score for $\boldsymbol{\psi}^*$ under model $\mathcal{B}$. When the conditional variance of $Y_i - \gamma(\mathbf{X}_i, R_i; \boldsymbol{\psi})(W_i - \delta)$, given $(R_i, \mathbf{X}_i)$, is constant, the latter is obtained by setting $q(\mathbf{X}_i)$ equal to

$$q_{opt}(\mathbf{X}_i) = E\left\{Y_i - \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})(W_i - \delta)|\mathbf{X}_i, R_i\right\}$$

and $\mathbf{d}(R_i, \mathbf{X}_i)$ equal to $\mathbf{d}_{opt}(R_i, \mathbf{X}_i) = \mathbf{g}_{opt}(R_i, \mathbf{X}_i) - \mathrm{E}\left\{\mathbf{g}_{opt}(R_i, \mathbf{X}_i)|\mathbf{X}_i\right\}$ with

$$\mathbf{g}_{opt}(R_i, \mathbf{X}_i) = E\left\{ \frac{\partial \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})(W_i - \delta)}{\partial \boldsymbol{\theta}} |\mathbf{X}_i, R_i \right\}$$

Theorem 1.

1. Under regularity conditions, the solution $\hat{\boldsymbol{\psi}}(\mathbf{d}, q)$ to (7) satisfies $\sqrt{n}\left(\hat{\boldsymbol{\psi}}(\mathbf{d}, q) - \boldsymbol{\psi}^*\right) \rightarrow N\left(0, \Gamma(\mathbf{d}, q)\right)$ in distribution, where

$$\Gamma(\mathbf{d}, q) = E^{-1}\left\{ \frac{\partial \mathbf{U}_i(\mathbf{d}, q; \boldsymbol{\psi}^*)}{\partial \boldsymbol{\psi}} \right\} \mathrm{Var}\{\mathbf{U}_i(\mathbf{d}, q; \boldsymbol{\psi}^*)\} E^{-1\prime}\left\{ \frac{\partial \mathbf{U}_i(\mathbf{d}, q; \boldsymbol{\psi}^*)}{\partial \boldsymbol{\psi}} \right\} \tag{8}$$

with $\mathbf{d}(R_i, \mathbf{X}_i) = (\mathbf{d}_\psi(R_i, \mathbf{X}_i), d_\delta(R_i, \mathbf{X}_i))$ and

$$\begin{aligned} \mathbf{U}_i(\mathbf{d}, q; \boldsymbol{\psi}) = & \left[\mathbf{d}_\psi(R_i, \mathbf{X}_i) - \frac{E\left\{\mathbf{d}_\psi(R_i, \mathbf{X}_i)\gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})\right\}}{E\left\{d_\delta(R_i, \mathbf{X}_i)\gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})\right\}} d_\delta(R_i, \mathbf{X}_i) \right] \\ & \times \left[Y_i - \gamma(\mathbf{S}_i, R_i; \boldsymbol{\psi})(W_i - \delta) - q(\mathbf{X}_i) \right] \end{aligned}$$

2. The average error δ^* is not root-n estimable at $\boldsymbol{\psi}^* = 0$.

3. For arbitrary $(\mathbf{d}, q)$, $\Gamma(\mathbf{d}_{opt}, q_{opt}) \leq \Gamma(\mathbf{d}, q)$ where $A \leq B$ is defined as $A - B$ being semi-positive definite.

Part 1 of Theorem 1 confirms that the solution $\hat{\boldsymbol{\psi}}(\mathbf{d}, q)$ to (7) is a root-n CAN estimator of $\boldsymbol{\psi}^*$. This is even so at $\boldsymbol{\psi}^* = 0$ where δ^* is not root-n estimable by the fact that the expected derivative of the estimating function w.r.t. δ is zero at $\boldsymbol{\psi}^* = 0$. Theorem 1 also shows how to calculate the efficient score $\mathbf{U}_i(\mathbf{d}_{opt}, q_{opt}; \boldsymbol{\psi})$ for $\boldsymbol{\psi}^*$ in model $\mathcal{A}$. For example, with binary R_i, $\mathbf{X}_i = \mathbf{T}_i$, $\gamma(\mathbf{S}_i, R_i; \psi) = \psi R_i$ and assuming homoscedasticity and constant randomization probabilities $\pi = P(R_i = 1)$, the semi-parametric efficient score for ψ^* is

$$(R_i - \pi) \left[E(W_i | R_i = 1, \mathbf{X}_i) - E\left\{E(W_i | R_i = 1, \mathbf{X}_i)\right\}\right] \left\{Y_i - \psi(W_i - \delta)R_i - q_{opt}(\mathbf{X}_i)\right\}$$

This score differs from the efficient score in the absence of biased measurement error (i.e. assuming that $\delta^* = 0$) in that it carries the additional term $E\left\{E(W_i | R_i = 1, \mathbf{X}_i)\right\}$, which corrects for estimation of the error mean. This term reduces the variance of the estimating functions and, as such, encodes efficiency loss. Specifically, note that the efficient score becomes 0 when the instrument T is uncorrelated with the observed exposure, and hence that ψ^* is not root-n estimable in that case. By the same token, weak instruments for the measurement error (i.e. instruments which are weakly correlated with observed exposure) yield unstable effect estimates.

2.3 Bias-variance Trade-off

The anticipated loss of efficiency of the error-adjusted estimator raises the question whether the bias correction of the previous section is meaningful. To this end, we investigate the bias-variance trade-off of using the error-adjusted instead of the standard unadjusted estimator for the causal effect ψ^*. To obtain tractable

expressions for the mean-squared error of both estimators, we assume that $Z \sim N(\mu_z, \sigma_z^2), T|Z \sim N(\nu_0 + \nu_1 Z, \sigma_{t|z}^2), Y_0|Z,T \sim N(\alpha_0 + \alpha_1 Z + \alpha_2 T, \sigma_0^2)$ and that $Y = Y_0 + (\psi + \epsilon)RZ$ with $\epsilon|Y_0, Z, T \sim N(0, \sigma^2)$.

Under the working assumption of no systematic measurement error (i.e. fixing $\delta^* = 0$ in equation (7) and not estimating it), the efficient score for ψ^* is $U_u(\psi) = (0.5 - R)\mathrm{E}(W|T, R = 1)\{Y - \psi RW - E(Y|R = 0, T)\}$ in model $\mathcal{A}$ with $\mathbf{X}_i = \mathbf{T}_i$ under the above data-generating mechanism. It follows after some algebra that the solution $\hat{\psi}_u$ to $\sum_{i=1}^n U_{ui}(\psi) = 0$ has bias which can be approximated by

$$\mathrm{E}^{-1}\left(\frac{\partial U_u(\psi)}{\partial \psi}\right)\mathrm{E}\{U_u(\psi)\} = \frac{\psi\delta(\mu_z + \delta)}{\sigma_z^2 - \sigma_{z|t}^2 + (\mu_z + \delta)^2}$$

where $\sigma_{z|t}^2 = \sigma_z^2\sigma_{t|z}^2/(\nu_1^2\sigma_z^2 + \sigma_{t|z}^2)$ is the conditional variance of Z given T, and asymptotic variance given by

$$\frac{1}{n}\left[\frac{4\sigma_0^2 + 4\alpha_1^2\sigma_{z|t}^2 + 2\psi^2\sigma_u^2 + \psi^2\delta^2}{\sigma_z^2 - \sigma_{z|t}^2 + (\mu_z + \delta)^2} + \frac{\psi^2\delta^2(\sigma_z^2 - \sigma_{z|t}^2)}{\left\{\sigma_z^2 - \sigma_{z|t}^2 + (\mu_z + \delta)^2\right\}^2}\right]$$

Allowing for systematic measurement error, the efficient estimator $\hat{\psi}_c$ for ψ^* under model $\mathcal{A}$ has no asymptotic bias and asymptotic variance which equals

$$\frac{1}{n}\frac{\sigma_0^2 + \alpha_1^2\sigma_{z|t}^2 + 0.5\psi^2\sigma_u^2}{0.5^2(\sigma_z^2 - \sigma_{z|t}^2)}$$

Note that the bias and asymptotic variance of the estimators is inversely proportional to the multiple correlation coefficient for the regression of Z on T, but becomes infinite for the error-adjusted estimator when Z and T are uncorrelated.

Figure 1 shows the range of values δ for the average error under which the standard estimator (which ignores measurement error) has smaller mean squared error than the error-adjusted estimator, in function of the sample size and the multiple correlation coefficient for the linear regression of Z on T. Specifically,

the values of δ comprised between the solid lines indicate data-generating mechanisms under which the standard estimator outperforms the error-adjusted estimator in terms of mean squared error. The figure was constructed using the values $\mu_z = 0.85, \sigma_z^2 = 0.11, \nu_0 = 0.75, \nu_1 = 0.12, \sigma_{t|z}^2 = 0.012, \alpha_0 = -4.4, \alpha_1 = 6.8, \alpha_2 = -13.7, \sigma_0^2 = 53.2, \sigma_u^2 = 0, \psi = -7.5$ and $\sigma^2 = 0$ which are reflective of the hypertension study that we will analyze in Section 4. The figure shows that at small sample sizes ($n = 105$), correction for systematic measurement error may lead to smaller mean squared error, but only when the systematic error component is substantial (i.e. of about the size of the average exposure μ_z) and, at the same time, the instrument T is strongly correlated with Z. Further note that bias correction using the error-adjusted estimator may be practical at moderate degrees of error and moderate correlations between T and Z, but only at very large sample sizes.

Figure 1 about here

3 INCORPORATING PRIOR INFORMATION

The previous results demonstrate the poor performance of the error-adjusted estimator, even in settings where the sample size is moderate and good (pre-exposure) predictors of the exposure are available. In particular, tests of the causal null hypothesis using this approach may be much less powerful than the standard test of the causal null (i.e. that R and Y are independent), which is immune to measurement error on the exposure. This is surprising, considering that the score test of $\boldsymbol{\psi}^* = 0$ under model $\mathcal{A}$ does not involve δ^* and hence that one need not correct for measurement error when testing the causal null hypothesis. Curiously, it follows that one can validly and efficiently test the causal null hypothesis without needing to correct for measurement error, but that a score test of $\boldsymbol{\psi}^* = \boldsymbol{\psi}_0$

with $\boldsymbol{\psi}_0$ arbitrarily close to (but different from) 0, would require correcting for measurement error and hence could imply a serious and sudden loss of power.

The root cause of this apparent discontinuity is the fact that, as shown in Part 2 of Theorem 1, δ^* is not root-n estimable at $\boldsymbol{\psi}^* = \mathbf{0}$ so that estimation of δ^* affects the distribution of the score test statistic, even though it gets multiplied by $\boldsymbol{\psi}^* = \mathbf{0}$ in the test statistic (i.e. even at the causal null hypothesis). In particular, it follows from the proof of Theorem 1 that $\sqrt{n}\left\{\hat{\delta}(\mathbf{d}, q) - \delta\right\}\boldsymbol{\psi}$, with $\hat{\delta}(\mathbf{d}, q)$ the solution for δ to (7), is bounded in probability with strictly positive variance for each value of $\boldsymbol{\psi}$, suggesting that $\hat{\delta}\boldsymbol{\psi}$ fluctuates around 0, even when $\boldsymbol{\psi} = \mathbf{0}$, with decreasing variance as the sample size increases.

Similar problems of inestimability at a local point in the parameter space have been noted in other measurement error problems (Gustafson, 2005). More general problems of inferring a parameter $\boldsymbol{\psi}^*$ when a nuisance parameter δ^* is only present under the alternative ($\boldsymbol{\psi}^* \neq \mathbf{0}$) have received some attention, mainly in the econometrics literature (Davies, 1977, 1987; Hansen, 1992; Andrews and Ploberger, 1994). To the best of our knowledge, attention has only been given to testing problems in which the test statistic involves a nuisance parameter which is unidentified at the null. Some of these approaches assume that the nuisance parameter lies within a known open set and base inference on the supremum of a score or likelihood ratio test statistic, taken over all values of the nuisance parameters in the chosen set (Davies, 1977, 1987). Andrews and Ploberger (1994) postulate a prior distribution for the nuisance parameter and base inference on the average of a score or likelihood ratio test statistic over the chosen prior distribution. Our problem is different (a) in that our main focus is on estimation rather than testing; and (b) that a score test for the causal null hypothesis does not involve the nuisance parameter. Nonetheless, inspired by the work of Davies (1977, 1987) and by

sensitivity analyses for IV-estimators with measurement error (Goetghebeur and Vansteelandt, 2005), we will consider estimation under the assumption that the average error δ^* lies within a known open set Δ. This strategy is motivated by the fact that (a) subject-matter experts often have a rough idea about the degree of mismeasurement (Gustafson, 2005); (b) this approach forces the estimate for δ^* to have bounded variation around the truth, contrary to what happens under the previous approach of Section 2.2; and (c) a little of prior information can often be a very good thing (Gustafson, 2005). Furthermore, even when the set Δ is chosen excessively wide, this approach will improve the performance of error-adjusted estimators for $\boldsymbol{\psi}^*$ dramatically by reducing variation in the estimates for the error bias, especially when $\boldsymbol{\psi}^*$ is close to zero (relative to the sample size).

3.1 Improved Error Adjustment

A first approach that we will consider under the assumption that $\delta^* \in \Delta =]\Delta_l, \Delta_u[$ is to solve equations (7) with δ replaced by $\{I(\lambda < 0)\Delta_l + I(\lambda > 0)\Delta_u\}\lambda/(1+|\lambda|)$ and λ unknown. This guarantees estimates for δ^* within the set Δ and will thus greatly improve the stability of estimators for the causal effect $\boldsymbol{\psi}^*$. A drawback which will become apparent in the simulation study of Section 5, is that tests of the causal null hypothesis may still loose substantial power under this approach due to the fact that also λ is not root-n estimable at $\boldsymbol{\psi}^* = \mathbf{0}$. To accommodate this, we will develop a second approach in this section, which we will recommend for data analysis. Specifically, we propose to trade bias for precision by solving a weighted average of the estimating functions for the standard SMM estimator and for the error-adjusted estimator of Section 2.3. Here, we choose to weight the estimating functions for the standard estimator proportional to the probability that the corresponding estimate for δ^* falls outside the chosen set Δ. The philosophy

behind this choice is that estimates for δ^* will not likely fall within the set Δ in situations where little information on the error mean is available. Hence more weight will be given to the standard unadjusted estimator in those cases.

For pedagogic purposes, we will explain our proposal for the case $\gamma(\mathbf{X}_i, R_i; \psi) = \psi R_i$. For notational convenience, we delete reference to the index functions (d, q) in the estimators. For each value ψ in a chosen grid, we calculate an estimator $\hat{\delta}(\psi)$ for δ^* which solves (7) for the given ψ with $d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)$ in place of $d(R_i, \mathbf{T}_i, \mathbf{X}_i)$. Next, we consider a weighted average of the estimating function $U_{\psi i}(\psi, \delta)$ for ψ^* (as defined in (7) with $d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)$ in place of $d(R_i, \mathbf{T}_i, \mathbf{X}_i)$), evaluated at the profile estimator $\delta = \hat{\delta}(\psi)$ and at $\delta = 0$, respectively:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \tilde{U}_i(\psi) \equiv \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \hat{P}\{\hat{\delta}(\psi) \in \Delta\} U_{\psi i}\{\psi, \hat{\delta}(\psi)\} + \hat{P}\{\hat{\delta}(\psi) \notin \Delta\} U_{\psi i}(\psi, 0)$$

In this expression, the weights involve the estimated probability $\hat{P}\{\hat{\delta}(\psi) \notin \Delta\}$ that $\hat{\delta}(\psi)$ falls outside the chosen interval $\Delta =]\Delta_l, \Delta_u[$. Using a similar development as in the proof of Theorem 1, this probability can be approximated by

$$P\{\hat{\delta}(\psi) \notin \Delta\} = 1 + \Phi\left(\frac{\Delta_l - \delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right) - \Phi\left(\frac{\Delta_u - \delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right) \tag{10}$$

with δ replaced by $\hat{\delta}(\psi)$ and $\sigma(\psi)$ replaced by a consistent estimator for the standard deviation of the scaled estimating function $E^{-1}\left[d_\delta(R, T, X)R\right] U_{i\delta}(\psi, \delta)$ for δ^*. We define the improved error-adjusted estimator $\tilde{\psi}$ for ψ^* as the value of ψ at which the score test (9) becomes zero. Curiously, this estimator assigns much weight to the standard estimating equations (which do not adjust for measurement error) when the error mean is estimated to be large. This is (a) because the philosophy behind the estimator is that large values for the error mean are indicative of imprecision; and (b) because the estimating functions are designed to equal the unadjusted estimating functions at the causal null hypothesis (see further).

Theorem 2. Under regularity conditions and for any fixed ψ, $\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\tilde{U}_i(\psi)\rightarrow N(0,\Sigma(\psi))$ in distribution, where $\Sigma(\psi)$ is the variance of

$$P\{\hat{\delta}(\psi)\in\Delta\}U_{i\psi}(\psi,\delta)+P\{\hat{\delta}(\psi)\notin\Delta\}U_{i\psi}(\psi,0)-\Big[P\{\hat{\delta}(\psi)\in\Delta\}$$
$$+\left\{\varphi\left(\frac{\Delta_l-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)-\varphi\left(\frac{\Delta_u-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)\right\}\frac{\sqrt{n}|\psi|\delta}{\sigma(\psi)}\Big]\frac{E\{d_\psi(R,T)R\}}{E\{d_\delta(R,T)R\}}U_{i\delta}(\psi,\delta)$$

Theorem 2 shows that for any fixed ψ, the score test (9) converges to a normal mean zero distribution. This can be used to construct $(1-\alpha)100\%$ confidence intervals for ψ^* as the range of values ψ_0 for ψ such that the two-sided score test based on (9) does not reject the null hypothesis $H_0:\psi^*=\psi_0$ at the $\alpha100\%$ significance level. To evaluate this score test, one may replace the variance of the score test statistic by the sample variance with $P\{\hat{\delta}(\psi)\in\Delta\}$ replaced by $\hat{P}\{\hat{\delta}(\psi)\in\Delta\},\delta$ by $\hat{\delta}(\psi)$ and $\sigma(\psi)$ by $\hat{\sigma}(\psi)$. The resulting confidence intervals have the desirable feature that, asymptotically, they exclude 0 if and only if the standard test of the causal null hypothesis (i.e., that $Y\perp\!\!\!\perp R$) rejects. Indeed, at the null hypothesis $\hat{P}\{\hat{\delta}(0)\notin\Delta\}\xrightarrow{p}1$ and hence the score test statistic becomes

$$\frac{1}{\sqrt{n}}\sum_{i=1}^{n}U_{\psi i}(\psi,0) \;=\; \frac{1}{\sqrt{n}}\sum_{i=1}^{n}d_\psi(R_i,\mathbf{T}_i,\mathbf{X}_i)\{Y_i-q(\mathbf{T}_i,\mathbf{X}_i)\}+o_p(1)$$

for an arbitrary mean zero function $d_\psi(R_i,\mathbf{T}_i,\mathbf{X}_i)$ conditional on $(\mathbf{T}_i,\mathbf{X}_i)$, which is asymptotically equivalent to a score test of the causal null hypothesis under the observed data model defined by restriction (5).

Unfortunately, the suggested confidence intervals are no uniform asymptotic confidence intervals. The reason is that, at each sample size, there exists a ψ^* depending on n which is sufficiently close to zero that the score test statistic (9) is significantly biased as a result of bias in the estimating functions of the standard unadjusted SMM estimator. Specifically, it follows from the proof of Theorem 2 that the improved error-adjusted estimator $\tilde{\psi}$ is asymptotically biased within root-n shrinking neighbourhoods of zero (i.e. when $\psi^*=k/\sqrt{n}$ for some constant k)

and may not converge to a normal distribution along such sequences. Curiously, $\tilde{\psi}$ is asymptotically unbiased and normally distributed along faster converging sequences (i.e. when $\psi^* = kn^{-a}$ for some constant k and $a > 1/2$) and in particular at $\psi^* = 0$. The reason is that, although the probability that $\hat{\delta}(\psi) \in \Delta$ now converges to 0 and hence $\tilde{\psi}$ is asymptotically equivalent to the standard unadjusted SMM estimator, ψ^* is sufficiently close to zero to make any bias in the estimator negligible. Likewise, $\tilde{\psi}$ is asymptotically unbiased and normally distributed along slower converging sequences (i.e. when $\psi^* = kn^{-a}$ for some constant k and $0 \leq a < 1/2$). The reason is that the probability of $\hat{\delta}(\psi) \in \Delta$ now converges to 1 so that the improved error-adjusted estimator is asymptotically equivalent to the error-adjusted estimator of Section 2.2, which is asymptotically unbiased.

The practical implication of the foregoing discussion is that the improved error-adjusted estimator $\tilde{\psi}$ and confidence intervals have no guaranteed performance in finite samples in the sense that, for each sample size, one can find a causal effect ψ^* which is close, but not too close to zero so that $\tilde{\psi}$ is significantly biased and that confidence intervals for ψ^* do not cover ψ^* at the nominal level. This local bias is the price we pay for estimators with smaller variability and limited loss of power for testing the causal null hypothesis. Because this problem only appears within 1 over root-n distances from zero and not within larger or shorter distances, we expect adequate performance in many practical situations. However, in view of this, we develop uniform asymptotic confidence intervals in the next section.

3.2 Uniform Asymptotic Confidence Intervals

Uniform asymptotic $(1 - \alpha)100\%$ confidence intervals are expected to have better finite-sample properties than the intervals of the previous section because they guarantee the existence of a minimal sample size such that, at larger sample

sizes, they cover ψ^* with at least $(1-\alpha)100\%$ chance regardless of the value of ψ^*. Following ideas in Robins (2005), we construct such intervals by first constructing, for each ψ, an asymptotic uniform $(1-\epsilon)100\%$ confidence interval $C(\psi)$ for δ^*, where the choice of $\epsilon < \alpha$ will be discussed later. Because we assume the parameter space for δ^* to be Δ, a conservative asymptotic interval $C(\psi)$ may be obtained as

$$\left\{\hat{\delta}(\psi) \pm z_{\epsilon/2}\frac{\hat{\sigma}(\psi)}{|\psi|\sqrt{n}}\right\} \cap \Delta$$

where $\hat{\sigma}(\psi)$ is a consistent estimator for $\sigma(\psi)$. It follows from Theorem 5.1 in Robins (2005) that an asymptotic uniform $(1-\alpha)100\%$ confidence interval for ψ^* may now be obtained as the set of ψ-values for which

$$\inf_{\delta \in C(\psi)} |\text{Var}^{-1/2}\{U_{\psi i}(\psi,\delta)\}\frac{1}{\sqrt{n}}\sum_{i=1}^{n} U_{\psi i}(\psi,\delta)| < z_{(\alpha-\epsilon)/2}$$

The optimal choice of ϵ that leads to confidence intervals of minimum length is difficult to determine (Robins, 2005). In this article, we propose to choose ϵ in function of ψ as $0.5\alpha|\psi|/(1+|\psi|)$. This choice guarantees that $C(\psi)$ will equal Δ for $\psi^* = 0$ and equal a $(1-\alpha/2)100\%$ confidence interval for δ^* at causal effects ψ^* far from 0. The philosophy behind this choice is that estimates for δ^* will be highly imprecise at causal effects close to zero and hence, given that the parameter space for δ^* is bounded, we expect no difference between 100% confidence intervals and $(1-\alpha)100\%$ confidence intervals for δ^* at $\psi^* = 0$. As such, we need not offer the significance level for ψ^* at small causal effects and will thus get narrower intervals in return. Specifically, the proposed confidence intervals have the feature that they involve no correction for measurement error at $\psi^* = 0$, which is desirable because there is no bias due to measurement error at $\psi^* = 0$.

4 DATA ANALYSIS

We analyze data from a placebo-controlled randomized hypertension trial which enrolled some 300 hypertensive patients (Goetghebeur and Lapp, 1997). After a run-in period of 4 weeks where all patients received placebo tablets, they were randomized to 4 weeks of one of two active treatments (A or B) or placebo. All treatments were prescribed at one tablet per day. Here, we analyze the subset of 105 patients randomized to A or placebo, for whom treatment compliance was electronically measured, ignoring 5 patients who had missing diastolic blood pressure or pill counts.

An intent-to-treat analysis reveals an average difference in blood pressure reduction of 7.5 mmHg (95% CI 4.0; 11.0) without adjustment. This reveals the effect of assignment to treatment A (instead of placebo) on expected diastolic blood pressure reduction from baseline (i.e. the time of randomization). Primary interest lies however in the effect of *received* treatment on average blood pressure reduction. We will therefore fit model (3) with Y_i the blood pressure reduction over the active study period, Z_i the average number of prescribed pills taken, and $\mathbf{X}_i$ the age of patient i. Assuming that compliance measurements are free of systematic error, we estimate that the average blood pressure reduction would have been 9.6 mmHg (95% CI 3.5; 11.8) smaller over the study period among those who choose to take on average one pill per day, had they not taken the exposure.

Because this study was not designed to correct for measurement error, no natural instrumental variables for the measurement error have been recorded. Our analysis is hence for illustrative purposes only and will use age as an instrumental variable in the measurement error analysis. Age was chosen because effect modification through age is not anticipated (nor observed) in this study population, which consists of middle aged hypertensive patients (5th, 95th percentiles: 41 and

69 years). A more adequate analysis would use placebo compliance during the run-in period (which was not recorded here) as an instrumental variable. Using the error-adjusted estimator of Section 2.2, we estimate a larger treatment effect of 27.0 mmHg (95% CI -91.2; 145.2). To improve this imprecise result, we impose the weak assumption that the average error is smaller than 0.25. We believe this assumption to be reasonable, given that the observed percentage of assigned dose taken (i.e. the observed exposure) is 0.85 (i.e., 85%) on average. Choosing $\Delta = [-0.25, 0.25]$ thus allows for 30% of the observed average exposure to be due to systematic error. Using the improved error-adjusted estimator for inference, we estimate a slightly smaller effect of 9.0 mmHg (95% CI 4.4; 17.4) as compared to the standard analysis. As predicted by the theory, the estimate is less precise than the unadjusted estimator, but still significantly different from 0 at the 5% significance level. The uniform asymptotic 95% confidence interval (2.7; 16.8) has a more guaranteed performance in finite samples. To investigate the sensitivity of our result to the choice of Δ, Figure 2 shows the improved error-adjusted estimate, along with uniform 95% confidence intervals in function of the maximum error mean Δ_u, with $\Delta = [-\Delta_u, \Delta_u]$, and reveals reasonable stability. Comparison with the sensitivity analysis results of Goetghebeur and Vansteelandt (2005) shows that the error-adjustment described in this article reduces uncertainty.

Figure 2 about here

5 SIMULATION STUDY

To investigate the behaviour of the error-adjusted estimators in finite samples with ψ^* possibly close to zero, we conducted simulation experiments. Each experiment was based on 5000 replications of random samples of size 105 (i.e. the

sample size of the blood pressure study) or 1000, generated as follows. In each experiment, the instrument T for the measurement error was normal with mean 0.83 and standard deviation 0.14 and R was independently generated from a Bernoulli distribution with success probability 0.5. The true exposure Z and exposure-free response were generated as $Z = T + 0.32\epsilon_Z$ and $Y_0 = -4.4 + 6.8Z - 7.3T + 7.3\epsilon_0$ for independent standard normal variates ϵ_Z, ϵ_0. Finally, we generated Y as $Y_0 + \psi RZ$ and the observed exposure W as $W = Z + U$ where $U \sim N(\delta, 0.01)$.

Table 1 about here

Table 1 summarizes the results for estimation of ψ using i) the standard IV estimator which ignores systematic measurement error (STD); (ii) the error-adjusted estimator of Section 3.1 (IV1); iii) the error-adjusted estimator of Section 3.3 which guarantees estimates for δ to stay within $\Delta = [\Delta_l, \Delta_u]$ with $\Delta_u = -\Delta_l$ equal to 0.5, 0.25 or 0.05, by defining $\delta = \{I(\lambda < 0)\Delta_l + I(\lambda > 0)\Delta_u\}\lambda/(1 + |\lambda|)$ for unknown λ (IV2); the improved error-adjusted estimator of Section 3.3 with the same choices for Δ (IV3). In addition, the table shows uniform asymptotic 95% confidence intervals (UI) corresponding to these choices. The results for the different estimators are as predicted by the theory. The error-adjusted estimator (IV1) is extremely variable at small sample sizes, but performs adequately at larger sample sizes, even at $\psi = 0$. Estimator (IV2) is less variable, although still substantially less precise than the standard unadjusted estimator. Figures 3 and 4 show that estimator (IV1) is normally distributed in moderate sample sizes, even at $\psi = 0$, but not in small samples. It also shows that the improved error-adjusted estimator (IV3) is much less variable than the error-adjusted estimator (IV1). While the former follows a normal distribution in small samples, deviations from normality appear in larger sample sizes as a result of convergence to a normal distribution not being uniform in ψ. By the same token, the improved error-adjusted estimator is more

biased than the error-adjusted estimator in larger samples, and even than the standard IV estimator in some scenarios. Informally, this happens because data sets which carry evidence for causal effects close to zero, yield estimated probabilities of $\hat{\delta}(\psi) \in \Delta$ close to zero. The bias then arises because the small estimated causal effects in such data sets will be more attracted towards the estimates obtained from a standard structural mean analysis (which ignores measurement error) than large estimated causal effects. Additional simulations (not displayed) have shown that, as predicted by the theory, this bias and deviation from normality disappears again in larger sample sizes. Furthermore, note that the confidence intervals for the improved error-adjusted estimator retain their coverage despite these deviations, although there is a tendency for the approach to be conservative. Finally, as predicted by the theory, the uniform confidence intervals are conservative and also wider on average than those obtained via the improved error-adjusted estimator.

The impact of narrower intervals $\Delta = [-0.25, 0.25]$ was large at small sample sizes, but moderate at large sample sizes. For instance, confidence intervals based on the improved error-adjusted estimator had an average length of 8.42 (instead of 13.3) and coverage of 97.0% (instead of 97.7%) in small samples and 4.35 (instead of 4.83) and 98.0% (instead of 97.8%), respectively, in large samples. The impact of $\Delta = [-0.05, 0.05]$ not including the error mean was to induce bias of the order of magnitude of the standard unadjusted estimator. The 95% confidence intervals based on the improved error-adjusted estimator and uniform 95% confidence intervals then no longer cover at the nominal rate. Coverage of those intervals was still better than the coverage of 95% confidence intervals based on the standard unadjusted estimator, but at the expense of being wider.

Figures 3 and 4 about here

6 CONCLUSIONS

We have proposed a general procedure to correct IV estimators for systematic error in the exposure when an instrumental variable for the measurement error is available. This procedure complements the sensitivity analysis approach of Goetghebeur and Vansteelandt (2005) and is especially attractive when the instrumental variables assumption (A4) is likely to be met. This is the case in placebo-controlled randomized trials with noncompliance where measurements on run-in placebo compliance may very well meet assumption (A4). With concern for compliance mismeasurement, recording run-in compliance may thus be favourable.

On theoretical grounds and on the basis of simulation experiments, we recommend the improved error-adjusted estimator of Section 3.1. This estimator was designed so that adjustment for measurement error does not compromise the power of tests of the causal null. This is attractive, knowing that standard tests of the causal null hypothesis (i.e., that the instrument R is independent of outcome) ignore exposure measurements and are thus valid in the presence of measurement error. Because the proposed estimator does not converge uniformly to a normal distribution, we recommend the uniform confidence intervals of Section 3.2.

For illustrative purposes, we have developed this work under structural mean models which assume linear exposure effects that are not modified by pre-exposure covariates. Extensions to linear structural mean models that allow for effect modification by baseline covariates are methodologically straightforward, but computationally more demanding. Finally, we believe our results to be more broadly useful as they suggest, in line with Gustafson (2005), that incorporating a little prior information on a weakly identified nuisance parameter may yield substantial efficiency improvements for the target parameter. In addition, they indicate how such prior information may be adopted.

REFERENCES

Andrews, D.W.K., and Ploberger, W. (1994). Optimal tests when a nuisance parameter is present only under alternative. *Econometrica* **62**, 1383-1414.

Angrist, J.D., Imbens, G.W. and Rubin, D.B. (1996). Identification of causal effects using instrumental variables (with discussion). *J. Am. Statist. Assoc.* **91**, 444-455.

Buzas, J.S. and Stefanski, L.A. (1996). Instrumental variable estimation in generalized linear measurement error models. *J. Am. Statist. Assoc.* **91**, 999-1006.

Casas, J.P., Bautista, L.E., Smeeth, L., et al. (2005). Homocysteine and stroke: evidence on a causal link from Mendelian randomisation. *Lancet*, **365**, 224232.

Carroll, R.J., Ruppert, D., Crainiceanu, C.M., et al. (2004). Nonlinear and nonparametric regression and instrumental variables. *J. Am. Statist. Assoc.* **99**, 736-750.

Carroll, R.J., Ruppert, D., Stefanski, L.A. and Crainiceanu, C.M. (2006). *Measurement Error in Nonlinear Models: A Modern Perspective, Second Edition.* CRC Press.

Davies, R.B. (1977). Hypothesis testing when a nuisance parameter is present only under alternative. *Biometrika* **64**, 247-254.

Davies, R.B. (1987). Hypothesis-testing when a nuisance parameter is present only under the alternative. *Biometrika* **74**, 33-43.

Dunn, G. (1999). The problem of measurement error in modelling the effect of compliance in a randomized trial. *Statist. Med.* **18**, 2863-2877.

Goetghebeur, E. and Lapp, K. (1997). The effect of treatment compliance in a placebo-controlled trial: regression with unpaired data. *Appl. Statist.* **46**, 351-364.

Goetghebeur, E. and Vansteelandt, S. (2005). Structural mean models for compliance analysis in randomized clinical trials and the impact of errors on measures of

exposure. *Statist. Meth. Med. Res.* **14**, 397-415.

Greenland, S. (2000). An introduction to instrumental variables for epidemiologists. *Int. J. Epidem.* **29**, 722-729.

Gustafson, P. (2005). On model expansion, model contraction, identifiability and prior information: two illustrative scenarios involving mismeasured variables. *Statist. Science* **20**, 111-129.

Hansen, B.E. (1992). The likelihood ratio test under nonstandard conditions - testing the Markov switching model of GNP. *J. Appl. Econom.* **7**, S61-S82.

Hernan, M.A. and Robins, J.M. (2006). Instruments for causal inference - An epidemiologist's dream? *Epidemiology* **17**, 360-372.

Leigh, J.P. and Schembri, M. (2004). Instrumental variables technique: cigarette price provided better estimate of effects of smoking on SF-12. *J. Clin. Epidem* **57**, 284-293.

Martens, E.P., Pestman, W.R., de Boer, A., Belitser, S.V. and Klungel, O.H. (2006). Instrumental variables application and limitations. *Epidemiology* **17**, 260-267.

Robins, J.M. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Commun. Statist.: Theory Meth.* **23**, 2379-2412.

Robins, J.M. (2005). Optimal structural nested models for optimal sequential decisions. In D.Y. Lin and P.J. Haegerty, editors, Proceedings of the second Seattle Symposium in Biostatistics, pages 189326. Springer Verlag, New York.

Robins, J. 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.

Serfling, R.J. (1980). *Approximation Theorems of Mathematical Statistics.* Wiley: New York.

Sheehan, N.A. and Didelez, V. (2005). Mendelian randomisation and instrumental variables for causal inference. *Genetic Epidemiology* **29**, 277-277.

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. *J. Am. Statist. Assoc.* **99**, 16-25.

Ten Have, T.R., Joffe, M.M., Lynch, K.G., Brown, G.K., Maisto, S.A. and Beck, A.T. (2007). Causal Mediation Analyses with Rank Preserving Models. *Biometrics*, Published article online: 20-Mar-2007, doi: 10.1111/j.1541-0420.2007.00766.x

Urquhart, J. and De Klerk, E. (1998). Contending paradigms for the interpretation of data on patient compliance with therapeutic drug regimens. *Statist. Med.* **17**, 251-267.

van der Vaart, A.W. (1998). *Asymptotic Statistics*. Cambridge University Press: Cambridge.

Vansteelandt, S. and Goetghebeur, E. (2005). Sense and sensitivity when correcting for observed exposures in randomized clinical trials. *Statist. Med.* **24**, 191-210.

APPENDIX: PROOFS

Proposition 1. Model $\mathcal{A}$ implies model $\mathcal{B}$ because

$$\begin{aligned} E\left(Y_{i0}|\mathbf{X}_i, R_i\right) &= E\left\{Y_i - \gamma(\mathbf{S}_i, R_i; \psi^*) Z_i | \mathbf{X}_i, R_i\right\} \\ &= E\left[Y_i - \gamma(\mathbf{S}_i, R_i; \psi^*)\left\{W_i - \delta(\mathbf{X}_i, R_i)\right\} | \mathbf{X}_i, R_i\right] \end{aligned}$$

by (A3) and because $E\left(Y_{i0}|\mathbf{X}_i, R_i\right) = E\left(Y_{i0}|\mathbf{X}_i\right)$ by (A1). Note that this does not require assumptions about the conditional association between Y_i and W_i, given Z_i, suggesting that this continues to hold when measurement error is differential.

To show that (6) is the only restriction (other than (5)) imposed on the observed data law, we proceed as in Robins and Rotnitzky (2004) by exhibiting for any observed data law satisfying (5) and (6), a joint law of the full data $(Y, \{Y_{rz}, \forall r, z\}, Z, W, R, X, T)$ satisfying the restrictions of model $\mathcal{A}$, where Y_{rz} is the potential outcome that would have been observed for given subject following exposure to $(R, Z) = (r, z)$, all other experimental conditions being the same as in the considered study. Given $(R = r, Z = z, W = w, X = x, T = t, Y = y)$, we define $Y_{rz} = y$ to satisfy (A2). We set $f(Z|R = r, W = w, X = x, T = t, Y = y)$ equal to an arbitrary density with conditional mean $w - \delta$. We define $f(Y_{r0}|R = r, Z = z, W = w, X = x, T = t, Y = y)$ to be an arbitrary density with conditional mean $y - \gamma(x, r; \psi^*)z$. In addition, given $(Z = z, W = w, X = x, T = t, Y = y)$, we set $Y_{r0} = Y_{r'0} \equiv Y_0$ for each (r, r') to satisfy (A1). By (6), the conditional distribution of Y_0 then also satisfies $E(Y_0|X = x, T = t, R) = E(Y_0|X = x, T = t)$ for each (x, t). Remaining features of the full data density can be chosen arbitrarily.

Theorem 1. Let for simplicity of exposition, but without loss of generality, $\gamma(\mathbf{X}_i, R_i; \psi) = \psi R_i$. Define $U_{i\delta} = d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)\,[Y_i$
$-\psi(W_i - \delta)R_i - q(\mathbf{T}_i, \mathbf{X}_i)]$ and $U_{i\psi} = d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)\,[Y_i - \psi(W_i - \delta)R_i - q(\mathbf{T}_i, \mathbf{X}_i)]$ the estimating functions for δ^* and ψ^*, respectively. Then standard asymptotic theory for M-estimators (van der Vaart, 1998) and Taylor expansions show that

$$\begin{aligned} 0 &= \frac{1}{\sqrt{n}}\sum_{i=1}^{n} U_{i\delta} + E\left(\frac{\partial U_{i\delta}}{\partial \psi}\right)\sqrt{n}(\hat{\psi} - \psi^*) + E\left(\frac{\partial U_{i\delta}}{\partial \delta}\right)\sqrt{n}(\hat{\delta} - \delta^*) \\ &\quad + \frac{1}{2}E\left(\frac{\partial^2 U_{i\delta}}{\partial \psi \partial \delta}\right)\sqrt{n}(\hat{\psi} - \psi^*)(\hat{\delta} - \delta^*) + o_p(1) \end{aligned} \tag{11}$$

from which

$$\begin{aligned} \sqrt{n}(\hat{\delta} - \delta^*)\frac{\hat{\psi} + \psi^*}{2} &= o_p(1) - E^{-1}\left\{d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i\right\} \\ &\quad \times \left[\frac{1}{\sqrt{n}}\sum_{i=1}^{n} U_{i\delta} - E\left\{d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta^*)R_i\right\}\sqrt{n}(\hat{\psi} - \psi^*)\right] \end{aligned}$$

Plugging this into a first order Taylor expansion of $U_{i\psi}$, shows that $\sqrt{n}(\hat{\psi} - \psi^*)$ equals

$$- \left[E\left\{ d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta^*)R_i \right\} - \frac{E\left\{ d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\}}{E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\}} E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta)R_i \right\} \right]^{-1}$$
$$\times \left[\frac{1}{\sqrt{n}} \sum_{i=1}^{n} U_{i\psi} - \frac{E\left\{ d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\}}{E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\}} U_{i\delta} \right] + o_p(1)$$

It follows that $\sqrt{n}(\hat{\psi} - \psi) = O_p(1)$ and that Part 1 of Theorem 1 holds.

Note that the last 2 terms in (11) can be replaced by $E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\} \left\{ \psi + O_p(n^{-1/2}) \right\}$ $\times \sqrt{n}(\hat{\delta} - \delta)$, from which $\sqrt{n}(\hat{\delta} - \delta)\psi = \sqrt{n}(\hat{\delta} - \delta)(\hat{\psi} + \psi^*) \left\{ 1/2 + o_p(1) \right\}$ equals

$$- \left[E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\} - \frac{E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta^*)R_i \right\}}{E\left\{ d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta^*)R_i \right\}} E\left\{ d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)R_i \right\} \right]^{-1}$$
$$\times \left[\frac{1}{\sqrt{n}} \sum_{i=1}^{n} U_{i\delta} - \frac{E\left\{ d_\delta(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta^*)R_i \right\}}{E\left\{ d_\psi(R_i, \mathbf{T}_i, \mathbf{X}_i)(W_i - \delta^*)R_i \right\}} U_{i\psi} \right] + o_p(1)$$

The latter expression is bounded in probability (under standard regularity conditions). It follows that, as ψ^* goes to zero with increasing sample size, $\hat{\delta}$ does not converge to δ^* at root-n rate and hence is not uniformly root-n consistent. In particular, there is no root-n consistent estimator of δ^* under model $\mathcal{A}$ at $\psi^* = 0$.

Part 2 of Theorem 1 is immediate from Robins (1994).

Proof of Theorem 2. Let for simplicity of exposition, but without loss of generality, $\gamma(\mathbf{X}_i, R_i; \psi) = \psi R_i$. Then standard asymptotic theory for M-estimators (van der Vaart, 1998) and Taylor expansions of the estimating functions (9) for ψ^* w.r.t. $\hat{\delta}(\psi)$ shows that (9) equals

$$\frac{1}{\sqrt{n}} \sum_{i=1}^{n} P\{\hat{\delta}(\psi) \in \Delta\} U_{i\psi}(\psi, \delta) + P\{\hat{\delta}(\psi) \notin \Delta\} U_{i\psi}(\psi, 0) + o_p(1) - \Big[P\{\hat{\delta}(\psi) \in \Delta\}$$
$$+ \left\{ \varphi\left(\frac{\Delta_l - \delta}{\sigma(\psi)/(\sqrt{n}|\psi|)} \right) - \varphi\left(\frac{\Delta_u - \delta}{\sigma(\psi)/(\sqrt{n}|\psi|)} \right) \right\} \frac{\sqrt{n}|\psi|\delta}{\sigma(\psi)} \Big] \frac{E\left\{ d_\psi(R, T)R \right\}}{E\left\{ d_\delta(R, T)R \right\}} U_{i\delta}(\psi, \delta) \quad (12)$$

That the remainder term converges to zero in probability for any fixed ψ can be seen because, for some $\tilde{\delta}$ on the open line segment between $\hat{\delta}(\psi)$ and δ^* (under standard regularity conditions which include uniform convergence of $n^{-1} \sum_{i=1}^{n} U_{i\psi}(\psi, \delta)$

w.r.t. δ), the remainder term equals

$$\left[P_{\delta=\tilde{\delta}}\{\hat{\delta}(\psi)\in\Delta\}E\left\{\frac{\partial^2}{\partial\delta^2}U_{i\psi}(\psi,\tilde{\delta})\right\}+2\frac{\partial}{\partial\delta}P_{\delta=\tilde{\delta}}\{\hat{\delta}(\psi)\in\Delta\}E\left\{\frac{\partial}{\partial\delta}U_{i\psi}(\psi,\tilde{\delta})\right\}\right.$$
$$\left.+\frac{\partial^2}{\partial\delta^2}P_{\delta=\tilde{\delta}}\{\hat{\delta}(\psi)\in\Delta\}E\left\{U_{i0}(\psi)-U_{i\psi}(\psi,\tilde{\delta})\right\}\right]\frac{\sqrt{n}}{2}\{\hat{\delta}(\psi)-\delta^*\}^2+o_p(1)$$

Here, the first term is zero. Because $E\left\{\partial U_{i\psi}(\psi,\tilde{\delta})/\partial\delta\right\}=O_p(1)\psi$ under standard regularity conditions and $\sqrt{n}\{\hat{\delta}(\psi)-\delta^*\}^2=O_p(1)n^{-1/2}\psi^{-2}$, the second term is

$$O_p(1)\left\{\varphi\left(\frac{\Delta_l-\tilde{\delta}}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)-\varphi\left(\frac{\Delta_u-\tilde{\delta}}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)\right\}\frac{1}{\sigma(\psi)}=o_p(1)$$

for any fixed ψ. Because $E\left\{U_{i0}(\psi)-U_{i\psi}(\psi,\tilde{\delta})\right\}=O_p(1)\delta\psi$, the third term is

$$\left\{\varphi\left(\frac{\Delta_l-\tilde{\delta}}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)-\varphi\left(\frac{\Delta_u-\tilde{\delta}}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)\right\}\frac{n|\psi|\psi\tilde{\delta}}{\sigma(\psi)^3}(\Delta_l-\Delta_u)=o_p(1)$$

for any fixed ψ because $x^a\varphi(x)\to 0$ as $x\to\infty$ for arbitrary $a>0$.

Because the estimating functions in (12) have mean and variance depending on the sample size, we use the triangular array Central Limit Theorem (Serfling, 1980, p.31) to derive the asymptotic distribution of (9) for fixed ψ. Application of this Theorem shows that for arbitrary fixed ψ, the estimating functions in (9) are asymptotically normally distributed under the weak regularity condition that the standard deviation of the estimating functions $\tilde{U}_i(\psi)$, as defined by (12), is bounded (i.e. $O(1)$) and that asymptotically $E\|\tilde{U}_i(\psi)-E\{\tilde{U}_i(\psi)\}\|^k=o(n^{k/2-1})$. Because for any fixed $\psi^*\neq 0$ and $\delta^*\in\Delta=]\Delta_l,\Delta_u[$, $P\{\hat{\delta}(\psi^*)\in\Delta\}$ converges to 1, it follows under these conditions that $n^{-1/2}\sum_{i=1}^n\tilde{U}_i(\psi^*)$ will be asymptotically normally distributed with mean zero and finite variance, which is given by the variance of (12). Within faster than root-n shrinking neighbourhoods of zero (i.e. if $\psi^*=kn^{-a}$ for some constant k and $a>1/2$), the remainder term in the Taylor series expansion is still $o_p(1)$. Further, $P\{\hat{\delta}(\psi^*)\in\Delta\}$ converges to 0 and $U_0(\psi^*)$ has mean converging to zero at 1 over n^a-rate. It then again follows that

$n^{-1/2}\sum_{i=1}^{n}\tilde{U}_i(\psi^*)$ is asymptotically normally distributed with mean zero and finite variance. Finally, within 1 over root-n shrinking neighbourhoods of zero (i.e. if $\psi^* = kn^{-1/2}$ for some constant k), the remainder term in the Taylor series expansion is bounded in probability, but not $o_p(1)$. The significant contribution of the squared term $\sqrt{n}\{\hat{\delta}(\psi^*)-\delta^*\}^2$ implies that $n^{-1/2}\sum_{i=1}^{n}\tilde{U}_i(\psi^*)$ may not converge to a normal distribution, nor to a mean zero distribution along such sequences. The implications of this will be discussed in the next paragraph.

To gain insight into the asymptotic distribution of $\tilde{\psi}$ (rather than its estimating function), we make a further Taylor series expansion of the estimating functions, evaluated at $\tilde{\psi}$. This shows that for any fixed ψ

$$
\begin{aligned}
0 = {} & \frac{1}{\sqrt{n}}\sum_{i=1}^{n} P\{\hat{\delta}(\psi)\in\Delta\}U_{i\psi}(\psi,\delta) + P\{\hat{\delta}(\psi)\notin\Delta\}U_{i\psi}(\psi,0) + o_p(1) - \Big[P\{\hat{\delta}(\psi)\in\Delta\} \\
& + \left\{\varphi\left(\frac{\Delta_l-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right) - \varphi\left(\frac{\Delta_u-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)\right\}\frac{\sqrt{n}|\psi|\delta}{\sigma(\psi)}\Big]\frac{E\{d_\psi(R,T)R\}}{E\{d_\delta(R,T)R\}}U_{i\delta}(\psi,\delta) \\
& + \left(P\{\hat{\delta}(\psi)\in\Delta\}E\left\{\frac{\partial}{\partial\psi}U_{i\psi}(\psi,\delta)\right\} + P\{\hat{\delta}(\psi)\notin\Delta\}E\left\{\frac{\partial}{\partial\psi}U_{i\psi}(\psi,0)\right\}\right. \\
& + \left\{\varphi\left(\frac{\Delta_l-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right) - \varphi\left(\frac{\Delta_u-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)\right\}\frac{\sqrt{n}|\psi|\delta(\Delta_l-\Delta_u)}{\sigma(\psi)}E\{d_\psi(R,T)R\} \\
& - \left[P\{\hat{\delta}(\psi)\in\Delta\} + \left\{\varphi\left(\frac{\Delta_l-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right) - \varphi\left(\frac{\Delta_u-\delta}{\sigma(\psi)/(\sqrt{n}|\psi|)}\right)\right\}\frac{\sqrt{n}|\psi|\delta}{\sigma(\psi)}\right]\frac{E\{d_\psi(R,T)R\}}{E\{d_\delta(R,T)R\}} \\
& \left.\times E\{d_\delta(R,T)R(W-\delta)\}\right)\sqrt{n}(\tilde{\psi}-\psi) \qquad (13)
\end{aligned}
$$

That the remainder term converges to zero in probability for any fixed ψ can be seen using a similar derivation as before. We conclude that, up to an $o_p(1)$ term and for fixed ψ, $\sqrt{n}(\tilde{\psi}-\psi)$ is a linear transformation of $n^{-1/2}\sum_{i=1}^{n}\tilde{U}_i(\psi)$ and thus shares its asymptotic properties. Specifically, within faster and slower than 1 over root-n shrinking neighbourhoods of zero (and in particular at arbitrary fixed ψ), $\sqrt{n}(\tilde{\psi}-\psi)$ is asymptotically normally distributed with mean zero and finite variance under weak regularity conditions. Within 1 over root-n neighbourhoods of zero, $\sqrt{n}(\tilde{\psi}-\psi)$ may be asymptotically biased and not normally distributed.

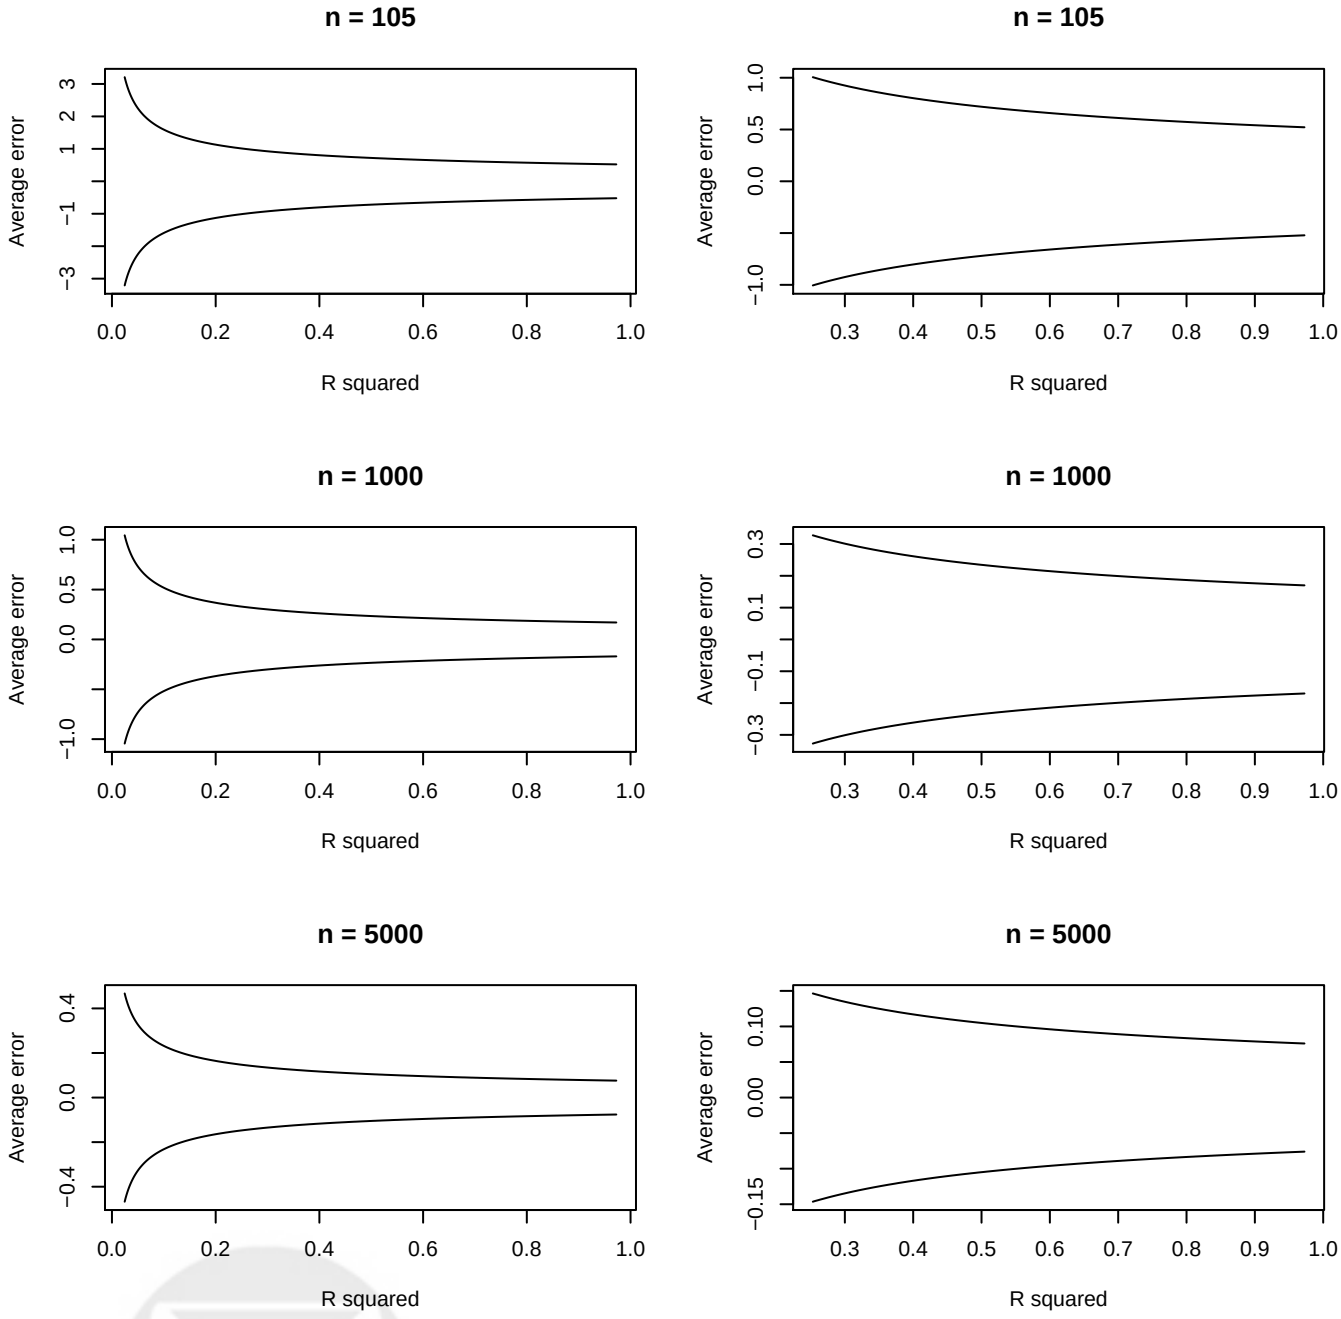

Figure 1: Curves indicating the tuples (R^2, δ) where the standard SMM estimator and the error-adjusted instrumental variable estimator have the same mean squared error, for different sample sizes $n = 105, 1000$ and 5000 and with R^2 equalling the multiple correlation coefficient for the regression of Z on T. Left: for R^2 from 0 to 1; Right: for R^2 from 0.25 to 1.

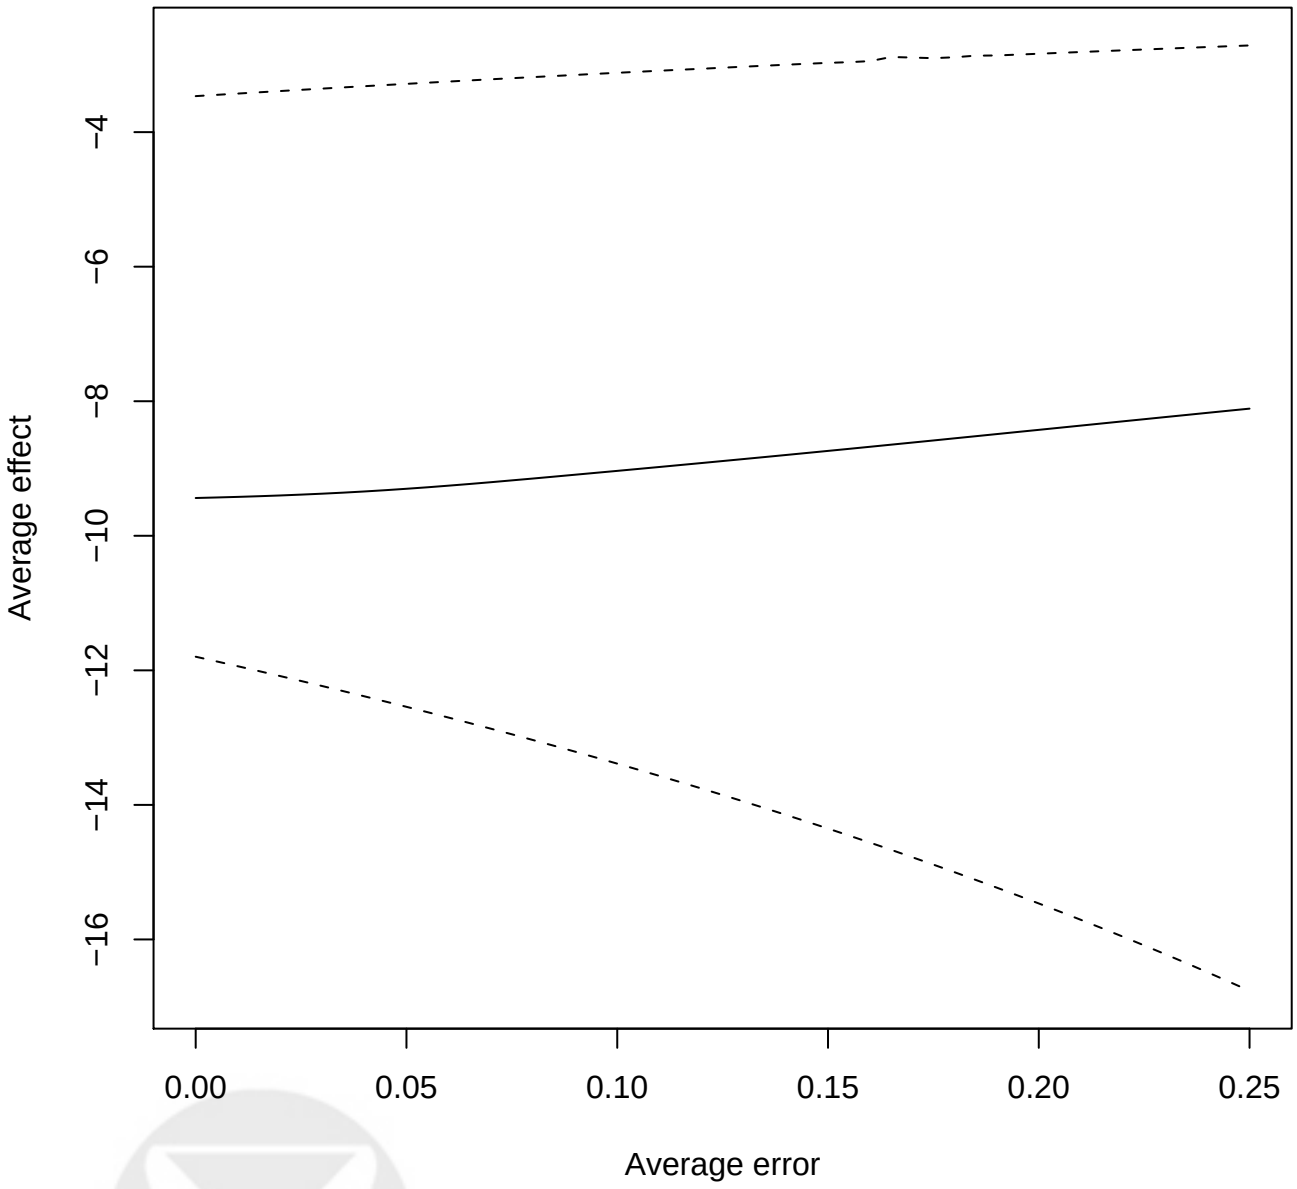

Figure 2: Improved error-adjusted estimate, along with uniform 95% confidence intervals in function of the maximum error mean Δ_u, with $\Delta = [-\Delta_u, \Delta_u]$.

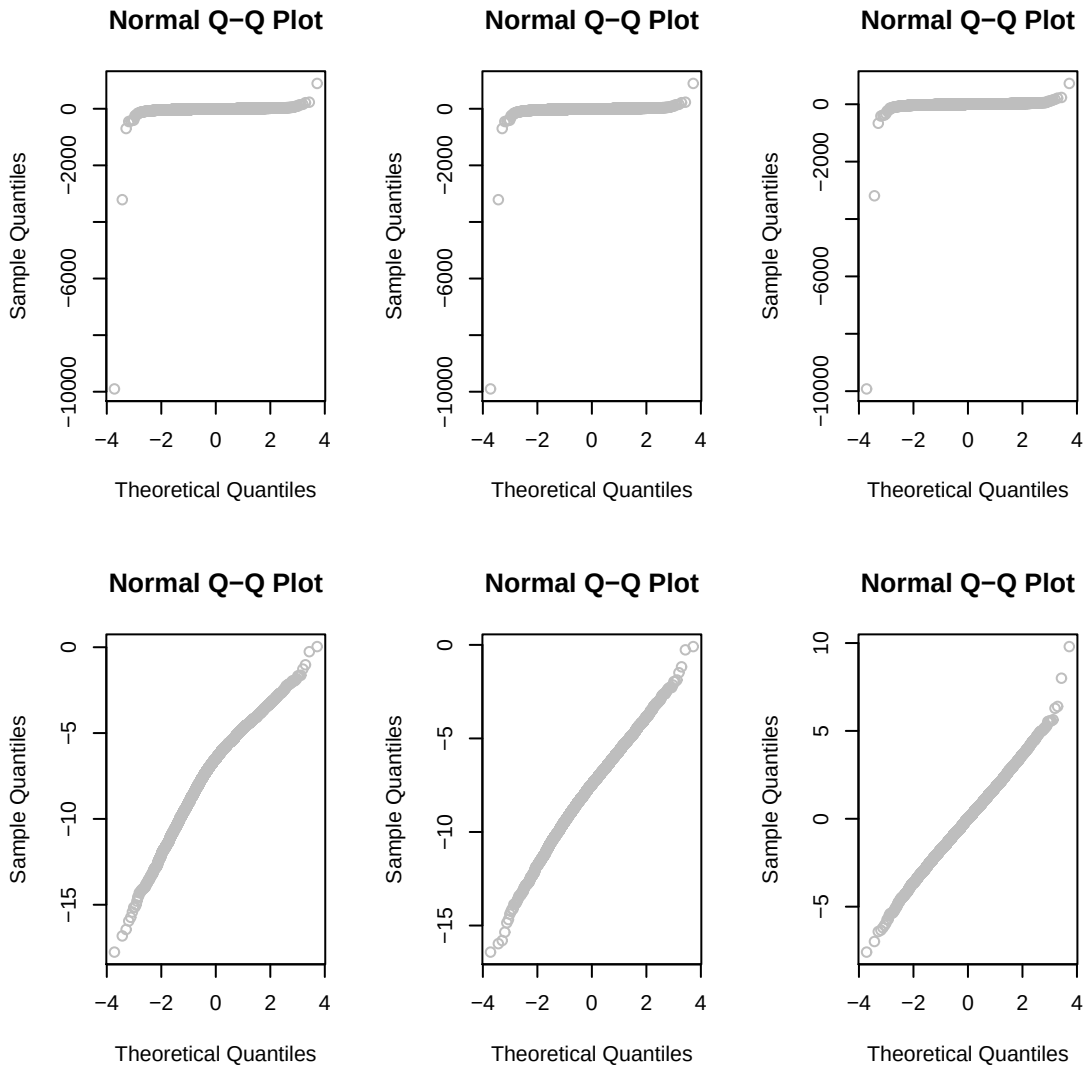

Figure 3: QQ-plots for $n = 105$. Row 1: error-adjusted estimator IV1; Row 2: improved error-adjusted estimator IV3; Column 1: $\psi = -7.5, \delta = 0.15$; Column 2: $\psi = -7.5, \delta = 0$; Column 1: $\psi = 0, \delta = 0$.

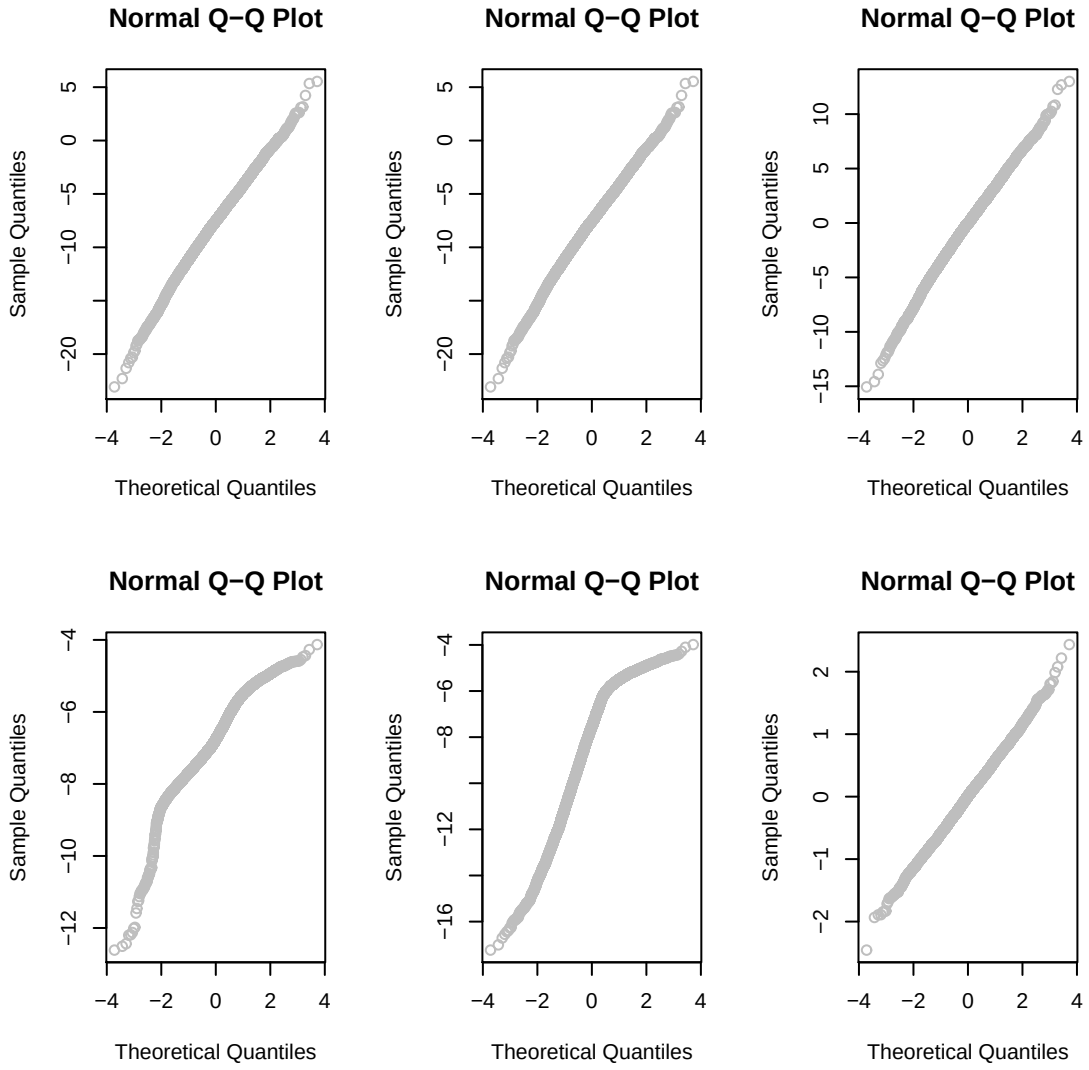

Figure 4: QQ-plots for $n = 1000$. Row 1: error-adjusted estimator IV1; Row 2: improved error-adjusted estimator IV3; Column 1: $\psi = -7.5, \delta = 0.15$; Column 2: $\psi = -7.5, \delta = 0$; Column 1: $\psi = 0, \delta = 0$.

Table 1: Bias of the different estimators and coverage and average length of corresponding 95% confidence intervals.

Δ	n	ψ	δ	Bias				Coverage					Average length CI				
				STD	IV1	IV2	IV3	STD	IV1	IV2	IV3	UI	STD	IV1	IV2	IV3	UI
0.5	105	-7.5	0.15	1.11	-3.77	-2.65	0.68	86.8	96.5	99.8	97.7	99.8	5.87	3039	56.2	13.3	18.8
0.5	105	-7.5	0	-0.020	-3.77	-2.31	-0.046	93.7	96.5	99.9	98.7	100	6.96	3039	55.7	10.8	29.2
0.5	105	0	0	-0.015	-3.63	-0.019	-0.0096	93.5	96.5	100	94.1	96.2	6.94	3027	63.8	8.88	21.1
0.5	1000	-7.5	0.15	1.13	-0.15	-0.28	0.81	36.4	95.1	99.9	97.8	100	1.90	14.0	14.0	4.83	11.4
0.5	1000	-7.5	0	-0.0048	-0.15	-0.52	-0.62	95.0	95.1	98.1	95.1	100	2.25	14.0	14.0	12.4	16.7
0.5	1000	0	0	-0.0042	-0.14	0.0032	-0.0036	94.9	95.2	100	95.0	96.2	2.24	14.0	13.9	2.78	5.86
0.25	105	-7.5	0.15	1.11	-3.77	0.59	1.06	86.8	96.5	100	97.0	99.5	5.87	3039	56.4	8.42	10.9
0.25	105	-7.5	0	-0.020	-3.77	-0.63	-0.062	93.7	96.5	100	98.8	99.9	6.96	3039	56.8	10.7	14.2
0.25	105	0	0	-0.015	-3.63	-0.013	-0.010	93.5	96.5	100	94.2	94.9	6.94	3027	67.8	8.88	11.6
0.25	1000	-7.5	0.15	1.13	-0.15	0.42	0.78	36.4	95.1	100	98.0	99.9	1.90	14.0	13.9	4.35	5.69
0.25	1000	-7.5	0	-0.0048	-0.15	-0.37	-0.19	95.0	95.1	100	95.8	99.6	2.25	14.0	14.0	5.83	7.61
0.25	1000	0	0	-0.0042	-0.14	-0.0012	-0.0036	94.9	95.2	100	95.0	95.5	2.24	14.0	13.9	2.78	3.43
0.05	105	-7.5	0.15	1.11	-3.77	1.07	1.11	86.8	96.5	100	91.3	94.8	5.87	3039	63.9	6.54	7.47
0.05	105	-7.5	0	-0.020	-3.77	-0.052	-0.015	93.7	96.5	100	96.3	98.2	6.96	3039	64.4	7.87	9.07
0.05	105	0	0	-0.015	-3.63	-0.014	-0.016	93.5	96.5	100	94.1	94.3	6.94	3027	67.9	7.53	8.51
0.05	1000	-7.5	0.15	1.13	-0.15	1.02	1.11	36.4	95.1	100	54.2	74.9	1.90	14.0	13.9	2.32	2.78
0.05	1000	-7.5	0	-0.0048	-0.15	-0.031	-0.0063	95.0	95.1	100	98.3	99.7	2.25	14.0	13.9	2.83	3.43
0.05	1000	0	0	-0.0042	-0.14	-0.0037	-0.0042	94.9	95.2	100	95.0	95.1	2.24	14.0	13.9	2.36	2.55