

How does a local instrumental variable method perform across settings with instruments of differing strengths? A simulation study and an evaluation of emergency surgery

Silvia Moler-Zapata¹  | Richard Grieve¹  | Anirban Basu^{2,3}  | Stephen O'Neill¹

¹Department of Health Services Research and Policy, London School of Hygiene and Tropical Medicine, London, UK

²Department of Pharmacy, and Departments of Health Services and Economics, The Comparative Health Outcomes, Policy, and Economics (CHOICE) Institute, University of Washington, Seattle, Washington, USA

³National Bureau of Economic Research, Cambridge, Massachusetts, USA

Correspondence

Silvia Moler-Zapata, Department of Health Services Research and Policy, London School of Hygiene and Tropical Medicine, 15-17 Tavistock Place, London WC1H 9SH, UK.
Email: silvia.moler@lshtm.ac.uk

Funding information

Medical Research Council/National Institute for Health Research, Better Methods Better Health, Grant/Award Number: 19091; National Institute for Health Research ARC North Thames

Abstract

Local instrumental variable (LIV) approaches use continuous/multi-valued instrumental variables (IV) to generate consistent estimates of average treatment effects (ATEs) and Conditional Average Treatment Effects (CATEs). There is little evidence on how LIV approaches perform according to the strength of the IV or with different sample sizes. Our simulation study examined the performance of an LIV method, and a two-stage least squares (2SLS) approach across different sample sizes and IV strengths. We considered four ‘heterogeneity’ scenarios: homogeneity, overt heterogeneity (over measured covariates), essential heterogeneity (unmeasured), and overt and essential heterogeneity combined. In all scenarios, LIV reported estimates with low bias even with the smallest sample size, provided that the instrument was strong. Compared to 2SLS, LIV provided estimates for ATE and CATE with lower levels of bias and Root Mean Squared Error. With smaller sample sizes, both approaches required stronger IVs to ensure low bias. We considered both methods in evaluating emergency surgery (ES) for three acute gastrointestinal conditions. Whereas 2SLS found no differences in the effectiveness of ES according to subgroup, LIV reported that frailer patients had worse outcomes following ES. In settings with continuous IVs of moderate strength, LIV approaches are better suited than 2SLS to estimate policy-relevant treatment effect parameters.

KEY WORDS

emergency surgery, instrument strength, instrumental variables, tendency to operate

1 | INTRODUCTION

The personalization of treatment choice can be informed by comparative effectiveness research that exploits the widespread availability of electronic health records (EHRs), but requires methods that address confounding and heterogeneity. For conventional linear Instrumental Variable (IV) methods, such as two-stage least squares (2SLS) to identify policy-relevant estimands such as the Average Treatment Effect (ATE) or Conditional Average Treatment Effects (CATEs), it is required that there is no essential heterogeneity (Heckman et al., 2006). Essential heterogeneity arises when treatment effects differ over levels of unmeasured confounders, in which case 2SLS no longer identifies the ATE, even if the instrument is strong and valid (Heckman

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. Health Economics published by John Wiley & Sons Ltd.

et al., 2006). Essential heterogeneity is a major concern in health care, as it is commonly the case that there are biological correlations between risk factors, some of which remain unobserved to the analyst.

In the presence of essential heterogeneity, Local Instrumental Variable (LIV) approaches can provide consistent estimates of the ATE and CATEs (Heckman & Vytlacil, 2005). LIV methods draw on theory about individual's choices to identify 'marginal treatment effects' (MTEs) for individuals at the 'margin of treatment choice' (Bjorklund & Moffitt, 1983; Heckman & Vytlacil, 1999). These MTEs are identified for individuals for whom the level of the IV is such that observed characteristics encouraging treatment (including the IV) and unobserved characteristics discouraging treatment are balanced, so there is equipoise about the treatment decision. Here, a small change (or nudge) in the level of a valid, continuous IV 'tips the balance' for the treatment decision for these marginal patients, without changing the distribution of the underlying risk factors. Therefore, comparing mean outcomes between two groups of patients only separated by a small change in the IV, identifies MTEs for individuals who *comply* with the change in treatment, due to that small change in the IV. A continuous instrument with sufficient support allows all individuals to be defined as 'compliers' at some level of the IV (Heckman & Vytlacil, 1999). Hence, given observed covariates, MTEs can be estimated along the continuum of the IV, and aggregated to provide CATEs and ATEs (Heckman & Vytlacil, 1999, 2001, 2005).

The theoretical properties of these LIV methods in settings with essential heterogeneity have been discussed by Heckman et al. (2006), Basu et al. (2007) and Angrist & Fernández-Val, 2011 *inter alia*. However, most simulation studies of IV methods only consider treatment effects that are homogeneous, or heterogenous according to measured factors (overt heterogeneity) (Martínez-Cambor et al., 2019; Terza et al., 2008a, 2008b). Studies that have considered essential heterogeneity, have found that 2SLS provides inconsistent estimates of the ATE (Basu, Coe, & Chapman, 2018; Brooks et al., 2018; Chapman & Brooks, 2016), whereas Basu (2014) reports that a LIV method could provide consistent estimates of the ATE and CATE in finite samples for different types of outcomes. LIV methods have now been applied across a multitude of settings including cardiovascular and bariatric surgery, universal child care programs and transfers to intensive care units (Basu, Jones, & Rosa Dias, 2018; Cornelissen et al., 2018; Grieve et al., 2019; Reynolds et al., 2021).

A major barrier to wider use of potentially valid IVs in general is that if the IV is only weakly associated with treatment assignment, then IV estimators can provide very biased and imprecise estimates (Bound et al., 1995; Nelson & Startz, 1990; Stock & Yogo, 2005). Weak IVs can also amplify the bias arising due to violations of the other assumptions (Bound et al., 1995; Small & Rosenbaum, 2008). While current practice tends to rely on the first-stage F-statistic exceeding the value of 10 (Staiger & Stock, 1997), recent developments in the weak identification literature for IV models have revealed the shortcomings of an unequivocal decision rule for assessing weak identification (Andrews et al., 2019; Keane & Neal, 2023; Lee et al., 2021; Moffitt & Zahn, 2022). For LIV to provide consistent, precise estimates of ATE or CATEs, requires a strong continuous/multi-valued IV with sufficient support to ensure that there is a level of the IV at which each unit 'complies' (i.e., is selected into treatment according to the level of the IV). However, no study has assessed the levels of IV strength that are required for an LIV estimator to perform well, nor how performance may differ according to the sample size available, in settings with essential heterogeneity.

This paper addresses this gap in the literature by contrasting LIV with the commonly used 2SLS estimator in Monte Carlo simulations, motivated by a case study which highlights typical issues pertaining to heterogeneity, sample size and IV strength. We simulate four scenarios: two of them under restrictive assumptions about heterogeneity (A: homogeneity; B: overt heterogeneity), one where treatment effects are allowed to be heterogenous according to an unmeasured confounder (C: essential heterogeneity), and one where both forms of heterogeneity are present (D: overt and essential heterogeneity). Across all scenarios, ATE and CATE are the parameters of interest.

This paper is structured as follows. In Section 2, we outline the motivating example. In Section 3, we define the estimands and identification assumptions for 2SLS and LIV and present the methods for the simulation study. In Section 4, we present the results of the simulation study and the case study. In Section 5, we discuss how this study adds to the literature and the implications for further research.

2 | MOTIVATING EXAMPLE: THE ESORT STUDY

The ESORT (Emergency Surgery OR noT) study evaluated the effectiveness of ES for acute gastrointestinal conditions. The primary outcome of the study was the number of 'days alive and out of hospital' (DAOH) at 90-days (see (Hutchings et al., 2022) for details), which encompasses mortality and total length of hospital stay (LOS). The study exemplifies the key issues that arise when applying IV methods to EHR data to provide policy-relevant estimates of comparative effectiveness (ESORT Study Group, 2020; Hutchings et al., 2021, 2022). Patients presented as emergency admissions and were selected for either ES, or alternative non-emergency surgery (NES) interventions such as medical management or delayed surgery,

according to unmeasured characteristics such as the severity of the disease, and hence unmeasured confounding and essential heterogeneity were major concerns.

The ESORT study followed Keele et al. (2018) and developed a continuous preference-based IV for ES receipt to evaluate the effectiveness of ES for three acute gastrointestinal conditions: acute appendicitis, gallstone disease and abdominal wall hernia, using routine hospitalization data from the hospital episode statistics (HES) inpatient database in England. The IV was the hospital's tendency to operate (TTO), a proxy measure of the hospital's latent preference for ES, defined as the proportion of eligible emergency admissions in each of 174 hospitals who had ES in the year preceding each admission. Given a relevant IV, two main assumptions need to hold: (i) conditional on the variables included in the models, the hospital's TTO was not correlated with the patient's outcome except through treatment assignment, (ii) it does not increase the probability of treatment for an individual at some value of the IV, but decrease it for higher values. The study design had some important features to support this assumption. First, in this emergency setting, patients were unlikely to select the hospital according to quality of care. Second, the study only included direct admissions to hospital, so there was no scope to transfer the patient according to physician or patient choice. Third, information was collated on a rich set of proxies for the hospital's quality of acute care, including rates of mortality and emergency admissions in previous years, which were included in the models as fixed effects. Fourth, observed covariates, were balanced across all levels of the TTO, which helped support the requisite assumption that the IV also balanced unmeasured confounders (Hutchings et al., 2022; Moler-Zapata et al., 2022). The requisite assumption that the IV has a monotonic effect on treatment receipt could not be formally tested on the data. However, it was deemed plausible in this setting, as it seems unlikely that there are patients who would receive ES when admitted to hospitals with low TTO but receive NES when admitted into a hospital with high TTO.

The ESORT study highlighted several outstanding concerns pertaining to IV methods in general, and LIV approach in particular. While the study reported estimates of the ATE, from the outset, there was policy interest in estimating the CATEs, according to baseline covariates including age, number of comorbidities, and levels of frailty. While the sample sizes for each condition, were relatively large, they also differed across conditions, from 268,144 (appendicitis) and 240,977 (gallstone disease), to 106,432 (hernia) patients. There were also differences in the strength of the IV with F-statistics ranging from 141 (acute appendicitis), 739 (hernia) to 9053 (gallstone disease). Hence, the ESORT study further motivated the interest in what strength of continuous IV was required to provide unbiased, efficient estimates of policy relevant estimands such as CATEs in settings with essential heterogeneity, and according to different sample sizes.

3 | METHODS

3.1 | Instrumental variables methods

Throughout we use the Neyman-Rubin potential outcomes framework (Neyman, 1990; Rubin, 1974). Let Y denote the outcome, D denote the treatment status, and Z denote the IV. Let D_Z denote the potential treatment status that would be observed if Z would be set to $Z = z$, and Y_D denote the potential outcome that would be observed if D would be set to $D = d$, with $d \in \{0, 1\}$, such that we observe (Y_D, D_Z, Z) for each individual. For each patient, let $Y_1 = \mu_1(X_O, X_U, \vartheta)$ and $Y_0 = \mu_0(X_O, X_U, \vartheta)$ denote the potential outcomes, where X_O is the vector of observed covariates, X_U is a vector of unmeasured confounders, and ϑ captures all the remaining unobserved random variables. Throughout, we assume exogeneity of the covariates (A1), so that the treatment assignment is the only source of endogeneity, such that $(X_O, X_U) \perp \vartheta$ and $X_O \perp X_U$.

3.1.1 | Identification assumptions

Angrist et al. (1993) defined a series of structural assumptions for the identification of the LATE. Here, following Abadie (2003) and Tan (2006) we make the following assumptions which are the conditional version of the assumptions outlined by Angrist et al. (1993):

(A2)	Unconfoundedness of Z	$(Y_{d_z}, D_z) \perp Z X_O$
(A3)	Exclusion restriction	$Y_{d_z} = Y_d$ with probability 1
(A4)	Relevance	$0 < P(Z = z) < 1$
(A5)	Monotonicity	If $z' > z$ then $D_{z'} \geq D_z$ with probability 1
(A6)	Stable unit treatment value assumption	$D = D_Z$ and $Y = Y_D$

Assumption (A2) requires that Z is as good as randomly assigned within levels of X_O . Assumption (A3) rules out the possibility that Z has a direct effect on the outcome other than through D_z . Assumptions (A2) and (A3) ensure that the only effect of the Z on the outcome is through D_z . This is sometimes called the independence assumption. Assumption (A4) ensures that Z and D_z are correlated conditional on X_O . Assumption (A5) requires that an increase in Z always results in a higher or equal level of treatment assignment. Assumption (A6) requires that one individual's potential outcomes (Y_D) and treatments (D_z) are not influenced by other individuals' levels of Z (i.e., no interference), nor by how the instrument or treatment is delivered (i.e., no different versions of Z or D_z).

3.1.2 | Estimands

Imbens and Angrist (1994) and Angrist et al. (1993) show that, under the assumptions outlined above, the LATE can be defined as $\Delta^{LATE}(x_o, z, z') = E[Y_1 - Y_0 | X_O = x_o, D_z < D_{z'}]$ and is identified by the IV estimand:

$$\frac{E[Y | X_O = x_o, Z = z'] - E[Y | X_O = x_o, Z = z]}{E[D | X_O = x_o, Z = z'] - E[D | X_O = x_o, Z = z]}$$

Vytlacil (2002) and Tan (2006) showed that the independence (A2 and A3) and monotonicity assumptions (A5) of the LATE framework are equivalent to those imposed by a non-parametric selection model, where treatment assignment depends on whether a latent index ($\mu_D(X_O, Z)$) crosses a particular threshold (X_{U_D}):

$$D_z = 1 \{ \mu_D(X_O, Z) \geq X_{U_D} \}$$

where X_{U_D} is a random variable that captures X_U and all other factors influencing treatment assignment but not the outcomes. As in Heckman and Vytlacil (1999, 2001), we can rewrite this equation as $D_z = 1 \{ P(X_O, Z) > V \}$, where $V = F_{X_{U_D}} [X_{U_D} | X_O = x_o, Z = z]$ with $V \perp (Z, X_O)$ and $P(x_o, z) = F_{X_{U_D} | x_o, z} [\mu_D(X_O, Z)]$ is the propensity for treatment, and F represents a cumulative distribution function. Therefore, for any arbitrary distribution of X_{U_D} conditional on X_O and Z , by definition $V \sim Uniform[0, 1]$ conditional on X_O and Z . Then, the MTE can be defined as, $\Delta^{MTE}(x_o, p) = E(Y_1 - Y_0 | X_O = x_o, V = p)$ and Heckman and Vytlacil (1999, 2001) showed that, under the standard IV assumptions, it can be identified by:

$$\frac{\partial E_{\theta}(Y | X_O = x_o, Z = z)}{\partial p} = E_{\theta}[(Y_1 - Y_0) | X_O = x_o, V = p]$$

MTEs can be aggregated directly to obtain estimates of the ATE, as shown in Heckman et al. (2006). Basu (2014) showed that MTEs can be used to derive personalized treatment (PeT) effects for each individual that take into account the plausible range of values that V may take for each patient, in addition to their observed covariates, IV and actual treatment assignment (see Section 3.1.3). The rationale for this approach is that the treatment assignment status provides some information on X_{U_D} . For patients in the treatment group ($D_z = 1$), the propensity to choose treatment based on X_O and Z must outweigh the propensity to choose the comparator strategy based on X_{U_D} , that is, $P(x_o, z) > v$. For patients in the comparator strategy ($D_z = 0$), the opposite is true. The PeT effect for an individual is obtained by averaging the MTEs corresponding to that individual's level of X_O and Z over those values of unobserved variables that are compatible with that patient's treatment assignment. Hence, $\Delta^{PeT}(x_o, p, D) = E(Y_1 - Y_0 | X_O = x_o, P(z, x_o) > v)$ for individuals with $D_z = 1$ and $\Delta^{PeT}(x_o, p, D) = E(Y_1 - Y_0 | X_O = x_o, P(z, x_o) < v)$ for individuals with $D_z = 0$.

All of the treatment effect estimands, including ATE and CATEs, can be derived by appropriately aggregating the PeT effects since these are defined at the individual level (see Section 3.1.3).

3.1.3 | Estimation methods

Two-stage least squares estimator

2SLS is a common approach to the implementation of IV methods that consistently estimates the ATE parameter under homogeneity, or the LATE parameter under essential heterogeneity given a binary IV. Under assumptions (A1)-(A6), the 2SLS (Wald) estimator involves: (i) estimating $E[D_Z | X_O, Z]$ by regressing D_z on X_O and Z , and (ii) estimating $E[Y_D | D_z, X_O, Z]$ by regressing on X_O and $\hat{E}[D_Z | X_O, Z]$. When the instrument is continuous, 2SLS reports a weighted average of LATEs, which requires careful interpretation (Baiocchi et al., 2014).

Local Instrumental Variables estimator: Estimating PeT effects

Basu (2014, 2015) describe in detail the series of steps required to estimate PeT effects using the LIV methodology. Briefly, D_z is regressed on Z and X_O , as above, using appropriate methods for binary outcomes and the propensity for treatment $p(x_O, z)$ is estimated. Next, Y is regressed on X_O and a function of $\hat{p}(x_O, z)$ including interactions with X_O . The approach outlined in Basu (2014) involves differentiating the outcome model $g(Y)$ by $\hat{p}(x_O, z)$. Next, PeT effects for each individual can be obtained by performing numerical integration, with MTE ($\partial \hat{g}(Y)/\partial \hat{p}$) evaluated by replacing \hat{p} using 1000 random draws of $u \sim \text{unif}(\min(\hat{p}(x_O, z)), \max(\hat{p}(x_O, z)))$. Then, $D^* = \Phi^{-1}\{\hat{p}(x_O, z)\} + \Phi^{-1}(1 - u)$ can be computed. Personalized treatment effects can be computed by averaging $\partial \hat{g}(Y)/\partial \hat{p}$ over values of u for which $D^* > 0$ if $D = 1$; or over values of $D^* \leq 0$ if $D = 0$. Finally, averaging PeT effects over all of the observations provides an estimate of the ATE for the population, and over strata of X_O gives the CATE for the subpopulation of interest. Standard errors can be computed using bootstrap methods (Basu, 2015). We now consider the design of the simulation study to contrast the relative performance of the LIV and 2SLS approaches.

3.2 | Simulation study

Motivated by the gaps in the extant literature, and the motivating example, this simulation study was designed to consider the relative performance of 2SLS and LIV approaches across settings that differed with respect to the form of heterogeneity, the sample size and the strength of the IV. We report the performance of the methods in a Monte Carlo Simulation study according to their mean bias (%) and Root Mean Squared Error (RMSE) for each estimand (ATE and CATE).

3.2.1 | Data generating process

We create 5000 datasets each containing $N = \{5000, 10000, 50000\}$ units, of which 50% are assigned to the treated group. The data generating process (DGP) includes one observed (X_O) and one unmeasured (X_U) covariate. We draw X_O , X_U and the instrument, Z from normal distributions with mean 0, and standard deviation 3. Three subgroups of interest are defined by whether the individuals' values for X_O are more than 0.5 standard deviations below or above its mean.

Treatment model

The treatment assignment is determined by the latent variable D^* , defined as:

$$D^* = \delta_D + 3X_O - 3X_U + \delta_Z Z + (4 - \delta_Z)\epsilon_D$$

where ϵ_D has a normal distribution with mean 0 and standard deviation, 1. Treatment is then determined as $D = 1$ if $D^* > 0$ and $D = 0$ otherwise. The parameters δ_Z and δ_D are chosen to ensure the average F-statistic, F_{IV} , across the datasets equals the desired level $F_{Target} = \{10, 25, 50, 100, 500, 1000\}$. F_{IV} is the Cragg and Donald (1993) F-statistic computed in each dataset as,

$$F_{IV} = (N - df_m - 1) * \frac{\sigma_{noIV}^2 - \sigma_{IV}^2}{\sigma_{IV}^2}$$

where σ_{noIV}^2 and σ_{IV}^2 indicate the residual variance from regressing D on X_O in a model without interactions with or without including the IV respectively, and df_m is the number of parameters in the model excluding the IV (i.e., $df_m = 2$ here). For a given F-statistic, a larger sample size implies a lower compliance rate, which in turn will imply a weaker instrument. At low compliance rates, the RMSE of IV estimates can increase substantially (Little et al., 2009). We estimate the compliance rate for each sample size and F-statistic, by contrasting treatment uptake at the 1st and 99th percentiles of the IV.

Outcome model

The outcome models under treatments (Y_1) and control (Y_0) can be written as:

$$Y_0 = \beta_0 + \beta_1 X_O + \beta_2 X_U + \epsilon_{Y_0}$$

$$Y_1 = (\beta_0 + \tau_0) + (\beta_1 + \tau_1) X_O + (\beta_2 + \tau_2) X_U + \epsilon_{Y_1}$$

Implying the treatment effect is $\tau = E(Y_1 - Y_0) = \tau_0 + \tau_1 X_O + \tau_2 X_U$. Specifically we define the outcome under control as follows:

$$Y_0 = -10 - 10X_O + 10X_U + N(0, 1)$$

TABLE 1 Definition of the simulation scenarios.

Scenario	Sample size	Target F-statistic	τ_0	τ_1	τ_2
Scenario A: Homogeneity	All sample sizes	All target F-statistic values	50	0	0
Scenario B: Overt heterogeneity	$(n = \{5000, 10000, 50000\})$	$(F_{Target} = \{10, 25, 50, 100, 500, 1000\})$	40	20	0
Scenario C: Essential heterogeneity			40	0	20
Scenario D: Overt and essential heterogeneity			20	20	20

Note: For each particular scenario, across all sample sizes and target F-statistic values, the form of treatment effect heterogeneity is defined by the values of τ_0 , τ_1 and τ_2 .

We consider 4 scenarios for the outcome under treatment, Y_1 . In Scenario A, effects are homogeneous ($\tau = 50$). In Scenario B, effects are heterogeneous but depend only on observed confounders (overt heterogeneity) ($\tau = 40 + 20X_O$). In Scenario C, X_U influences both the treatment assignment and the gains from treatment ($\tau = 40 + 20X_U$). In this Scenario, there is essential heterogeneity but no overt heterogeneity. Finally, in Scenario D there is both overt and essential heterogeneity ($\tau = 20 + 20X_O + 20X_U$). Table 1 displays the parameter values for each scenario. The parameter combinations of interest consist of combinations of $n = \{5000, 10000, 50000\}$ and $F_{Target} = \{10, 25, 50, 100, 500, 1000\}$.

For each parameter combination for each scenario, we create 5000 datasets using the DGP described above and estimate the treatment effects as described below.

3.2.2 | Implementation of methods

For the 2SLS model, we control for X_O and instrument D by Z . To capture heterogeneity, we also include an interaction between X_O with D , and instrument this with interactions of Z and X_O . To obtain effect estimates, we use the recycled predictions approach, whereby the two potential outcomes (Y_0 and Y_1) are predicted from the second stage model after setting $D = 0$ or $D = 1$ and the interaction $X_O * D = 0$ or X_O (Basu & Rathouz, 2005). The individual level effect is then estimated as $\hat{\tau} = \hat{Y}_1 - \hat{Y}_0$, allowing us to calculate the ATE, and CATEs for the three subgroups (CATE₁, CATE₂, and CATE₃).

For the LIV approach, we first estimate the propensity for treatment conditional on X_O and Z , and in the second stage outcome model we include X_O , D , the estimated propensity score, \hat{p} , $\hat{p} * X_O$ and \hat{p}^2 . We then estimate PeT effects for each individual as described in Basu (2015) using the petiv command in Stata. The estimated PeT effects are then aggregated to obtain estimates of the ATE, CATE₁, CATE₂, and CATE₃. Before applying either method, we remove observations at those levels of the estimated propensity score where there is insufficient overlap (Basu, 2015).

4 | RESULTS

4.1 | Simulation study

Figures 1–4 present mean (%) bias in the ATE and CATE estimates (Figures 1 and 2, respectively) and the corresponding plots for RMSE (Figures 3 and 4, respectively). The results for the three subgroups showed similar patterns, and hence, for brevity, we only report the results for one of them.

In settings with homogenous treatment effects, or with overt heterogeneity, levels of bias in the ATE estimates were generally low (<5%). When the F-statistic was below 50 or the sample size was smaller ($n = 5000$), the bias for the 2SLS estimates was somewhat higher (5%–10%) (Figure 1). In settings with essential heterogeneity, 2SLS reports relatively high (>10%) levels of mean bias across almost all combinations of IV strength and sample size. The levels of mean bias are only similar between the methods when the target F statistic is high (>100). For 2SLS, the confidence intervals (CI) around the estimates of mean bias are generally wide. For 2SLS to provide estimates with moderate to small levels of bias, with narrow CI around those estimates required an F statistic of at least 100 and a sample size of 50,000. The LIV estimator reports low levels of bias in ATE estimates across all scenarios, aside from those with both a smaller sample size ($n = 5000$) and a F-statistic of 25 or less (Figure 1).

The bias plots for the CATE estimates have a somewhat similar pattern, although for this estimand the 2SLS estimator reports high levels of mean bias even in settings with overt heterogeneity, unless the sample size is relatively large ($n = 50,000$)

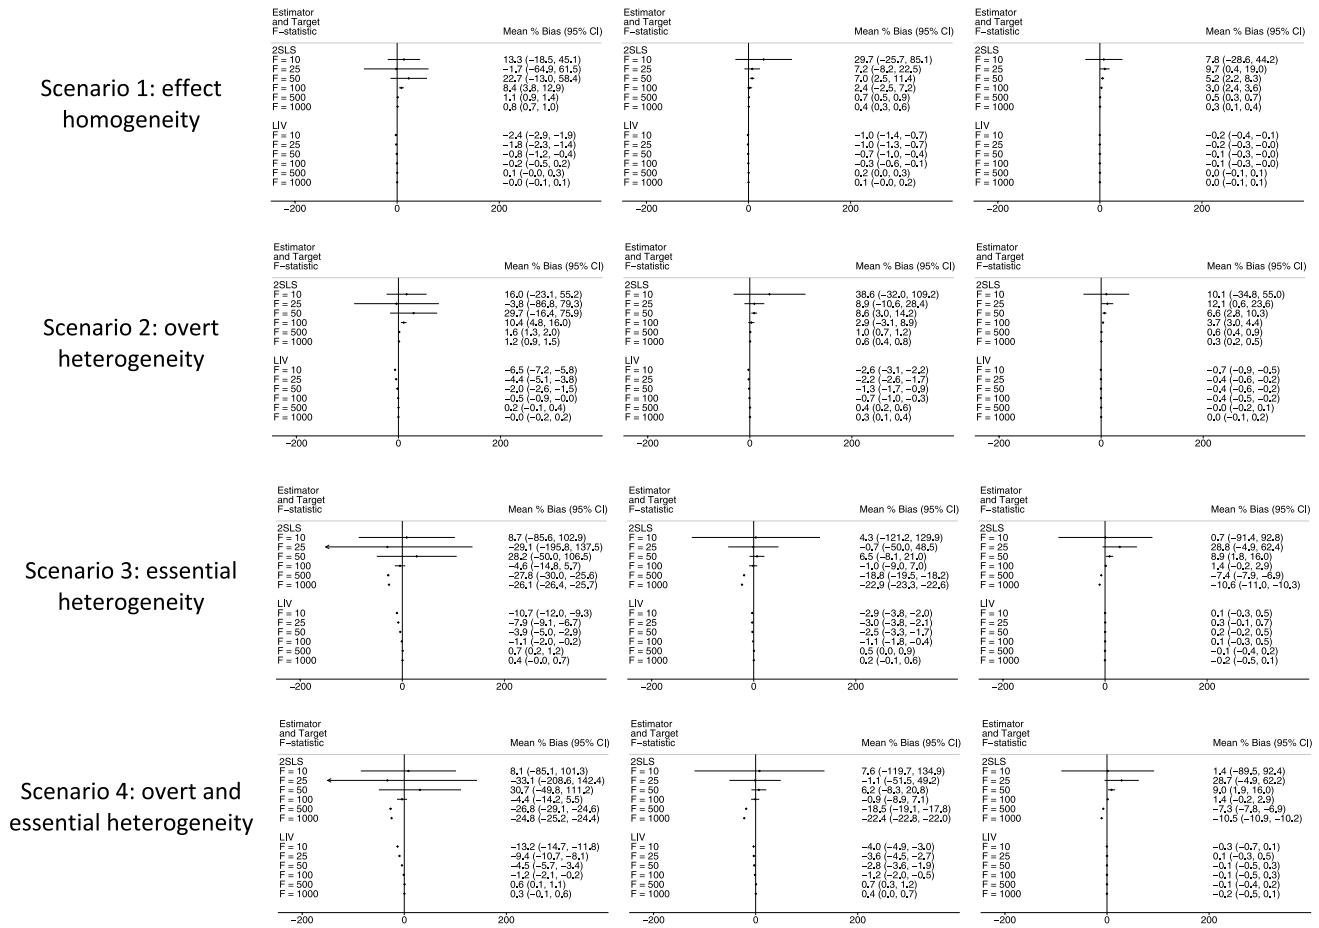


FIGURE 1 Bias plot for Average Treatment Effect (ATE) estimates across scenarios, with sample sizes of 5000 (left), 10000 (middle) and 50,000 (right). 2SLS, Two-Stage Least Squares; CI, Confidence Interval; LIV, Local Instrumental Variables.

and/or the F-statistic is above 100 (Figure 2). The LIV estimator reports lower levels of bias than 2SLS across the majority of scenarios.

In general, for both methods, across most scenarios, for a given sample size, the levels of mean (%) bias decrease at higher levels of the F-statistic (Figure 2). The RMSE in the estimates of the ATE are substantially lower for the LIV than the 2SLS estimator, except for those settings with an F-statistic of 500 or 1000 (Figure 3).¹ For the CATE, in general, the RMSE estimates mirror the bias results, in that they are substantially lower across all settings for LIV (Figure 4).

Compliance rates for a given F-statistic were sensitive to the sample size available (see Table 2 below). For a sample size of 5000, increasing the F-statistic from 10 to 1000 increases the compliance rate from 8% to 73%, while for a sample size of 50,000, the compliance rate only increases from 3% to 29%.

4.2 | Case study

4.2.1 | Case study: Implementation of 2SLS and LIV approaches

LIV estimated PeT effects of ES versus NES on DAOH at 90 days, for each individual allowing for treatment effect heterogeneity and confounding. These PeT effects were aggregated to report the effects of ES overall, and for each pre-specified subgroup of interest. Since DAOH at 90 days was left skewed due to the maximum being 90 days, we rescaled this to lie between 0 and 1 (90-DAOH)/90) and effects were then rescaled back to the original scale. Probit regression models were used to estimate the initial propensity score (first stage), while GLMs were applied to the outcome data, with the most appropriate family and link function chosen according to RMSE, with Hosmer-Lemeshow and Pregibon tests also used to check model fit and

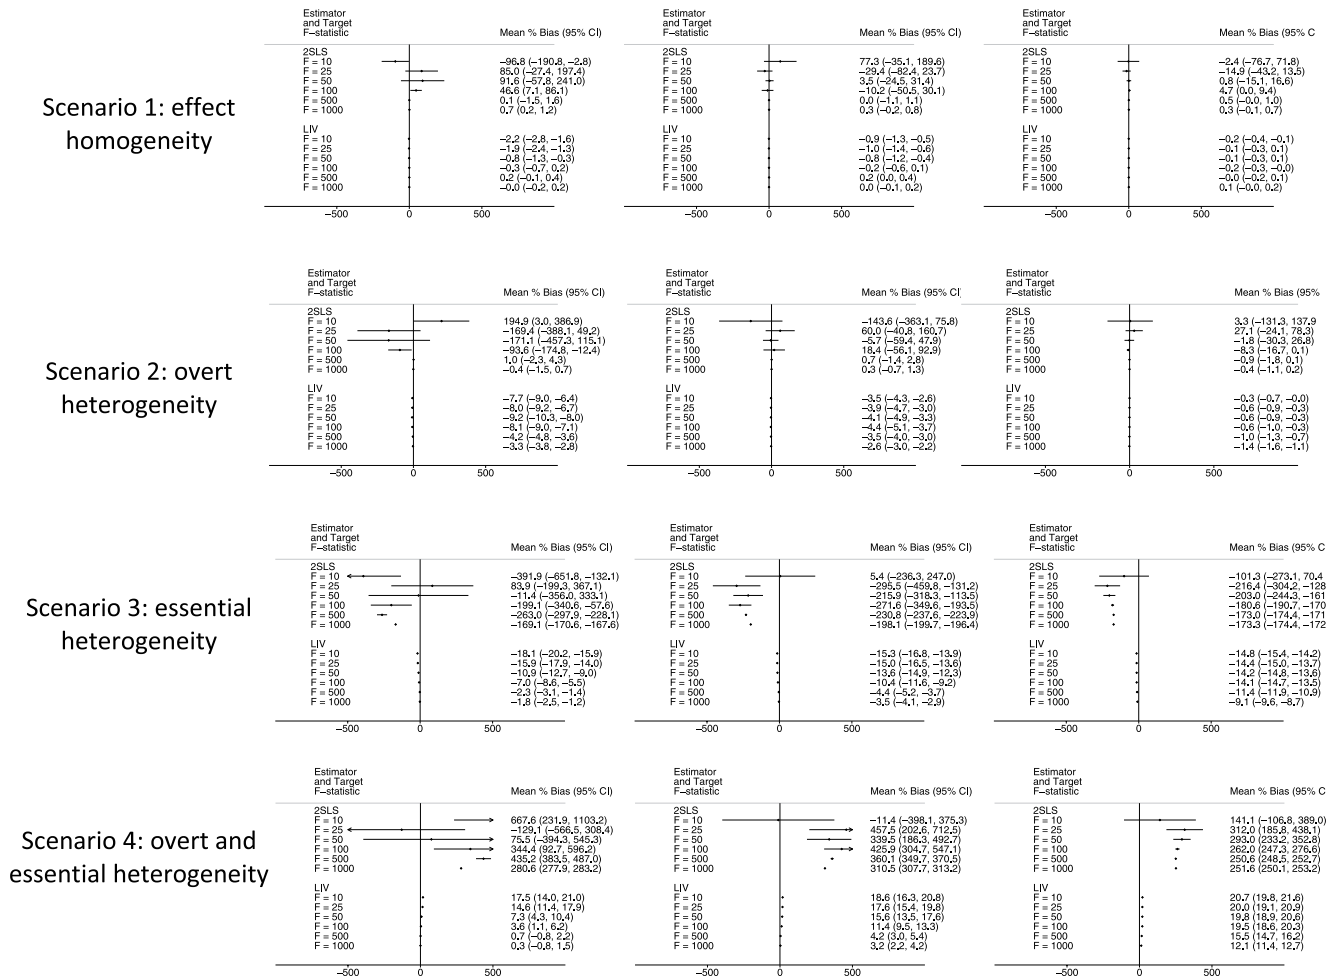


FIGURE 2 Bias plot for Conditional Average Treatment Effect (CATE) estimates across scenarios, with sample sizes of 5000 (left), 10,000 (middle) and 50,000 (right). 2SLS, Two-Stage Least Squares; CI, Confidence Interval; LIV, Local Instrumental Variables.

appropriateness (Hosmer & Lemeshow, 2000; Pregibon, 1980). The logit link and binomial family were selected for all three conditions. Models at both stages adjusted for baseline measures, time period, and proxies for hospital quality, defined by rates of emergency readmission and mortality in 2009–10 (time constant), and in the year prior to the specific admission concerned (time-varying).

Estimates of mean differences in DAOH between the comparison groups, overall and for pre-specified subgroups (CATEs) were reported with standard errors and CI obtained with the non-parametric bootstrap (300 replications), allowing for the clustering of individuals within hospitals. The 2SLS approach used the same model specification and selection (including covariates used for confounding adjustment) to report estimates overall and for subgroups.

4.2.2 | Case study: Results

The study reported somewhat similar that for both methods the 95% CIs surrounding the mean differences included zero (Figure 5). Beneath this overall result, the LIV approach reported evidence that the effectiveness of ES was heterogeneous according to pre-specified subgroups. In particular, for all three conditions, ES led to lower DAOH for patients who had severe levels of frailty, and for those with acute appendicitis, ES was less effective for older patients (aged 80–84) or those with three or more comorbidities. By contrast, the 2SLS approach, which failed to account for unobserved heterogeneity (e.g., disease severity), did not report any substantive differences in relative effectiveness according to patient subgroup (Figure 5).

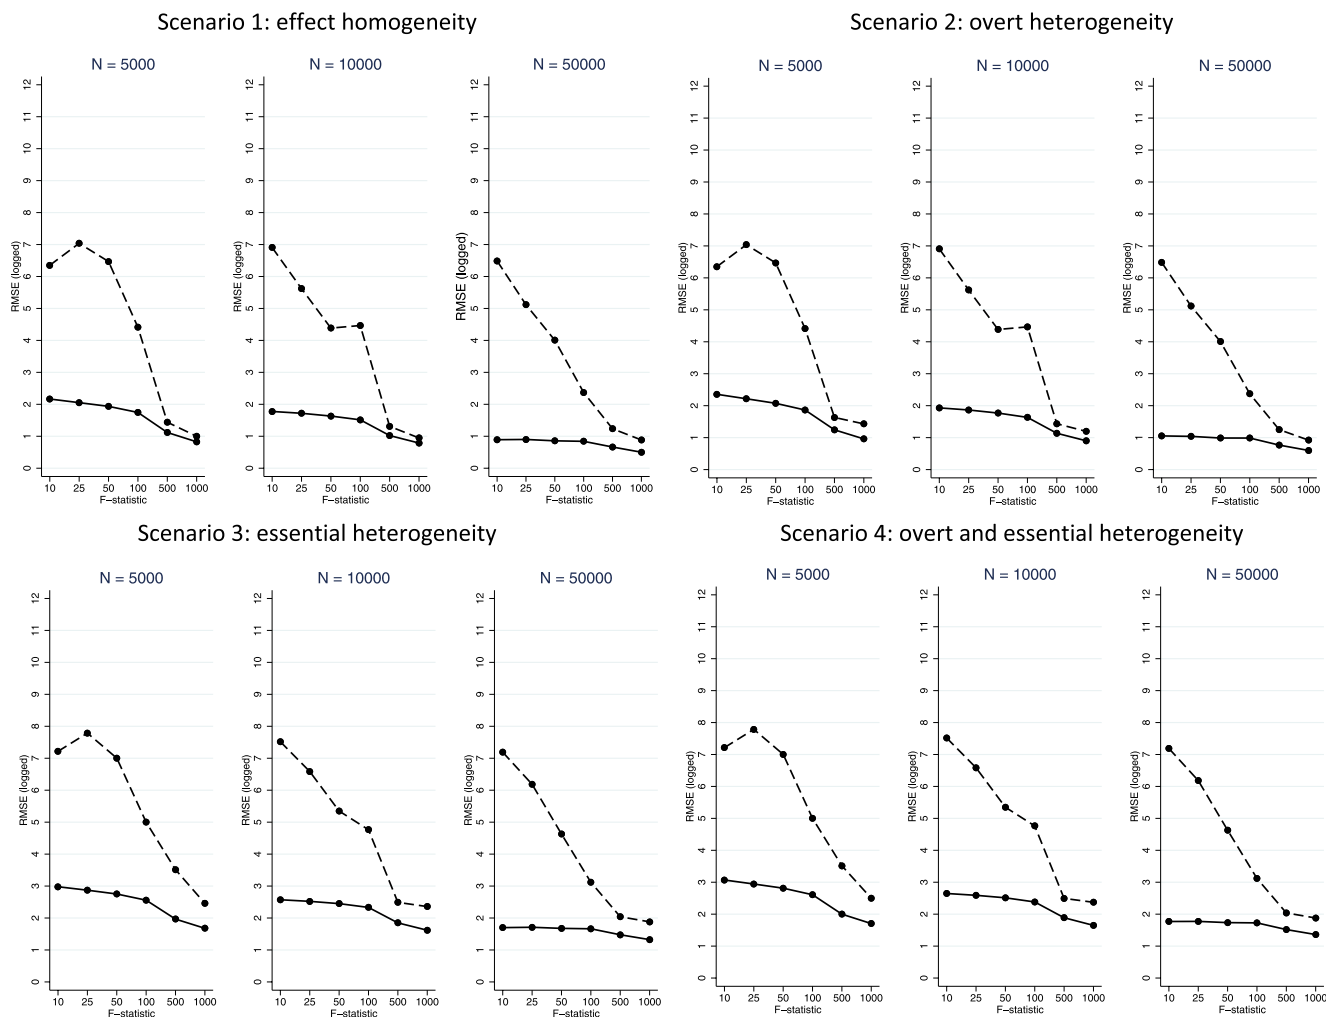


FIGURE 3 Root Mean Squared Error (RMSE) plots for Average Treatment Effect (ATE) estimates from 2SLS (dashed line) and LIV (solid line) across the scenarios, with sample sizes (N) of 5000 (left), 10,000 (middle) and 50,000 (right). [Colour figure can be viewed at wileyonlinelibrary.com]

5 | DISCUSSION

This paper formally assessed the performance of the LIV methodology developed by Heckman and Vytlačil (1999, 2001) and further extended by Basu (2014) to provide policy relevant estimates of ATE and CATE in settings that differed according to the form of heterogeneity, the sample size, and level of IV strength. We contrasted the performance of LIV with that of the widely-used 2SLS approach. The scenarios considered in the simulation study were directly motivated by gaps in the literature and by a comparative effectiveness study that used LIV in evaluating ES for three acute gastrointestinal conditions for subgroups of prime policy relevance. In the case study, overt and essential heterogeneity were important concerns, amid differing levels of IV strength and sample sizes, and these issues motivated the scenario of prime interest for the simulation study (Scenario D). However, we also considered scenarios, which can, in principle provide accurate estimates of ATE and CATEs with conventional IV methods such as 2SLS (Scenarios A and B). We compared the performance of the two methods, according to bias and statistical efficiency (RMSE).

Four preliminary findings of the simulation study are worth emphasizing. First, our results suggest that while LIV performs better according to increasing levels of IV strength and sample size, this estimator reports relatively low levels of bias in estimates of the ATE and CATEs across all scenarios including those with essential heterogeneity. These findings compliment those of Basu (2014) in evaluating the reliance of the estimator on the relevance condition as well as the consistency of the estimator, but also by considering a wider range of assumptions about heterogeneity.

Second, our results suggest that 2SLS reports biased estimates of the ATE and CATEs in the presence of essential heterogeneity (Scenarios C and D), except in those cases where the instrument is very strong (F -statistic > 500), and sample size is

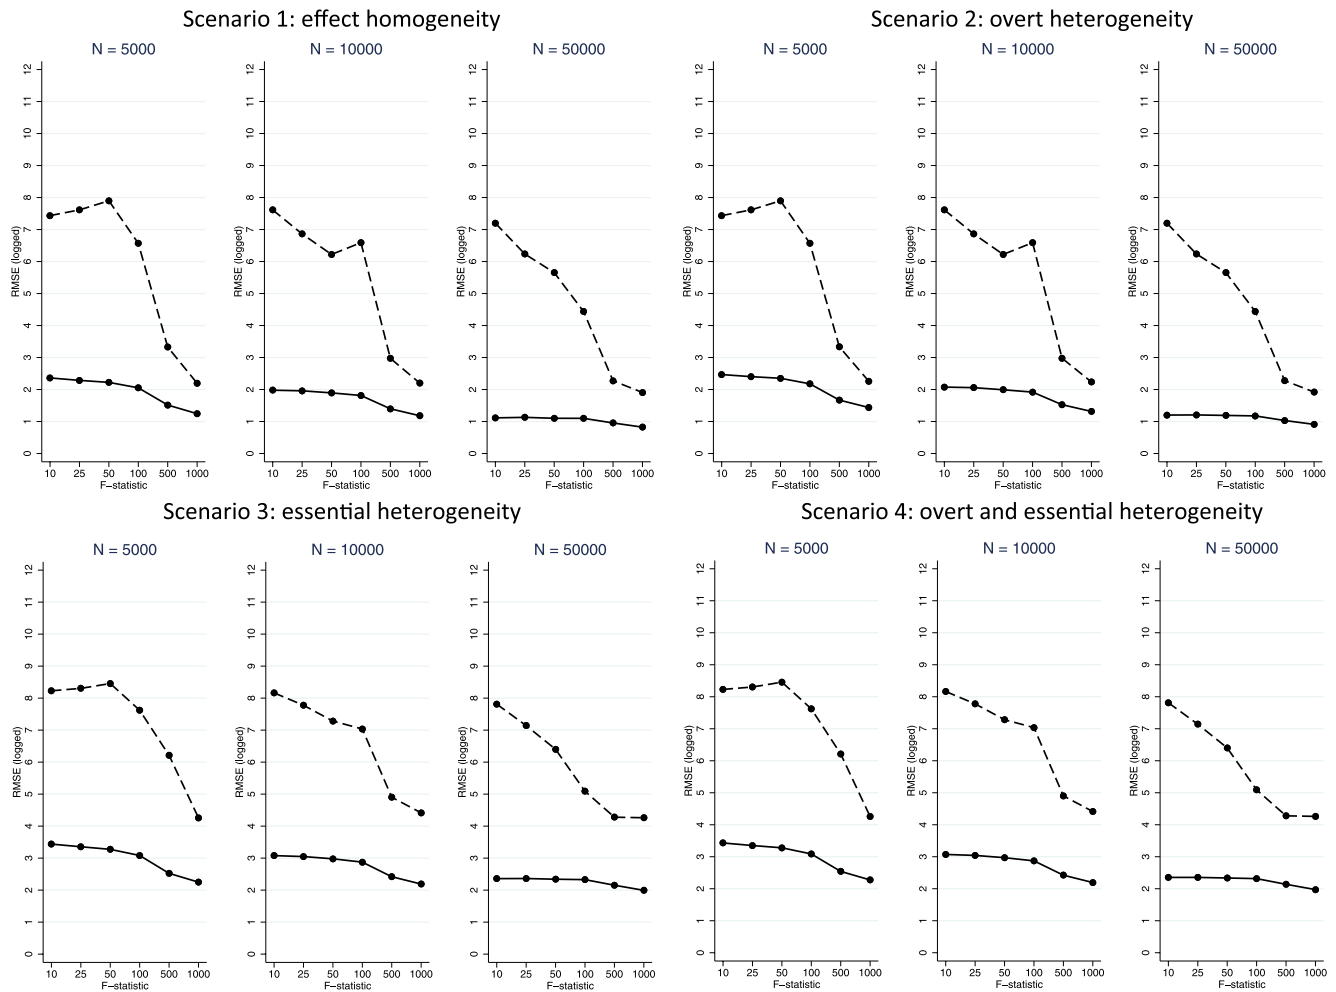


FIGURE 4 Root Mean Squared Error (RMSE) plots for Conditional Average Treatment Effect (CATE) estimates from 2SLS (dashed line) and LIV (solid line) across the scenarios, with sample sizes (N) of 5000 (left), 10,000 (middle) and 50,000 (right). [Colour figure can be viewed at wileyonlinelibrary.com]

Target F-statistic	$N = 5000$	$N = 10,000$	$N = 50,000$
10	8%	6%	3%
25	13%	9%	5%
50	18%	13%	6%
100	26%	20%	9%
500	56%	42%	21%
1000	73%	57%	29%

TABLE 2 Compliance rate by sample size (N) and F-statistic.

Note: Each cell shows the compliance rate for each value of the Target F-statistic calculated as the difference in treatment probability between highest and lowest quintile of the instrument, Z .

large ($n = 50,000$). These results are consistent with previous findings that 2SLS estimates cannot generally be extrapolated to broader populations beyond the compliers unless restrictive assumptions are made about the heterogeneity of treatment effects (Brooks et al., 2018; Chapman & Brooks, 2016). In scenarios where essential heterogeneity is absent (scenarios A and B), 2SLS reports biases larger than 5% at levels of $F < 50$, unless the sample size is very large ($n = 50,000$). It should be noted that at those values of F the distribution of the mean bias is quite skewed, and this is reflected in the wide CIs around the estimates of mean bias (see Figures 1 and 2). These findings suggest that, in studies without essential heterogeneity, and with large samples and a sufficient strong IV, 2SLS is a simple alternative to LIV.

As the sample size increases the magnitude of the bias is reduced for both methods, but for 2SLS an F statistic of at least 50 (depending on the sample size) is required for estimates of mean bias and the accompanying CI to be $<5\%$. This finding




FIGURE 5 Mean differences in days alive and out of hospital (DAOH) between emergency surgery (ES) and non-emergency surgery for appendicitis (left), gallstone disease (center) and hernia (right) subgroups. 2SLS, two-stage least squares; CI, Confidence Interval; LIV, Local Instrumental variables.

lends support to existing guidance suggesting that the requisite magnitude of the F statistic depends on other factors such as the sample size (Hirano & Porter, 2015; Keane & Neal, 2023). This finding further emphasizes the inadequacy of guidance resting solely on a ‘rule of thumb’ for a single setting, the target F-statistic, and highlights the importance of these wider considerations of the likely form of heterogeneity, and sample size, as well as the F statistic when interpreting a study’s results.

Thirdly, while 2SLS can reliably estimate CATEs in the presence of effect homogeneity or overt heterogeneity given a sufficiently strong IV or large enough sample, in the presence of essential heterogeneity, as theory would suggest, 2SLS can give extremely biased estimates of CATEs, and so in settings where essential heterogeneity is anticipated, 2SLS should not be used to estimate CATEs. In contrast, the LIV method provided estimates with low bias in the presence of overt and/or essential heterogeneity, provided the F-statistic was greater than 50. Interestingly, for the estimates of the CATEs, we find that as the sample size increases, an increase in the F-statistic is less beneficial in mitigating bias and reducing RMSE, in line with the observation that a given increase in the F-statistic has less impact on compliance rates at larger sample sizes.

Finally, LIV generally reported lower levels of RMSE than 2SLS, in particular for estimating the CATEs. However, it is important to note that here the propensity score and outcome models underlying the LIV method are correctly specified, and that performance may deteriorate where this is not the case. Data adaptive approaches could prove useful where model specification is not known.

The findings from the simulation study are informative in interpreting the CATE estimates in the ESORT study. The results offer reassurance that in such settings where essential heterogeneity would appear inevitable, that a LIV approach can provide unbiased estimate of policy-relevant estimands such as CATE, with sample sizes and F-statistics smaller than those of the ESORT study. Here, the LIV approach was able to report relative effectiveness according to subgroup, and the finding that for patients with high levels of frailty ES was not cost-effective (or cost-effective) is of potential importance. Notwithstanding inevitable concerns about multiple testing, and it should be noted that no formal adjustment for this was made, it is of potential interest for policy, that the finding that the intervention was not cost-effective for the subgroup with severe frailty was replicated across all five conditions in the original study. Further research to test this hypotheses in different settings is now warranted (Moler-Zapata et al., 2022).

This study has several strengths. First, it builds on insights and hypotheses raised by a large observational study using EHRs from England. The ESORT study illustrates the main challenges of using LIV methods for comparative effectiveness research

and its findings in relation to IV strength, sample size requirements directly informed the scenarios considered in the simulation study. Second, while the uptake of LIV methods has been limited almost entirely to settings with essential heterogeneity, the simulation study considers different forms of heterogeneity of treatment effects as well as the scenario where treatment effects are assumed to be homogeneous in the study population. Future work will expand the simulation study to incorporate other well-known issues of IVs methods, including the challenges in applying IV estimation methods to non-linear outcome data (Clarke & Windmeijer, 2010; Vansteelandt et al., 2011). Previous research has shown that the power of 2SLS conveyed by conventional F-statistic values is low (Keane & Neal, 2023; Y. Lee et al., 2020). In this future work, we will therefore consider the implications of sample size and instrument strength for the power of LIV analyses and confidence interval coverage. Future work will also formally assess whether imbalances in treatment assignment rates are detrimental to consistency and power of LIV inferences. This is an important concern for applied work using EHRs. For instance, the observed difference in the prevalence of ES and NES in ESORT (90/10 in the cohort with appendicitis) could reduce the power of the analysis (Walker et al., 2017). Finally, the current implementation of LIV requires a continuous IV which may not be available in some settings, and remains unknown the extent to which the method may produce biased or inefficient estimates if the underlying assumptions with respect to the treatment assignment or outcome models are violated. Further research can build on the simulation study described to assess these potentially important issues for practice.

ACKNOWLEDGMENTS

We would like to acknowledge Professor Luke Keele for helpful discussions, and other members of the ESORT study team. We would like to acknowledge the feedback from those who participated in the sessions, and particularly the paper discussants at the Summer 2022 Health Economics Study Group conference and 29th European Workshop On Econometrics and Health Economics, Dr Manuel Gomes and Dr Eugenio Zucchelli. This report is independent research supported by the National Institute for Health Research ARC North Thames. The views expressed in this publication are those of the author(s) and not necessarily those of the National Institute for Health Research or the Department of Health and Social Care.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from NHS Digital but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available.

ETHICS STATEMENT

The research was approved by the London School of Hygiene and Tropical Medicine ethics committee (Ethics Reference no: 21687).

ORCID

Silvia Moler-Zapata  <https://orcid.org/0000-0003-4733-5601>

Richard Grieve  <https://orcid.org/0000-0001-8899-1301>

Anirban Basu  <https://orcid.org/0000-0003-4238-7402>

ENDNOTE

¹ For example, in scenario C, for a level of IV strength when $N = 5000$ and $F = 10$, the logged RMSE for the ATE is 2.98 for LIV and 7.91 for 2SLS. These findings correspond to values of a normalised RMSE, which relates the (unlogged) RMSE to the true ATE, of 0.49 and 67.96 for LIV and 2SLS, respectively.

REFERENCES

- Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics*, 113(2), 231–263. [https://doi.org/10.1016/S0304-4076\(02\)00201-4](https://doi.org/10.1016/S0304-4076(02)00201-4)
- Andrews, I., Stock, J. H., & Sun, L. (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics*, 11(1), 727–753. <https://doi.org/10.1146/annurev-economics-080218-025643>
- Angrist, J., Imbens, G., & Rubin, D. (1993). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434), 444–455. <https://doi.org/10.1080/01621459.1996.10476902>
- Angrist, J. D., & Fernández-Val, I. (2011). Extrapolating: External validity and overidentification in the LATE framework. *Advances in Economics and Econometrics: Tenth World Congress*, 3, 401–434. <https://doi.org/10.1017/CBO9781139060035.012>

- Baiocchi, M., Cheng, J., & Small, D. S. (2014). Instrumental variable methods for causal inference. *Statistics in Medicine*, 33(13), 2297–2340. <https://doi.org/10.1002/sim.6128>
- Basu, A. (2014). Estimating person-centered treatment (PeT) effects using instrumental variables: An application to evaluating prostate cancer treatments. *Journal of Applied Econometrics*, 29(4), 671–691. <https://doi.org/10.1002/jae>
- Basu, A. (2015). Person-centered treatment (PeT) effects: Individualized treatment effects using instrumental variables. *STATA Journal*, 15(2), 397–410. <https://doi.org/10.1177/1536867x1501500204>
- Basu, A., Coe, N. B., & Chapman, C. G. (2018a). 2SLS versus 2SRI: Appropriate methods for rare outcomes and/or rare exposures. *Health Economics*, 27(6), 937–955. <https://doi.org/10.1002/hec.3647>
- Basu, A., Heckman, J. J., Navarro-Lozano, S., Urzúa, S., & Urzua, S. (2007). Use of instrumental variables in the presence of heterogeneity and self-selection: An application to treatments of breast cancer patients. *Health Economics*, 16(2007), 1133–1157. <https://doi.org/10.1002/hec.1291>
- Basu, A., Jones, A. M., & Rosa Dias, P. (2018b). Heterogeneity in the impact of type of schooling on adult health and lifestyle. *Journal of Health Economics*, 57, 1–14. <https://doi.org/10.1016/j.jhealeco.2017.10.007>
- Basu, A., & Rathouz, P. J. (2005). Estimating marginal and incremental effects on health outcomes using flexible link and variance function models. *Biostatistics*, 6(1), 93–109. <https://doi.org/10.1093/biostatistics/kxh020>
- Bjorklund, A., & Moffitt, R. (1983). *Estimation of wage gains and welfare gains from self-selection models*. IUI Working Paper, No. 105.
- Bound, J., Jaeger, D. A., & Baker, R. M. (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association*, 90(430), 443–450. <https://doi.org/10.1080/01621459.1995.10476536>
- Brooks, J. M., Chapman, C. G., & Schroeder, M. C. (2018). Understanding treatment effect estimates when treatment effects are heterogeneous for more than one outcome. *Applied Health Economics and Health Policy*, 16(3), 381–393. <https://doi.org/10.1007/s40258-018-0380-z>
- Chapman, C. G., & Brooks, J. M. (2016). Treatment effect estimation using nonlinear two-stage instrumental variable estimators: Another cautionary note. *Health Services Research*, 51(6), 2375–2394. <https://doi.org/10.1111/1475-6773.12463>
- Clarke, P., & Windmeijer, F. (2010). Instrumental variable estimators for binary outcomes. CMPO Working Paper Series No. 10/239 Instrumental Retrieved from <http://www.bristol.ac.uk/cmppo/Tel>
- Cornelissen, T., Dustmann, C., Raute, A., & Schönberg, U. (2018). Who benefits from universal child care? Estimating marginal returns to early child care attendance. *Journal of Political Economy*, 126(6), 2356–2409. <https://doi.org/10.1086/699979>
- Cragg, J. G., & Donald, S. G. (1993). Testing identifiability and specification in instrumental variable models. *Econometric Theory*, 9(2), 222–240. <https://doi.org/10.1017/s0266466600007519>
- ESORT Study Group. (2020). *Emergency surgery or NoT (ESORT) study*. Study Protocol. Retrieved from <https://www.lshtm.ac.uk/media/38711>
- Grieve, R., O'Neill, S., Basu, A., Keele, L., Rowan, K. M., & Harris, S. (2019). Analysis of benefit of intensive care unit transfer for deteriorating ward patients: A patient-centered approach to clinical evaluation. *JAMA Network Open*, 2(2), 1–13. <https://doi.org/10.1001/jamanetworkopen.2018.7704>
- Heckman, J. J., Urzua, S., & Vytlacil, E. (2006). Understanding instrumental variables in models with essential heterogeneity. In *NBER working paper No. 12574* (Vol. 88). <https://doi.org/10.1162/rest.88.3.389.3>
- Heckman, J. J., & Vytlacil, E. (2005). Structural equations, treatment effects and econometric policy evaluation. *Econometrica*, 73(3), 669–738. <https://doi.org/10.1111/j.1468-0262.2005.00594.x>
- Heckman, J. J., & Vytlacil, E. J. (1999). Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences of the United States of America*, 96(8), 4730–4734. <https://doi.org/10.1073/pnas.96.8.4730>
- Heckman, J. J., & Vytlacil, E. J. (2001). Policy-relevant treatment effects. *The American Economic Review*, 91(2), 107–111. <https://doi.org/10.1257/aer.91.2.107>
- Hirano, K., & Porter, J. R. (2015). Location properties of point estimators in linear instrumental variables and related models. *Econometric Reviews*, 34(6–10), 720–733. <https://doi.org/10.1080/07474938.2014.956573>
- Hosmer, D. W., & Lemeshow, S. (2000). *Applied logistic regression* (2nd ed.). Wiley.
- Hutchings, A., Moler-Zapata, S., O'Neill, S., Smart, N., Hinchliffe, R., Cromwell, D., & Grieve, R. (2021). Variation in the rates of emergency surgery amongst emergency admissions to hospital for common acute conditions. *BJS Open*, 5(6). <https://doi.org/10.1093/bjsopen/zrab094>
- Hutchings, A., O'Neill, S., Lugo-Palacios, D., Moler Zapata, S., Silverwood, R., Cromwell, D., Keele, L., Bellingan, G., Moonesinghe, S. R., Smart, N., Hinchliffe, R., & Grieve, R. (2022). Effectiveness of emergency surgery for five common acute conditions: An instrumental variable analysis of a national routine database. *Anaesthesia*, 77(8), 865–881. <https://doi.org/10.1111/anae.15730>
- Imbens, G. W., & Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2), 467. <https://doi.org/10.2307/2951620>
- Keane, M., & Neal, T. (2023). Instrument strength in IV estimation and inference: A guide to theory and practice. *Journal of Econometrics*, xxxx. <https://doi.org/10.1016/j.jeconom.2022.12.009>
- Keele, L., Sharoky, C. E., Sellers, M. M., Wirtalla, C. J., & Kelz, R. R. (2018). An instrumental variables design for the effect of emergency general surgery. *Epidemiologic Methods*, 7(1). <https://doi.org/10.1515/em-2017-0012>
- Lee, D., McCrary, J., Moreira, M. J., & Porter, J. R. (2021). *Valid T-ratio inference for IV*. National Bureau of Economic Research Working Paper Series. No. W29124. <https://doi.org/10.2139/ssrn.3901588>
- Lee, Y., Kennedy, E. H., & Mitra, N. (2020). Doubly robust nonparametric instrumental variable estimators for survival outcomes.
- Little, R. J., Long, Q., & Lin, X. (2009). A comparison of methods for estimating the causal effect of a treatment in randomized clinical trials subject to noncompliance. *Biometrics*, 65(2), 640–649. <https://doi.org/10.1111/j.1541-0420.2008.01066.x>

- Martínez-Cambor, P., MacKenzie, T. A., Staiger, D. O., Goodney, P. P., & James O'Malley, A. (2019). An instrumental variable procedure for estimating Cox models with non-proportional hazards in the presence of unmeasured confounding. *Journal of the Royal Statistical Society - Series C: Applied Statistics*, 68(4), 985–1005. <https://doi.org/10.1111/rssc.12341>
- Moffitt, R. A., & Zahn, M. V. (2022). The marginal labor supply disincentives of welfare: Evidence from administrative barriers to participation.
- Moler-Zapata, S., Grieve, R., Lugo-Palacios, D., Hutchings, A., Silverwood, R., Keele, L., Kircheis, T., Cromwell, D., Smart, N., Hinchliffe, R., & O'Neill, S. (2022). Local instrumental variable methods to address confounding and heterogeneity when using electronic health records: An application to emergency surgery. *Medical Decision Making*, 0(0), 1010–1026. <https://doi.org/10.1177/0272989X221100799>
- Nelson, C. R., & Startz, R. (1990). Some further results on the exact small sample properties of the instrumental variables estimator. *Econometrica*, 58(4), 967–976. <https://doi.org/10.2307/2938359>
- Neyman, J. (1990). On the application of probability theory to agricultural experiments. *Statistical Science*, 5, 463–480.
- Pregibon, D. (1980). Goodness of link tests for generalized linear models. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 29(1), 14–15. <https://doi.org/10.2307/2346405>
- Reynolds, K., Barton, L. J., Basu, A., Fischer, H., Arterburn, D. E., Barthold, D., Courcoulas, A., Crawford, C. L., Kim, B. B., Fedorka, P. N., Mun, E. C., Murali, S. B., Zane, R. E., & Coleman, K. J. (2021). Comparative effectiveness of gastric bypass and vertical sleeve gastrectomy for hypertension remission and relapse: The ENGAGE CVD study. *Hypertension*, 78(4), 1116–1125. <https://doi.org/10.1161/HYPERTENSIONAHA.120.16934>
- Rubin, D. B. (1974). Estimating causal effects of treatment in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688–701. <https://doi.org/10.1037/h0037350>. http://www.fsb.muohio.edu/lij14/420_paper_Rubin74.pdf
- Small, D. S., & Rosenbaum, P. R. (2008). War and wages: The strength of instrumental variables and their sensitivity to unobserved biases. *Journal of the American Statistical Association*, 103(483), 924–933. <https://doi.org/10.1198/016214507000001247>
- Staiger, D., & Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3), 557. <https://doi.org/10.2307/2171753>
- Stock, J. H., & Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. W. K. Andrews & J. H. Stock (Eds.), *Identification and inference for econometric models: Essays in honor of Thomas J. Rothenberg* (pp. 80–108). Cambridge University Press.
- Tan, Z. (2006). Regression and weighting methods for causal inference using instrumental variables. *Journal of the American Statistical Association*, 101(476), 1607–1618. <https://doi.org/10.1198/016214505000001366>
- Terza, J. V., Basu, A., & Rathouz, P. J. (2008a). Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling. *Journal of Health Economics*, 27(3), 531–543. <https://doi.org/10.1016/j.jhealeco.2007.09.009>
- Terza, J. V., Bradford, W. D., & Dismuke, C. E. (2008b). The use of linear instrumental variables methods in health services research and health economics: A cautionary note. *Health Services Research*, 43(3), 1102–1120. <https://doi.org/10.1111/j.1475-6773.2007.00807.x>
- Vansteelandt, S., Bowden, J., Babanezhad, M., & Goetghebeur, E. (2011). On instrumental variables estimation of causal odds ratios. *Statistical Science*, 26(3), 403–422. <https://doi.org/10.1214/11-STS360>
- Vytlačil, E. (2002). Independence, monotonicity and latent index models: An equivalence result. *Econometrica*, 70(1), 331–341. <https://doi.org/10.1080/0305006790150101>
- Walker, V. M., Davies, N. M., Windmeijer, F., Burgess, S., & Martin, R. M. (2017). Power calculator for instrumental variable analysis in pharmacoepidemiology. *International Journal of Epidemiology*, 46(5), 1627–1632. <https://doi.org/10.1093/ije/dyx090>

How to cite this article: Moler-Zapata, S., Grieve, R., Basu, A., & O'Neill, S. (2023). How does a local instrumental variable method perform across settings with instruments of differing strengths? A simulation study and an evaluation of emergency surgery. *Health Economics*, 32(9), 2113–2126. <https://doi.org/10.1002/hec.4719>