

NBER WORKING PAPER SERIES

ONE INSTRUMENT TO RULE THEM ALL:
THE BIAS AND COVERAGE OF JUST-ID IV

Joshua Angrist
Michal Kolesár

Working Paper 29417
<http://www.nber.org/papers/w29417>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2021, Revised August 2022

We thank Ahmet Gulek and Luther Yap for expert research assistance. Thanks also go to Tim Armstrong, Isaiah Andrews, Brigham Frandsen, Guido Imbens, Mike Keane, Dave Lee, Whitney Newey, and Steve Pischke for helpful discussions and insightful comments. Kolesár acknowledges support by the Sloan Research Fellowship and by the National Science Foundation Grant SES-22049356. The views expressed here are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Joshua Angrist and Michal Kolesár. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

One Instrument to Rule Them All: The Bias and Coverage of Just-ID IV
Joshua Angrist and Michal Kolesár
NBER Working Paper No. 29417
October 2021, Revised August 2022
JEL No. C21,C26,C31,C36,J08

ABSTRACT

We revisit the finite-sample behavior of just-identified instrumental variables (IV) estimators, arguing that in most microeconomic applications, just-identified IV bias is negligible and the usual inference strategies likely reliable. Three widely-cited applications are used to explain why this is so. We then consider pretesting strategies of the form $t_1 > c$, where t_1 is the first-stage t-statistic, and the first-stage sign is given. Although pervasive in empirical practice, pretesting on the first-stage F-statistic exacerbates bias and distorts inference. We show, however, that median bias is both minimized and roughly halved by setting $c = 0$, that is by screening on the sign of the estimated first stage. This bias reduction is a free lunch: conventional confidence interval coverage is unchanged by screening on the estimated first-stage sign. To the extent that IV analysts sign-screen already, these results strengthen the case for a sanguine view of the finite-sample behavior of just-ID IV.

Joshua Angrist
Department of Economics, E52-436
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
angrist@mit.edu

Michal Kolesár
Department of Economics
278 Julis Romo Rabinowitz Building
Princeton University
Princeton, NJ 08544-1021
mkolesar@princeton.edu

1 Introduction

The heavily over-identified two-stage least squares (2SLS) estimates reported in Angrist (1990) and Angrist and Krueger (1991, AK91) sparked a flood of interest in the finite-sample behavior of instrumental variables (IV) estimates. Kicked off by Bekker (1994) and Bound et al. (1995), attention to bias in 2SLS estimates with many weak instruments has since become a staple of applied microeconomics. The fact that the finite-sample distribution of 2SLS estimates is shifted towards the corresponding ordinary least squares (OLS) probability limit is especially worrying. IV is often motivated by the belief that OLS estimates are compromised by omitted variable bias (OVB). The IV analyst hopes, therefore, that when IV and OLS are close, this signals minimal OVB rather than substantial finite-sample bias in the IV estimates.

The heavily over-identified models that motivated 1990s IV critiques involve many weak instruments, roped together by 2SLS in an effort to estimate a single causal effect with acceptable precision. Strikingly, Bound et al. (1995) show that in the specifications reported in AK91 that interact quarter of birth dummies with 10 year-of-birth and 50 state-of-birth dummies to generate 180 instruments, replacing real quarter of birth dummies with dummies randomly drawn yields 2SLS estimates and standard errors much like those generated by the real thing.¹ But most IV studies (including Angrist (1990) and AK91) report just-identified IV (just-ID IV) estimates computed with a single instrument along with estimates from over-identified models. Our analysis comes in the wake of renewed interest in the finite-sample properties of just-ID IV, an interest reflected in Andrews and Armstrong (2017), Lee et al. (2020), and Keane and Neal (2021), among others.

We argue here that in typical microeconomic applications of just-ID IV, conventional IV estimates and t -tests are compromised little by failures of the usual asymptotic theory. Our analysis builds on the (approximate) finite-sample normality of reduced-form and first-stage estimators (in the argot of classical simultaneous equations models, these are both estimated “reduced forms”). This modeling framework parallels that used in Andrews et al. (2019) and earlier theoretical investigations of weak instrument problems. The normality of reduced-form estimates is justified by conventional asymptotic reasoning, as well as by the local-to-zero asymptotic sequence used in Staiger and Stock (1997) and Stock and Yogo (2005), in which the first stage shrinks to zero at a rate inversely proportional to the square root of the sample size.

Our setup has two free parameters: the correlation between structural and first-stage residuals (henceforth, “endogeneity”) and the population first-stage F -statistic. This fact lends itself to the construction of easily-interpreted rejection contours characterizing conventional second-stage t -tests and confidence interval coverage rates. We see, for example, that for endogeneity less than about 0.76, 95% confidence interval coverage is distorted by no more than 5% for *any* population F . This is explained by the fact that, even as bias increases when the first stage gets weaker, second-stage precision falls. In contrast with the over-identified case, conventional just-ID IV standard errors reflect this, and confidence intervals widen accordingly. This fact keeps interval coverage high unless

¹This “fake instruments” simulation was originally suggested by Alan Krueger. Although not an empirical study, Bekker (1994) is likewise motivated by a heavily over-identified specification in Angrist (1990) that uses 73 draft lottery dummies plus interaction terms as instruments for Vietnam-era veteran status. This application is featured at the end of Bekker’s paper, and, originally, in an Amsterdam bar in 1992, where Paul Bekker first confronted Angrist with claims of bias.

endogeneity is extraordinarily high.

What range of values for ρ is relevant? Three applications are used to calibrate endogeneity: AK91, the Angrist and Evans (1998, AE98) IV estimates using a dummy for same-sex sibships as an instrument for family size, and the Angrist and Lavy (1999, AL99) fuzzy regression discontinuity estimates of class size effects on student learning. These studies span a range of OVB scenarios, from modest (for most of the AK91 estimates), to substantial (in AE98, where OLS exceeds IV by about 50%), to dramatic (in AL99, where IV exceeds small, insignificant OLS estimates by an order of magnitude). Yet, the absolute value of estimated endogeneity is no more than 0.47 in these applications, and over 0.4 only for a single specification and sample. Although three examples do not make a theorem, we argue that the features of these studies that limit endogeneity are common to empirical strategies designed to estimate causal effects or to mitigate attenuation bias in models with measurement error.²

Evidence on the reliability of conventional just-ID IV inference notwithstanding, IV practitioners have come to see weak instruments as problematic in just-identified as well as over-identified models. Responding to Bound et al. (1995) and Staiger and Stock (1997), analysts now routinely report first-stage t - and F -statistics, a practice that’s hard to argue with. Yet, Hall et al. (1996) and others since note that requiring, say, a first-stage F -statistic greater than 10 when instruments are truly weak often does more harm than good. Requiring first-stage estimates to meet a prespecified cutoff amounts to imposition of a pretest that distorts sampling distributions and makes conventional confidence intervals misleading. This is the IV version of the general pretesting problem highlighted by Leeb and Pötscher (2005).³

Analyses of pretest bias to date apply when the analyst is agnostic as to the sign of the first stage. Yet, Andrews and Armstrong (2017) note that the typical just-ID IV scenario includes a theoretical sign restriction. This leads us to consider pretesting strategies that maintain a sign restriction on the population first stage. Specifically, we examine pretesting rules of the form $t_1 > c$, where t_1 is the first-stage t -statistic and c is a constant chosen by the analyst. This examination leads to a novel theoretical result: the median bias of just-ID IV conditional on $t_1 > c$ is minimized by choosing $c = 0$, that is, by screening on the sign of the estimated first-stage. Moreover, median bias of just-ID IV is roughly halved by requiring the estimated first-stage to have the expected (that is, population) sign (we focus on median bias because the conventional just-ID IV estimator has no moments).

Surprisingly, pre-screening on the estimated first-stage sign is also shown to be virtually costless: rejection contours for a sign-screened estimator differ little from those obtained without screening. The upshot is that sign-screening mitigates the already-modest bias of just-ID IV without degrading coverage. To the extent that such screening is a feature of modern empirical work, reported IV estimates reflect the impressively minimal bias characteristic of sign-screened IV. Our theoretical results on the bias-minimizing and bias-mitigating consequences of requiring $t_1 > 0$ therefore strengthen the case for a sanguine view conventional inference for just-ID IV.

Finally, the theorem establishing bias mitigation from screening on an estimated first stage provides an interesting contrast to Andrews and Armstrong (2017), which shows how to use a sign restriction

²The spirit of this argument differs from that in Stock and Yogo (2005), which focuses on worst-case rejection rates over all possible endogeneity values. Lee et al. (2020) and Keane and Neal (2021), discussed further in Section 3 below, consider though largely downplay restrictions on endogeneity.

³Andrews et al. (2019) survey recent empirical scholarship to demonstrate the empirical relevance of pretest bias in IV applications.

on the *population* first stage to construct an unbiased just-ID IV estimator. We show that conditional on the sign of the *estimated* first stage, this estimator, denoted $\hat{\beta}_U$, is no longer unbiased. Rather, $\hat{\beta}_U$ is unbiased by virtue of the fact that it averages two conditional estimators, each biased but in opposite directions. If, as seems likely, empirical sign-screening is endemic in applied microeconomics, this would seem to reduce the appeal of $\hat{\beta}_U$ for empirical practice.

The next section details the just-ID IV setup assuming normally distributed first-stage and reduced-form estimates and derives an expression for endogeneity in terms of OLS OVB. Section 3 reviews the relationship between t -test rejection rates and the parameters that govern the normal model. This section also explains why endogeneity in applied microeconometrics is unlikely to be high enough for conventional IV inference to mislead. Section 4 presents our theoretical results on first-stage screening for just-ID IV. Section 5 concludes. Proofs and details behind numerical calculations appear in the appendix.

2 Setup

The sample is assumed to consist of n units indexed by i , with data on outcome variable, Y_i , a scalar treatment variable, D_i , a vector of covariates, X_i , and a scalar instrument, Z_i . Population regressions of outcome and treatment on the instrument and covariates define the reduced form and first stage. These are written as follows:

$$Y_i = Z_i\delta + X_i'\psi_1 + u_i, \quad (1)$$

$$D_i = Z_i\pi + X_i'\psi_2 + v_i. \quad (2)$$

The parameter of interest is $\beta = \frac{\delta}{\pi}$, the ratio of the reduced-form and first-stage regression coefficients on Z_i . Provided that the instrument, Z_i , satisfies an exclusion restriction and is relevant (i.e. $\pi \neq 0$), this parameter captures the causal effect of D_i on Y_i . More generally, if treatment effects are heterogeneous and a monotonicity condition holds, β is a weighted average of individual causal effects (Angrist & Imbens, 1995; Imbens & Angrist, 1994). While treatment effect heterogeneity affects the interpretation of β , heterogeneity has no bearing on the behavior of the estimators and inference procedures considered in our just-ID IV setting.

Let $\hat{\delta} = \sum_{i=1}^n \tilde{Z}_i Y_i / \sum_{i=1}^n \tilde{Z}_i^2$ and $\hat{\pi} = \sum_{i=1}^n \tilde{Z}_i D_i / \sum_{i=1}^n \tilde{Z}_i^2$ denote OLS estimates of δ and π , where \tilde{Z}_i is the residual from a regression of Z_i on X_i . Under mild regularity conditions that allow the errors (u_i, v_i) to be non-normal, heteroskedastic, and serially or cluster-dependent, $(\hat{\delta}, \hat{\pi})$ is consistent and asymptotically normal as $n \rightarrow \infty$, with an asymptotic covariance matrix that can be consistently estimated. Importantly, this holds regardless of the strength of the instrument. We therefore follow Andrews et al. (2019) and earlier analyses of weak instrument problems by assuming this large-sample approximation holds exactly. Specifically, we assume:

$$\begin{pmatrix} \hat{\delta} \\ \hat{\pi} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \pi\beta \\ \pi \end{pmatrix}, \Sigma = \begin{pmatrix} \sigma_{\hat{\delta}}^2 & \sigma_{\hat{\delta}\hat{\pi}} \\ \sigma_{\hat{\delta}\hat{\pi}} & \sigma_{\hat{\pi}}^2 \end{pmatrix} \right), \quad (3)$$

with a known covariance matrix, Σ . This distributional assumption is implied by the Staiger and Stock

(1997) weak-instrument asymptotic sequence (see Andrews et al. (2019, Section 3.2) for additional discussion and references). Finite-sample results under eq. (3) can therefore be seen as asymptotic under the Staiger and Stock (1997) sequence.

Equation (3) is our only substantive restriction; this assumption allows us to focus on the weak instrument problem, separating this from other finite-sample problems, such as the effect of high-leverage observations on the quality of the normal approximation to the distribution of the OLS estimators $(\hat{\delta}, \hat{\pi})$ and the challenge of standard-error estimation with clustered data.⁴ With (3) as foundation, we derive finite-sample properties of the IV estimator:

$$\hat{\beta}_{IV} = \frac{\hat{\delta}}{\hat{\pi}}, \quad (4)$$

and the null rejection rate for the corresponding Wald test. The latter is based on the t -statistic centered at the parameter of interest, β , divided by the estimated IV standard error, $\hat{\sigma}_{IV}$:

$$t_W = \frac{\hat{\beta}_{IV} - \beta}{\hat{\sigma}_{IV}}; \quad \hat{\sigma}_{IV}^2 = \frac{\sigma_{\hat{\delta}}^2 - 2\sigma_{\hat{\delta}\hat{\pi}}\hat{\beta}_{IV} + \sigma_{\hat{\pi}}^2\hat{\beta}_{IV}^2}{\hat{\pi}^2}, \quad (5)$$

where $\hat{\sigma}_{IV}^2$ estimates the asymptotic variance of $\hat{\beta}_{IV}$ under standard $n \rightarrow \infty$ asymptotics. The corresponding theoretical variance is $\sigma_{IV}^2 = (\sigma_{\hat{\delta}}^2 - 2\sigma_{\hat{\delta}\hat{\pi}}\beta + \sigma_{\hat{\pi}}^2\beta^2)/\pi^2$. In a homoskedastic model with constant causal effects, this simplifies to the familiar formula

$$\sigma_{IV}^2 = \frac{\sigma_{\epsilon}^2}{nE[\ddot{Z}_i^2]\pi^2},$$

where \ddot{Z}_i is the residual from the population projection of Z_i on X_i , and σ_{ϵ}^2 is the variance of the residual in the structural equation,

$$Y_i = D_i\beta + X_i'(\psi_1 - \psi_2\beta) + \epsilon_i, \quad (6)$$

that motivates IV estimation in the classic linear set-up (the structural residual is $\epsilon_i = u_i - v_i\beta$).

The Anderson and Rubin (1949, AR) statistic for just-ID IV, which can be written:

$$t_{AR} = \frac{\hat{\delta} - \hat{\pi}\beta}{\sqrt{\sigma_{\hat{\delta}}^2 - 2\sigma_{\hat{\delta}\hat{\pi}}\beta + \sigma_{\hat{\pi}}^2\beta^2}}, \quad (7)$$

is a useful complement to t_W . The AR statistic differs from t_W in that it replaces $\hat{\beta}_{IV}$ with the null value of β in the formula for $\hat{\sigma}_{IV}^2$: in the context of the just-ID IV model described by eq. (3), the square of t_{AR} equals the Lagrange multiplier statistic testing whether $\frac{\hat{\delta}}{\hat{\pi}}$ equals β .⁵ The AR test has size undistorted by weak instruments under the Staiger and Stock (1997) sequence. Moreover, in our just-ID IV context, an AR test is optimal among unbiased tests (Moreira, 2009). AR tests

⁴Young (2021) discusses these problems in an IV context.

⁵Since the moment restrictions in the IV model are linear, Proposition 3 in Newey and West (1987) implies that t_{AR}^2 is also the relevant likelihood ratio statistic. See Anderson and Rubin (1949) for the general form of the AR statistic in over-identified models with a fixed number of instruments and Mikusheva and Sun (2021) for an adaptation to models with many weak instruments.

are also compelling by virtue of the fact that, when testing $\beta = 0$, t_{AR} is the t -statistic for the associated reduced form. It's hard to imagine a convincing case for statistical significance of an IV estimate when the associated reduced form is statistically indistinguishable from zero. We focus here on a conventional Wald-type t -statistic because the conventional test—and the associated confidence intervals—are ubiquitous in applied work.⁶

Given the assumption of a known covariance matrix for the first-stage and reduced-form estimates, both t_W and $\hat{\sigma}_{IV}^2$ depend on the data only through $(\hat{\delta}, \hat{\pi})$. These have distributions determined by the two unknown parameters, π and β . It's illuminating, however, to reparametrize in terms of instrument strength and the degree of endogeneity (a reparametrization adopted in Staiger and Stock (1997) and Lee et al. (2020), among others). The first parameter in this scheme, denoted $E[F]$, is defined as:

$$E[F] = \pi^2 / \sigma_{\hat{\pi}}^2 + 1.$$

Because $E[F]$ is the expectation of $F = \hat{\pi}^2 / \sigma_{\hat{\pi}}^2$, the F -statistic testing $\pi = 0$, it's sometimes called the *population first-stage F-statistic*, a term adopted here. Since π is a scalar, $E[F] = E[t_1]^2 + 1$, where $t_1 = \hat{\pi} / \sigma_{\hat{\pi}}$ is the first-stage t -statistic.

The second parameter is defined as:

$$\rho = \text{cor}(\hat{\delta} - \hat{\pi}\beta, \hat{\pi}) = \frac{\sigma_{\hat{\pi}}}{\sqrt{\sigma_{\hat{\delta}}^2 - 2\beta\sigma_{\hat{\delta}\hat{\pi}} + \sigma_{\hat{\pi}}^2\beta^2}} \times (\sigma_{\hat{\delta}\hat{\pi}} / \sigma_{\hat{\pi}}^2 - \beta). \quad (8)$$

With independent heteroskedastic errors, ρ is also given by $\text{cor}(\ddot{Z}_i\epsilon_i, \ddot{Z}_iv_i)$. When, in addition, the errors (u_i, v_i) are homoskedastic, $\rho = \text{cor}(\epsilon_i, v_i)$, where ϵ_i is the structural residual in (6). We therefore refer to ρ as (the degree of) *endogeneity*.⁷

With weak instruments as well as homoskedastic error terms, ρ is proportional to the bias of the OLS estimand. This can be seen by using the first-stage and reduced-form equations to write the OLS slope coefficient, β_{OLS} , as follows:

$$\begin{aligned} \beta_{OLS} &= \frac{E[\ddot{D}_i Y_i]}{E[\ddot{D}_i^2]} = \frac{E[\ddot{Z}_i^2]\pi^2\beta + E[u_i v_i]}{E[\ddot{Z}_i^2]\pi^2 + \sigma_v^2} \\ &= R_p^2\beta + (1 - R_p^2)\frac{E[u_i v_i]}{\sigma_v^2}; \quad R_p^2 = \frac{E[\ddot{Z}_i^2]\pi^2}{E[\ddot{Z}_i^2]\pi^2 + \sigma_v^2}, \end{aligned} \quad (9)$$

where \ddot{D}_i is the residual from a population regression of D_i on X_i , and $\sigma_v^2 = E[v_i^2]$. The weight multiplying β in (9), denoted R_p^2 , is the population partial R^2 generated by adding the instrument to the first-stage regression. When the instrument is weak, R_p^2 is close to zero, and (9) is approximately $E[u_i v_i] / \sigma_v^2$. The OLS estimand likewise converges to $E[u_i v_i] / \sigma_v^2$ in the Staiger and Stock (1997) weak-

⁶Since Wald and Lagrange multiplier-type tests are first-order asymptotically equivalent, the t_W vs t_{AR} distinction is asymptotically negligible. In finite samples, inference based on these two statistics may differ. While rejection rates for the two tests cannot be ranked in general, confidence intervals derived by inverting t_{AR} have infinite length when $F < z_{1-\alpha/2}^2$ (i.e., less than about 4 for 95% intervals). Finite AR intervals are always wider than the corresponding intervals based on t_W (see, e.g., Lee et al., 2020).

⁷This simplification is obtained using the fact that, under homoskedasticity, the variance of v_i is $\sigma_v^2 = \sigma_{\hat{\pi}}^2 \cdot nE[\ddot{Z}_i^2]$ and the variance of ϵ_i is $\sigma_{\epsilon}^2 = (\sigma_{\hat{\delta}}^2 - 2\beta\sigma_{\hat{\delta}\hat{\pi}} + \sigma_{\hat{\pi}}^2\beta^2) \cdot nE[\ddot{Z}_i^2]$, with $\text{cov}(v_i, \epsilon_i) = (\sigma_{\hat{\delta}\hat{\pi}} - \beta\sigma_{\hat{\pi}}^2) \cdot nE[\ddot{Z}_i^2]$. The homoskedastic formula for the variance of ϵ_i also leads yields the simplification of the formula for σ_{IV}^2 noted above.

instrument sequence (which takes $\pi \rightarrow 0$). This in turn equals $\sigma_{\hat{\delta}\hat{\pi}}/\sigma_{\hat{\pi}}^2$ under homoskedasticity, so the second term on the right-hand side of (8),

$$\sigma_{\hat{\delta}\hat{\pi}}/\sigma_{\hat{\pi}}^2 - \beta = \beta_{WOLS} - \beta, \quad (10)$$

is the weak-instrument OVB of OLS (where we've introduced the notation β_{WOLS} for $\sigma_{\hat{\delta}\hat{\pi}}/\sigma_{\hat{\pi}}^2$). Moreover, when $\pi = 0$, it follows from (3) that $\beta_{WOLS} - \beta$ is the median bias of $\hat{\beta}_{IV}$, a result requiring no independence or heteroskedasticity assumptions on the errors in (1) and (2).⁸ Thus, ρ also measures endogeneity in the sense that it's proportional to the median bias of the IV estimator when the instrument is irrelevant.

3 Rejection Rates in Theory and Practice

We're interested in the bias of $\hat{\beta}_{IV}$ and in t -test rejection rates when the null hypothesis is true. Beginning with the latter, the null rejection rate for a two-sided t -test with level α is the probability that the absolute value of a t -statistic, $|t_W|$, exceeds $z_{1-\alpha/2}$, the $1-\alpha/2$ quantile of a standard normal distribution. This is:

$$R_W = P_{E[F],\rho}(|t_W| > z_{1-\alpha/2}),$$

where $P_{E[F],\rho}$ is the distribution of t_W parameterized by $E[F], \rho$. We evaluate R_W by numerical integration, a computation detailed in Appendix A.2.

Summarizing the behavior of a conventional 5% nominal test, Panel (a) in Figure 1 depicts rejection rates for t_W as a contour plot indexed by ρ and $E[F]$. The figure shows that rejection rates greatly exceed the nominal 5% rate only if the instrument is weak (i.e., $E[F]$ is close to 1) *and* endogeneity is high. In particular, if $|\rho| < 0.76$, rejection rates are below 10%, regardless of the strength of the first stage. If $|\rho| < 0.565$, the nominal 5% test under-rejects (this cutoff is also noted in Lee et al. (2020)). A simple corollary, further substantiated below, is that the coverage of conventional nominal 95% confidence intervals for $\hat{\beta}_{IV}$ is likely to be satisfactory in most applications.

The modest over-rejection seen in Figure 1 is explained by a signal feature of just-ID IV: the bias of $\hat{\beta}_{IV}$ rises as the instrument grows weaker, but precision falls apace. The IV standard error reflects this lack of precision well enough that, unless endogeneity is egregious, inference is distorted little. This contrasts with overidentified 2SLS with many weak instruments (as in Bekker (1994) and Bound et al. (1995)), where, bias notwithstanding, the usual standard errors for 2SLS remain small enough for the t -statistic to be misleading.

Our conclusions here also contrast with those drawn in Stock and Yogo (2005) and Lee et al. (2020) regarding the reliability of inference based on a conventional just-ID IV t -statistic. Although Lee et al. (2020) present a similar plot, both studies emphasize worst-case rejection rates over ρ , for a given $E[F]$. As can be seen in our Figure 1, this worst-case rejection rate occurs at $|\rho| = 1$. In

⁸Assumption (3) implies that we can write reduced form and first stage estimates as $\hat{\delta} = \pi\beta + (\sigma_{\hat{\delta}\hat{\pi}}/\sigma_{\hat{\pi}})\mathcal{Z}_\pi + (\sigma_\delta^2 - \sigma_{\hat{\delta}\hat{\pi}}^2/\sigma_{\hat{\pi}}^2)^{1/2}\mathcal{Z}_\delta$ and $\hat{\pi} = \pi + \sigma_{\hat{\pi}}\mathcal{Z}_\pi$, where \mathcal{Z}_δ and \mathcal{Z}_π are independent standard normal variables. When $\pi = 0$, therefore, $\hat{\beta}_{IV} = \frac{1}{\sigma_{\hat{\pi}}}(\sigma_\delta^2 - \sigma_{\hat{\delta}\hat{\pi}}^2/\sigma_{\hat{\pi}}^2)^{1/2}(\mathcal{Z}_\delta/\mathcal{Z}_\pi) + \beta_{WOLS}$, the median of which is β_{WOLS} since $\mathcal{Z}_\delta/\mathcal{Z}_\pi$ has a standard Cauchy distribution with zero median.

the same spirit, Keane and Neal (2021) highlights simulations showing that conventional just-ID IV t -tests can be misleading when endogeneity is very high. Sections 3.1 and 3.2 explain why we are not much concerned with high values of ρ .

Keane and Neal (2021) also observe that, since $\hat{\sigma}_{IV}$ and $\hat{\beta}_{IV}$ tend to be negatively correlated when ρ is positive, most false rejections occur when $\hat{\beta}_{IV} > \beta$. This, they argue, militates so strongly against t_W that conventional Wald tests are to be avoided even with a first-stage F in the thousands. As we see it, the fact that tests based on t_W have asymmetric power does not make conventional frequentist inference unreliable. The conventional standard for reliability of inference is the accuracy of confidence interval coverage, gauged without reference to the particular value of the estimate. Our analysis adheres to this standard.⁹

3.1 The Anatomy of Endogeneity

We put endogeneity in context using three applications. These are the AK91 study that launched the modern weak instruments literature, the AE98 estimates using a dummy for samesex sibships (of first- and second-born children) as an instrument for family size, and the AL99 fuzzy regression discontinuity estimates of class size effects. The AE98 and AL99 first-stage t -statistics exceed those for AK91 and are arguably out of the zone where an instrument might be considered weak. With a first-stage t -statistic of almost 8, the AK91 quarter-of-birth instrument also seems strong enough. But all three studies can be used to calibrate endogeneity and to document contextual features that constrain it.

Table 1 reports key statistics for specifications drawn from each study (some estimates in the table differ slightly from those in the original). The first row in Panel A shows estimates of the economic returns to schooling in the AK91 sample of men born 1920–29. Here, OLS and IV estimates equal 0.080 and 0.072, respectively. These are close, so endogeneity is small in this case, with an estimated ρ of only 0.043. Schooling returns estimated in the second AK91 sample, consisting of men born 1930–39, exhibit more OVB. In this sample, the IV estimate of 0.105 surprisingly exceeds the OLS estimate of 0.072 (IV estimation of the returns to schooling is usually motivated by a concern that omitted ability controls causes OLS estimates to be too large). Endogeneity is correspondingly larger at $\rho = -0.175$, but still well outside the danger zone depicted in Figure 1.¹⁰

The AK91, AE98, and AL99 studies span a range of OVB scenarios, from modest in the first AK91 sample, to substantial in AE98 (where OLS magnitudes consistently exceed IV by at least 50%), to dramatic in AL99 (where IV exceeds small, insignificant OLS estimates, mostly by an order of magnitude, and sometimes with a sign flip). Yet, the magnitude of endogeneity exceeds 0.40 in only one specification, that for reading scores in the AL99 discontinuity sample (which consists of classes

⁹To see why such conditional coverage can be problematic, consider the standard $(1 - \alpha)$ -level confidence interval $[U_{(n)}, U_{(n)}/\alpha^{1/n}]$ for the endpoint of a uniform distribution, $\text{Unif}(0, \theta)$, where $U_{(n)}$ is the largest order statistic for sample of size n . Here all false rejections occur when the minimum variance unbiased estimate $nU_{(n)}/(n + 1)$ (provided $\alpha < 1/e$) is below θ . Also, similar to the IV problem, the length of the interval is correlated with the point estimate; here, the correlation is 1.

¹⁰Estimates and standard errors for ρ in the table are computed as described in Appendix A.1. The estimates may differ from parameter ρ because $\hat{\beta}_{IV}$ is biased, and because of sampling variance. In the examples analyzed here, however, the instruments are not particularly weak, so the bias in estimated endogeneity is negligible. Standard errors for estimated endogeneity use the delta method; confidence interval endpoints for ρ based on inverting the AR test are virtually indistinguishable from these.

in schools with enrollment near the cutoff that determines class size). Just-ID IV inference in all three of these studies is therefore unlikely to be compromised by weak instruments.¹¹

Although the consistently moderate levels of endogeneity documented in Table 1 do not add up to a theorem, these applications have features in common with many IV-driven microeconomic investigations of causal effects. First, measured against endogenous variable standard deviations, first-stage coefficient estimates, denoted here by $\hat{\pi}$, are small. In particular, the range of first-stage impact in eq. (2) runs from under 4% of the endogenous variable standard deviation in AK91, to about 12% of the endogenous variable standard deviation in AE98, with the AL99 first stage in-between.

Most importantly, endogeneity in research on causal effects is often capped by the modest size of the causal effects of interest. To make this point, it's helpful to write ρ as a function of OVB. Using eqs. (9) and (10), we can express ρ under homoskedasticity as:

$$\begin{aligned} \rho &= \frac{\sigma_v}{\sigma_\epsilon}(\beta_{WOLS} - \beta) \\ &= \frac{\sigma_v}{\sigma_\epsilon} \left(\frac{\beta_{OLS} - \beta}{1 - R_p^2} \right) \approx \frac{\sigma_D}{\sigma_Y}(\beta_{OLS} - \beta). \end{aligned} \tag{11}$$

We can use this expression to compute ρ by replacing β with $\hat{\beta}_{IV}$. The relevance of this representation of ρ can be seen in the AE98 estimates of the effects of a third child on weeks worked by women aged 21–35 in the 1980 Census. Here, the first-stage partial R-square (R_p^2) is close to zero, while the difference between the conventional OLS estimate and the corresponding IV estimate is 3.42. The term multiplying this, $\frac{\sigma_v}{\sigma_\epsilon}$, is well-approximated by the ratio of the endogenous variable standard deviation to the dependent variable standard deviation (denoted $\frac{\sigma_D}{\sigma_Y}$), a ratio of about 0.022. The product of these two terms gives 0.075, equal to the value of ρ reported in the table for this sample.

Equation (11) suggests alternative bounds on endogeneity. For starters, in the AK91 scenario, it seems reasonable to assume that the (causal) economic returns to schooling are no more than double the OLS estimate and certainly positive. Under these restrictions, the descriptive statistics in Table 1, which approximate $\frac{\sigma_v}{\sigma_\epsilon}$ at around 5.2, suggest $|\rho|$ can be no more than about 0.41. Although substantial, this is well below the 0.76 threshold for coverage concerns. With β bounded below by zero, large magnitudes of ρ require β to far exceed β_{WOLS} . Only when the causal effect of schooling is triple the OLS estimate (so that OLS is *too small* by 0.16) does the endogeneity danger zone become relevant.¹²

Many microeconomic IV applications involve linear probability models in which causal effects are changes in probabilities. The AE98 estimates of the effect of the birth of a third child on female labor force participation in 1980, for example, range from roughly -0.18 for OLS to -0.12 for IV. Labor force participation rates for women with only two children run around 57%. Causal effects might therefore be as large as -0.57 , but no larger, since probabilities can't be negative. In this case,

¹¹As Table 1 shows, estimates of ρ computed under homoskedasticity, reported as $\hat{\rho}^*$, are almost identical to the heteroskedasticity-based estimates in the AK91 and AE98 samples. For the AL99 samples, endogeneity parameters computed using homoskedasticity are smaller.

¹²Keane and Neal (2021) consider bounds on ρ in the context of estimates of the economic returns to schooling motivated by the view that OLS returns should exceed causal effects. Although this seems defensible, it's worth noting that the literature surveyed by Card (2001) reports many IV estimates in excess of the corresponding OLS estimates, a pattern first highlighted by Lang (1993).

$\frac{\sigma_v}{\sigma_\epsilon}$ is about 1 (again, using standard deviations in the data rather than residuals), so $\beta_{OLS} - \beta$ can be no larger than $-0.18 + 0.57 = 0.39$, thereby bounding ρ at this value. This generous bound makes no use of the fact that selection bias is likely to make OLS estimates of family-size effects on female supply too large (in magnitude) rather than too small. Other applications with Bernoulli outcomes admit similar sorts of bounds.

A related argument, appropriate for models with continuous outcomes, shows endogeneity to be constrained by plausible values for causal effects measured in standard deviation units. This line of reasoning is especially apt for education research where standardized effect sizes are widely reported. The influential Tennessee STAR class size experiment analyzed in Krueger (1999), for instance, generated a reduction of 7 students per class, roughly one standard deviation of class size in the AL99 data. The STAR experiment yielded treatment effects of about 0.2σ , an impact typical of education interventions deemed to have been effective. At the same time, education researchers often view effect sizes as large as half a standard deviation in the outcome distribution as rare, if not implausible. Using the fact that $\frac{\sigma_v}{\sigma_\epsilon}$ is about equal to $(1 - R_p^2)$ in the AL99 data, the scenario of a half-standard deviation effect size generated by a one-standard deviation reduction in class size implies $\frac{\sigma_v}{\sigma_\epsilon} \frac{\beta}{1 - R_p^2} = -0.5$ on the second line of eq. (11). At the same time, OLS estimates of class size effects in AL99 are mostly zero (as is often found in class size research; see e.g., Hanushek (1986)), so the magnitude of endogeneity is capped at 0.5.

Contributing to all three of these empirically-grounded arguments is the fact that endogeneity under homoskedasticity can be split into the difference between two R-squared-like terms:

$$\rho \approx \frac{\sigma_D}{\sigma_Y}(\beta_{OLS} - \beta) = \frac{\sigma_D}{\sigma_Y}\beta_{OLS} - \frac{\sigma_D}{\sigma_Y}\beta.$$

The square of the first term, $(\frac{\sigma_D}{\sigma_Y}\beta_{OLS})^2$, is the variation in the dependent variable accounted for by D_i in an analysis-of-variance for Y_i . In microeconomic applications, this sort of R^2 term is mostly small, as is the causal analog that determines the square of the second term, $(\frac{\sigma_D}{\sigma_Y}\beta)^2$. The small size of these two R^2 terms limits the magnitude of the difference between them. Consistent with this claim to generality, the many IV estimates collated in Chernozhukov and Hansen (2008) likewise show modest endogeneity.

Finally, it's worth noting that these bounds on ρ are remarkably insensitive to the presumption of homoskedasticity. To see this, note that we can also bound ρ by plugging our estimates of components of covariance matrix Σ , which allow for heteroskedasticity or clustering, into the definition of ρ in eq. (8), and bounding the treatment effect β . Assuming (as above) that the returns to schooling lie between 0 and 0.16, effects of a third child on labor force participation lie between 0 and -0.57 , and class size effects lie between 0 and $0.5\sigma_Y/\sigma_D$, yields upper bounds on $|\rho|$ equal to 0.41, 0.36, and 0.57, respectively. These numbers are close to the corresponding homoskedasticity-based bounds.

3.2 When Measurement Error Motivates IV

In addition to IV for causal effects, a second major arena for microeconomic IV involves models with measurement error. Suppose the regression of interest is $Y_i = D_i^*\beta + X_i'\gamma + \eta_i$, where η_i is a residual uncorrelated with (D_i^*, X_i) by definition. The regressor D_i^* is unobserved; we see only a noisy measure,

$D_i = D_i^* + e_i$, where the measurement error, e_i , is assumed to be classical, that is uncorrelated with (D_i^*, X_i, η_i) . Replacing D_i^* with D_i yields the structural equation to be instrumented:

$$\begin{aligned} Y_i &= D_i\beta + X_i'\gamma + (\eta_i - e_i\beta) \\ &= D_i\beta + X_i'\gamma + \epsilon_i, \end{aligned}$$

where $\epsilon_i = \eta_i - e_i\beta$ is the structural residual. Given an instrument correlated with D_i^* and uncorrelated with ϵ_i , the coefficients of interest are consistently estimated by IV. The first stage in this scenario can be written as in (2), with first-stage residual, v_i .

To calibrate endogeneity in this model, note first that, given the classical measurement error assumption, $\text{cov}(v_i, \epsilon_i) = -\sigma_e^2\beta$. Under homoskedasticity, endogeneity squared can therefore be written:

$$\rho^2 = \frac{\sigma_e^4\beta^2}{\sigma_v^2\sigma_\epsilon^2} = \frac{\sigma_e^4\beta^2}{\sigma_v^2(\sigma_\eta^2 + \beta^2\sigma_e^2)} \leq \frac{\sigma_e^2}{\sigma_v^2} = \frac{1-r}{1-R_p^2}, \quad (12)$$

where $r = \sigma_{D^*}^2/\sigma_D^2$ denotes the reliability (or signal-to-noise ratio) of mismeasured D_i , after partialing out covariates.¹³ Although we can't speak to reliability across all fields, labor economists have collected evidence on the reliability of key variables of interest. These include schooling, earnings, hours worked, and hourly wages. Schooling often appears on the right-hand side of wage equations, while earnings, hours, and hourly wages are used in various configurations to estimate labor supply elasticities. The Angrist and Krueger (1999) summary of reliability estimates suggests $r \approx 0.9$ for schooling and $r \approx 0.8$ for earnings, falling to about 0.65 – 0.75 for hours worked and hourly wages. The lower end of this range may be more relevant for wage reliability after partialing out covariates.¹⁴ With $r = 0.65$ as a reasonably conservative value, we'd need to see an R_p^2 equal to at least 0.4 for ρ to reach 0.76. But $E[F] = \frac{nR_p^2}{1-R_p^2} + 1$, so, at this level of first-stage fit, $E[F]$ is way nowhere near the trouble zone for any sample size that's empirically relevant. This suggests that, unless reliability is unusually low, microeconomic measurement error can be expected to generate parameter combinations for which conventional IV inference is trouble-free.

4 Bias Under a Good Sign

Having made an empirical case for conventional inference with just-ID IV, we add a novel analytical argument. This builds on the idea that IV identification strategies are most credible when a strong institutional or theoretical foundation explains the first stage. These foundations usually imply a sign for π . In the AK91 application, for example, the quarter-of-birth first stage arises from the fact that children born later in the year enter school younger, and are therefore constrained by compulsory attendance laws to stay in school longer than those born earlier. The AE98 same-sex instrument for family size is predicated on parents' preference for mixed-sex sibships. The AL99 Maimonides Rule instrument for class size is derived from regulations that determine class size as a function of

¹³The first equality in (12) follows from the definition of correlation, the middle inequality uses the fact that σ_η^2 must be non-negative, and the last equality uses the definition of partial R^2 in eq. (9).

¹⁴See Table 11 in Angrist and Krueger (1999). The reliability of variables other than wages is less sensitive to covariate adjustment.

enrollment. In these and many other applied micro applications, institutions or preferences determine the sign of π .

4.1 Sign-Screened Bias and Coverage

We gauge estimator performance using median bias since the expectation of a just-ID IV estimator is undefined (2SLS moments exists only for over-identified models). Assuming the sign of π is known, the theorem below shows that $c = 0$ minimizes median bias of $\hat{\beta}_{IV}$ among screening rules of the form $t_1 > c$.

Theorem 1. *Consider the model in eq. (3), and suppose $\pi > 0$. The absolute value of the median bias of $\hat{\beta}_{IV}$ conditional on $t_1 > c$, $|\text{median}_{E[F],\rho}(\hat{\beta}_{IV} - \beta \mid t_1 > c)|$, is minimized at $c = 0$.*

Note that empirical sign-screening yields the greatest bias reduction uniformly over *all* parameter values $(E[F], \rho)$. In particular, empirical sign-screening reduces median bias relative to no screening, since the latter sets $c = -\infty$.

For intuition as to why sign-screening is optimal, note first that joint normality of first-stage and reduced form estimates implies that the distribution of $\hat{\beta}_{IV} - \beta$ conditional of t_1 is normal with a mean and median that can be written:

$$E[\hat{\beta}_{IV} - \beta \mid t_1] = (\beta_{WOLS} - \beta) \frac{t_1 - E[t_1]}{t_1}. \quad (13)$$

Suppose $\rho > 0$, so $\beta_{WOLS} - \beta > 0$ is positive. When t_1 is positive, eq. (13) implies that conditional median bias is increasing in t_1 . This implies that for any $a > 0$, screening on $t_1 > 0$ is better than screening on $t_1 > a$, since IV estimates in samples with $t_1 \in [0, a]$ are less biased than estimates with $t_1 > a$. To see why screening on $t_1 > 0$ is better than screening on $t_1 > a$ for $a < 0$, note that in this case, $(t_1 - E[t_1])/t_1 > 1$, so that conditional on t_1 , IV bias exceeds that of OLS, while, as we show in Theorem 2 below, median bias of IV conditional on $t_1 > 0$ is smaller than OLS bias. Inclusion of samples with negative t_1 therefore increases median bias. The upshot is an optimal screening cutoff of zero. Figure 2 demonstrates this for selected values of $E[F]$ (numerical computation of median bias is described in Appendix A.6). The kinks at zero in the figure reflect the fact that the median of $\hat{\beta}_{IV}$ conditional on t_1 is discontinuous at zero.

The fact that sign-screening affords *substantial* bias reduction is established by a theorem characterizing median IV bias scaled by the weak-IV bias of OLS. Rescaling simplifies bias formulas, while the relationship between conditional and unconditional bias stands without this.¹⁵ The theorem below gives a result for worst-case relative bias over ρ , which obtains in the limit as $|\rho| \rightarrow 0$ (this is not the same as relative bias when $\rho = 0$; with no endogeneity, both IV and OLS are unbiased, so that relative bias is discontinuous in ρ). The relationship between ρ and relative median bias derived here contrasts with that in Section 3, which shows higher endogeneity leads to worse coverage. This reversal reflects the fact that, although the bias of $\hat{\beta}_{IV}$ increases with endogeneity, OLS OVB increases faster. Worst-case relative bias is characterized by:

¹⁵Stock and Yogo (2005) use a similar rescaling, focusing on relative mean bias for 2SLS models with over-identifying restrictions.

Theorem 2. Consider the model in eq. (3), and suppose that $\pi > 0$. Then, the unconditional relative median bias of $\hat{\beta}_{IV}$ is given by

$$\sup_{\rho} \left| \frac{\text{median}_{E[F],\rho}(\hat{\beta}_{IV} - \beta)}{\beta_{WOLS} - \beta} \right| = \frac{\phi(E[t_1])}{E[t_1][\Phi(E[t_1]) - 1/2] + \phi(E[t_1])}. \quad (14)$$

Moreover, if $E[t_1] \geq 0.84$, the relative median bias of $\hat{\beta}_{IV}$ conditional on $\hat{\pi} > 0$ satisfies

$$\sup_{\rho} \left| \frac{\text{median}_{E[F],\rho}(\hat{\beta}_{IV} - \beta \mid \hat{\pi} > 0)}{\beta_{WOLS} - \beta} \right| = \frac{\phi(E[t_1])}{E[t_1]\Phi(E[t_1]) + \phi(E[t_1])}. \quad (15)$$

Equivalently, these expressions give the limit of relative unconditional and conditional median bias as $|\rho| \rightarrow 0$.

Given the assumptions of the theorem, IV with or without sign-screening has bias below the weak-instrument bias of OLS since the right-hand sides of both (14) and (15) are in $(0, 1)$ as long as $\pi > 0$. Moreover, the ratio of the two bias expressions in the theorem is close to 0.5 for all but the smallest values of $E[F]$. Specifically, the ratio of conditional to unconditional bias is:

$$\frac{0.5E[t_1]}{E[t_1]\Phi(E[t_1]) + \phi(E[t_1])}.$$

This quantity is within 1 percentage point of 0.5 once $E[t_1]$ greater than about 1.5 since the normal cdf is then close to one and the normal density close to zero.

Theorem 2 describes worst-case bias over ρ . Remarkably, however, the bias reduction from sign-screening varies little with the degree of endogeneity. This is documented in Figure 3, which plots relative bias as a function of the population first-stage F , using shading to show the variation in relative bias as a function of ρ (as in the previous figure, this figure plots numerical calculations detailed in Appendix A.6). We see, for example, that for $E[F]$ around 3.5, the ratio of sign-screened to unconditional bias varies between 0.5–0.52, converging quickly to 0.5 thereafter.¹⁶

Screening on the sign of the first-stage estimate clearly mitigates bias. But perhaps this sort of screening is costly on the inference side. Econometricians have long warned that screening on the first-stage F -statistic is a form of pretesting that may do more harm than good. Intuitively, when π is truly zero, large F -statistics overstate first-stage strength, leading to overly optimistic standard errors. This problem is compounded by the fact that large F s signal realizations in which the in-sample correlation between instruments and structural errors is largest, exacerbating the bias of $\hat{\beta}_{IV}$. Consequently, when instruments are truly weak, pretesting can lead to confidence intervals with very poor coverage. It's therefore worth investigating whether empirical sign-screening runs a similar risk.

As it turns out, pretesting concerns here are unfounded. By way of evidence on this point, Panel (b) in Figure 1 plots rejection contours for a conventional (second-stage) t -test conditional on $\hat{\pi} > 0$.

¹⁶Richardson (1968) shows that under homoskedasticity, the relative (mean) bias of over-identified 2SLS relative to (weak-instrument) OLS bias is unrelated to ρ .

That is, the figure plots contours for:

$$R_W^c = P_{E[F],\rho}(|t_W| > z_{1-\alpha/2} \mid \hat{\pi} > 0).$$

Comparison of the two panels in Figure 1 suggests sign-screening affects rejection rates little. For instance, the endogeneity threshold required to keep rejection rates below 10% is $|\rho| \leq 0.75$, close to the unconditional value of 0.76 otherwise required for this. This happy finding is explained by the fact that, when the instrument is very weak, sign-screening has two effects. On one hand, screening out wrong-signed first-stage estimates tends to overestimate of first-stage strength. At the same time, in contrast to screening on first-stage F , the bias of $\hat{\beta}_{IV}$ is reduced. These two effects are just about offsetting, so that the conditional rejection contours in panel (b) of Figure 1 are much like the unconditional contours plotted in panel (a).

4.2 Theoretical Sign Restrictions Only

As we noted at the outset, Andrews and Armstrong (2017) introduce an estimator that's unbiased given a sign restriction on the population first-stage coefficient, rather than the sign of estimated $\hat{\pi}$. This unbiased estimator is:

$$\hat{\beta}_U \equiv \hat{\tau}(\hat{\delta} - \beta_{WOLS}\hat{\pi}) + \beta_{WOLS} = t_1\mu(t_1)\hat{\beta}_{IV} + (1 - t_1\mu(t_1))\beta_{WOLS}, \quad (16)$$

where $\mu(x) = \frac{1-\Phi(x)}{\phi(x)}$ is the Mills' ratio of a standard normal random variable (ϕ and Φ denote the standard normal density and cdf, respectively). Estimator $\hat{\beta}_U$ is unbiased because, under normality of first-stage estimates and given $\pi > 0$,

$$E\left[\frac{\mu(t_1)}{\sigma_{\hat{\pi}}}\right] = \frac{1}{\pi},$$

that is, $\hat{\tau} \equiv \frac{\mu(t_1)}{\sigma_{\hat{\pi}}}$ is an unbiased estimator of the reciprocal of π . Moreover, $\hat{\delta} - \beta_{WOLS}\hat{\pi}$ and $\hat{\tau}$ are uncorrelated, since β_{WOLS} is the slope in the regression of the estimated reduced form on the estimated first stage. Unbiasedness then follows from the fact that $E[\hat{\delta} - \beta_{WOLS}\hat{\pi}] = (\beta - \beta_{WOLS})\pi$.

$\hat{\beta}_U$ has an interesting and counter-intuitive shrinkage interpretation when $t_1 > 0$. Observe that

$$0 \leq 1 - t_1\mu(t_1) \leq \frac{1}{t_1^2} \quad (17)$$

when $t_1 > 0$ (this is implied by a Mill's ratio inequality given in Feller (1968)). Thus, when the first stage is right-signed, the weights $t_1\mu(t_1)$ in eq. (16) lie between 0 and 1, and $\hat{\beta}_U$ shrinks the conventional IV estimate towards OLS.

The shrinkage interpretation of $\hat{\beta}_U$ seems surprising since shrinkage *towards* OLS increases bias. This fact is reconciled with the unbiasedness of $\hat{\beta}_U$ by the following theorem:

Theorem 3. *Consider the model in (3), and suppose that $\pi > 0$. Then, the mean bias of $\hat{\beta}_U$*

conditional on $t_1 > 0$ can be written:

$$E[\hat{\beta}_U - \beta \mid t_1 > 0] = \frac{0.5e^{-0.5E[t_1]^2}}{\Phi(E[t_1])}(\beta_{WOLS} - \beta),$$

while, conditional on $t_1 < 0$, relative mean bias is:

$$E[\hat{\beta}_U - \beta \mid t_1 < 0] = -\frac{0.5e^{-0.5E[t_1]^2}}{1 - \Phi(E[t_1])}(\beta_{WOLS} - \beta).$$

Note that the denominators of these expressions equal the probability t_1 is positive and negative, respectively. $\hat{\beta}_U$ is therefore unbiased because it averages conditional positive bias when $t_1 > 0$ and conditional negative bias when $t_1 < 0$.

An analyst who is prepared to sign the population first stage seems unlikely to ignore the sign of the estimated first stage. Yet, when it comes to $\hat{\beta}_U$, empirical sign-screening results in more bias not less. This would seem to strip $\hat{\beta}_U$ of its appeal. And use of median rather than mean bias to measure performance does not ameliorate this: Appendix A.6 shows that the conditional median bias of $\hat{\beta}_{IV}$ is always less than that of $\hat{\beta}_U$, and at least 50% smaller once $E[t_1] \geq 1$.¹⁷

5 Summary and Conclusions

Assuming reduced-form and first-stage estimates are approximately normally distributed, null rejection rates for conventional t -tests in just-ID IV models are distorted little unless endogeneity is exceptionally high. A corollary to this fact is the good coverage of conventional confidence intervals in the just-ID case. Three widely-cited applications, two of which exhibit considerable OLS OVB, are characterized by moderate endogeneity and consequently fall well inside the low-distortion just-ID IV comfort zone. We've argued that these three examples should be seen as representative rather than idiosyncratic: the structure of much applied micro research naturally bounds endogeneity.

We've also developed a new theoretical explanation of modest bias in just-ID IV. As Andrews and Armstrong (2017) note, in the most convincing applications of just-ID IV restrict the sign of the first stage. Unlike Andrews and Armstrong (2017), however, we impose the same sign restriction on the estimated as well as the theoretical first stage. In contrast to screening on first-stage F , which may do more harm than good, empirical sign-screening roughly halves the median bias of the IV estimator while leaving inference uncompromised. Since most analysts likely impose an estimated first-stage sign screen as a matter of course, the bias reduction sign-conditioning engenders is already reflected in published empirical work.

What practical lesson should we draw from this? In the context of the AK91, AE98, and AL99 studies, first-stage sign screening adds no action items to the empirical agenda. The first-stage estimates in these applications are robustly right-signed. In applications with weaker instruments than these, an empirical strategy that begins by examining the first-stage sign would seem to have no down-

¹⁷Andrews and Armstrong (2017) shows numerically that the unconditional median bias of $\hat{\beta}_U$ is smaller than that of $\hat{\beta}_{IV}$ when $E[F]$ is small, while this bias ranking reverses for larger $E[F]$. Andrews and Armstrong (2017) notes also that the median absolute deviation of $\hat{\beta}_U$ is always smaller than that of $\hat{\beta}_{IV}$. Our numerical calculations indicate that this fails to hold for all parameter values conditional on the estimated first stage sign.

side. Claims for credible causal evidence requires more than this, however. In AK91, for instance, the quarter-of-birth story holds water because schooling can be seen to move sharply up and down with quarter of birth as predicted by compulsory attendance laws, across 30 birth cohorts in three data sets, and because graduate degree completion, which should be changed little by compulsory attendance, indeed moves little with quarter of birth. Beyond statistical significance, this sort of coherence is an essential part of what gives any IV first stage its strength.

References

- Anderson, T. W., & Rubin, H. (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics*, 20(1), 46–63. <https://doi.org/10.1214/aoms/1177730090>
- Andrews, I., & Armstrong, T. B. (2017). Unbiased instrumental variables estimation under known first-stage sign. *Quantitative Economics*, 8(2), 479–503. <https://doi.org/10.3982/QE700>
- 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. D. (1990). Lifetime earnings and the Vietnam era draft lottery: Evidence from social security administrative records. *American Economic Review*, 80(3), 313–336. <https://www.jstor.org/stable/10.2307/2006669>
- Angrist, J. D., & Evans, W. N. (1998). Children and their parents' labor supply: Evidence from exogenous variation in family size. *American Economic Review*, 88(3), 450–477. <https://www.jstor.org/stable/116844>
- Angrist, J. D., & Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. 90(430), 431–442. <https://doi.org/10.1080/01621459.1995.10476535>
- Angrist, J. D., & Krueger, A. B. (1991). Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics*, 106(4), 979–1014. <https://doi.org/10.2307/2937954>
- Angrist, J. D., & Krueger, A. B. (1999). Empirical strategies in labor economics. In O. C. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (pp. 1277–1366). Elsevier. [https://doi.org/10.1016/S1573-4463\(99\)03004-7](https://doi.org/10.1016/S1573-4463(99)03004-7)
- Angrist, J. D., & Lavy, V. (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *The Quarterly Journal of Economics*, 114(2), 533–575. <https://doi.org/10.1162/003355399556061>
- Bekker, P. A. (1994). Alternative approximations to the distributions of instrumental variable estimators. *Econometrica*, 62(3), 657–681. <https://doi.org/10.2307/2951662>
- 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>

- Card, D. (2001). Estimating the return to schooling: Progress on some persistent econometric problems. *Econometrica*, 69(5), 1127–1160. <https://doi.org/10.1111/1468-0262.00237>
- Chernozhukov, V., & Hansen, C. (2008). The reduced form: A simple approach to inference with weak instruments. *Economics Letters*, 100(1), 68–71. <https://doi.org/10.1016/j.econlet.2007.11.012>
- Feller, W. (1968). *An introduction to probability theory and its application* (3rd ed., Vol. 1). Wiley.
- Hall, A. R., Rudebusch, G. D., & Wilcox, D. W. (1996). Judging instrument relevance in instrumental variables estimation. *International Economic Review*, 37(2), 283. <https://doi.org/10.2307/2527324>
- Hanushek, E. A. (1986). The economics of schooling: Production and efficiency in public schools. *Journal of Economic Literature*, 24(3), 1141–1177. <https://www.jstor.org/stable/2725865>
- Imbens, G. W., & Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2), 467–475. <https://doi.org/10.2307/2951620>
- Keane, M., & Neal, T. (2021). *A practical guide to weak instruments* (Working Paper). SSRN. <https://doi.org/10.2139/ssrn.3846841>
- Krueger, A. B. (1999). Experimental estimates of education production functions. *The Quarterly Journal of Economics*, 114(2), 497–532. <https://doi.org/10.1162/003355399556052>
- Lang, K. (1993). *Ability bias, discount rate bias, and the return to education* [Unpublished manuscript, Boston University].
- Lee, D. S., McCrary, J., Moreira, M. J., & Porter, J. (2020). *Valid t-ratio inference for IV*. arXiv: [2010.05058](https://arxiv.org/abs/2010.05058).
- Leeb, H., & Pötscher, B. M. (2005). Model selection and inference: Facts and fiction. *Econometric Theory*, 21(1), 21–59.
- Mikusheva, A., & Sun, L. (2021). Inference with many weak instruments. *Review of Economic Studies*.
- Moreira, M. J. (2009). Tests with correct size when instruments can be arbitrarily weak. *Journal of Econometrics*, 152(2), 131–140. <https://doi.org/10.1016/j.jeconom.2009.01.012>
- Newey, W. K., & West, K. D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 28(3), 777–787. <https://doi.org/10.2307/2526578>
- Owen, D. B. (1980). A table of normal integrals. *Communications in Statistics - Simulation and Computation*, 9(4), 389–419. <https://doi.org/10.1080/03610918008812164>
- Richardson, D. H. (1968). The exact distribution of a structural coefficient estimator. *Journal of the American Statistical Association*, 63(324), 1214–1226. <https://doi.org/10.1080/01621459.1968.10480921>
- Staiger, D., & Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3), 557–586. <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 Rothenberg* (pp. 80–108). Cambridge University Press. <https://doi.org/10.1017/CBO9780511614491.006>
- Young, A. (2021). *Leverage, heteroskedasticity and instrumental variables in practical application* [Unpublished working paper, LSE]. <https://personal.lse.ac.uk/YoungA/LeverageandIV.pdf>




Figure 1: Contour plot of the rejection rate of conventional t -test with nominal level $\alpha = 0.05$ as function of $E[F]$ and ρ . Panel (a) plots the unconditional rejection rate R_W . Panel (b) plots the rejection rate R_W^c conditional on $\hat{\pi} > 0$. See Appendix A.2 for computational details.

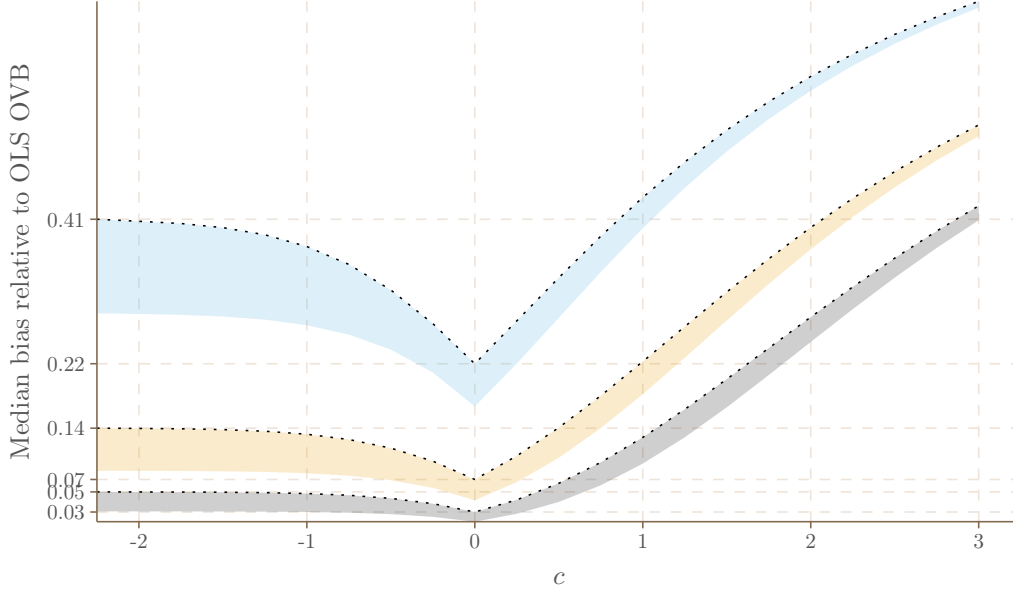


Figure 2: Relative median bias of $\hat{\beta}_{IV}$ conditional on $t_1 > c$, $\text{median}(\hat{\beta}_{IV} - \beta \mid t_1 > c) / |\beta_{WOLS} - \beta|$ as a function of cutoff c for selected values of $E[F]$. Shaded areas delineate the range of variation in the relative bias over possible values of ρ . Blue: $E[F] = 2$, orange $E[F] = 3.5$, gray: $E[F] = 5$. Dotted lines denote the limit of the relative bias as $\rho \rightarrow 0$. See Appendix A.6 for computational details.




Figure 3: Median bias of $\hat{\beta}_{IV}$ conditional on $t_1 > 0$ relative to unconditional median bias. Solid line denotes the ratio when $\rho = 1$, dotted line denotes the limit as $\rho \rightarrow 0$, and blue shaded area denotes the range of possible values over the possible values of ρ . Solid orange line marks a reference line at 0.5. See Appendix A.6 for computational details.

Table 1: Estimates and Endogeneity in Three IV Applications.

Sample	Outcome (1)	Treatment (2)	Instrument (3)	Estimates					
				OLS (4)	$\hat{\pi}$ (5)	$\hat{\delta}$ (6)	IV (7)	$\hat{\rho}$ (8)	$\hat{\rho}^*$ (9)
A. AK91									
Men born 1920–29 ($n = 247, 199$)	Log weekly wage [0.65]	Years of education [3.36]	Born in Q1 [0.43]	0.080 (0.0004)	-0.122 (0.016)	-0.009 (0.003)	0.072 (0.023)	0.043 (0.128)	0.044 (0.127)
Men born 1930–39 ($n = 329, 509$)	[0.68]	[3.28]	[0.43]	0.071 (0.0004)	-0.106 (0.013)	-0.011 (0.003)	0.105 (0.025)	-0.175 (0.121)	-0.171 (0.121)
B. AE98									
1980 Census, mothers aged 18–35 ($n = 394, 840$)	Worked for pay [0.50]	More than 2 kids [0.49]	Samesex [0.50]	-0.176 (0.002)	0.061 (0.001)	-0.007 (0.002)	-0.117 (0.025)	-0.058 (0.024)	-0.058 (0.024)
	Weeks worked [22.29]	[0.49]	[0.50]	-8.978 (0.071)	0.061 (0.001)	-0.340 (0.069)	-5.559 (1.118)	-0.075 (0.024)	-0.075 (0.024)
1990 Census, mothers aged 18–35 ($n = 380, 007$)	Worked for pay [0.48]	[0.48]	[0.50]	-0.164 (0.002)	0.062 (0.002)	-0.006 (0.002)	-0.092 (0.025)	-0.069 (0.024)	-0.069 (0.024)
	Weeks worked [22.87]	[0.48]	[0.50]	-8.649 (0.077)	0.062 (0.002)	-0.341 (0.073)	-5.462 (1.156)	-0.067 (0.024)	-0.067 (0.024)
C. AL99									
5th grade, full sample ($n = 2, 019$)	Reading [7.68]	Class size [6.55]	Maimonides' rule [6.11]	0.009 (0.034)	0.477 (0.041)	-0.125 (0.042)	-0.263 (0.094)	0.357 (0.075)	0.231 (0.045)
	Math [9.60]	[6.54]	[6.10]	0.036 (0.045)	0.477 (0.041)	-0.126 (0.056)	-0.264 (0.123)	0.315 (0.077)	0.189 (0.045)
5th grade, discontinuity sample ($n = 471$)	Reading [8.18]	[7.42]	[7.50]	-0.070 (0.050)	0.481 (0.057)	-0.197 (0.050)	-0.410 (0.118)	0.469 (0.092)	0.397 (0.067)
	Math [10.20]	[7.42]	[7.50]	0.090 (0.070)	0.481 (0.057)	-0.089 (0.072)	-0.185 (0.155)	0.322 (0.106)	0.249 (0.075)

Notes: This table reports IV and OLS estimates replicating the AK91, AE98, and AL99 studies discussed in the text. For each study, the table reports IV and OLS estimates from multiple samples, as well as the corresponding first-stage estimates, $\hat{\pi}$, reduced-form estimates, $\hat{\delta}$, and estimates of the endogeneity parameter, $\hat{\rho}$. The endogeneity estimate denoted $\hat{\rho}^*$ is computed assuming homoskedasticity. Standard errors appear in parentheses. These are robust for AK91 and AE98, and clustered on school for AL99. Standard deviations for the outcome, treatment, and instrument are reported in columns 1–3 in square brackets.

Appendix A Derivations and Proofs

The appendix uses the notation $\tilde{\beta}_{IV} = (\hat{\beta}_{IV} - \beta)/|\beta_{WOLS} - \beta|$ and $\tilde{\beta}_U = (\hat{\beta}_{IV} - \beta)/|\beta_{WOLS} - \beta|$ to denote the IV and the unbiased estimator, after centering and scaling by the weak-IV OVB of OLS. Also, we let $\omega = \rho/\sqrt{1 - \rho^2}$.

A.1 Estimating ρ

We estimate ρ (defined in eq. (8)) using first-stage and IV estimates, and the associated first-stage, reduced-form and IV standard errors. To see how this works, rewrite eq. (5) as:

$$\sigma_{\hat{\delta}\hat{\pi}} = \frac{\sigma_{\hat{\pi}}^2 \hat{\beta}_{IV}^2 - \hat{\pi}^2 \hat{\sigma}_{IV}^2 + \sigma_{\hat{\delta}}^2}{2\hat{\beta}_{IV}}. \quad (\text{A.1})$$

With this in hand for $\sigma_{\hat{\delta}\hat{\pi}}$, endogeneity can be computed as the sample analog of eq. (8), replacing β with $\hat{\beta}_{IV}$. The resulting endogeneity estimator is:

$$\hat{\rho} = \frac{\sigma_{\hat{\pi}}}{|\hat{\pi}|\hat{\sigma}_{IV}} \times (\sigma_{\hat{\delta}\hat{\pi}}/\sigma_{\hat{\pi}}^2 - \hat{\beta}_{IV}).$$

Under the normal model in eq. (3), this estimator depends on the data only through $\hat{\beta}_{IV}$, with the derivative given by $\partial\hat{\rho}/\partial\hat{\beta}_{IV} = \hat{\sigma}_{IV}^{-1} \cdot (\hat{\rho}^2 - 1)/|t_1|$. Hence, the delta-method standard error for $\hat{\rho}$ is simply $(1 - \hat{\rho}^2)/|t_1|$. Paralleling concerns with finite-sample coverage of the usual confidence interval for β , we might worry that confidence intervals for ρ based on delta-method standard errors suffer from undercoverage if endogeneity is high and the instruments are weak. We therefore also compute confidence sets for ρ by inverting the AR statistic, which are immune to this concern. Since ρ is monotone decreasing in β , if the AR confidence set is given by $[\beta_\ell, \beta_u]$, this leads to the confidence set for ρ given by $[\varrho(\beta_u), \varrho(\beta_\ell)]$, where $\varrho(\beta) = \frac{\sigma_{\hat{\pi}}}{\sqrt{\sigma_{\hat{\delta}}^2 - 2\beta\sigma_{\hat{\delta}\hat{\pi}} + \sigma_{\hat{\pi}}^2\beta^2}} \times (\sigma_{\hat{\delta}\hat{\pi}}/\sigma_{\hat{\pi}}^2 - \beta)$, and by $[-1, \varrho(\beta_u)] \cup [\varrho(\beta_\ell), 1]$ if the AR confidence set takes the form $(-\infty, \beta_\ell] \cup [\beta_u, \infty)$. As noted in the text, intervals computed in this way are virtually indistinguishable from the those based on the delta-method standard errors reported in Table 1.

A.2 t -Test Rejection Rates

This section writes the rejection probabilities of the t -test as an integral indexed by $(E[F], \rho)$. Stock and Yogo (2005) use Monte Carlo methods to compute unconditional rejection probabilities in a similar setup. The calculation described here is much faster. More importantly, it allows us to easily compute both unconditional rejection rates and rejection rates conditional on sign-screening.

Using eq. (8), and the fact that $\beta_{WOLS} - \beta$ and ρ have the same sign, we may write t_{AR} as

$$t_{AR} = \frac{(\hat{\delta} - \hat{\pi}\beta)|\rho|}{\sigma_{\hat{\pi}}|\beta_{WOLS} - \beta|}. \quad (\text{A.2})$$

Consequently,

$$\tilde{\beta}_{IV} = \frac{\hat{\delta} - \beta\hat{\pi}}{\sigma_{\hat{\pi}}t_1|\beta_{WOLS} - \beta|} = \frac{t_{AR}}{|\rho|t_1}. \quad (\text{A.3})$$

Thus,

$$t_W = \frac{\text{sign}(t_1)t_{AR}}{\sqrt{\frac{\sigma_{\hat{\delta}}^2/\sigma_{\hat{\pi}}^2 - 2\beta_{WOLS}\beta + \beta^2}{(\beta_{WOLS} - \beta)^2}\rho^2 + \frac{t_{AR}^2}{t_1^2} - 2\rho\frac{t_{AR}}{t_1}}} = \frac{\text{sign}(t_1)t_{AR}}{\sqrt{1 + t_{AR}^2/t_1^2 - 2\rho t_{AR}/t_1}} \quad (\text{A.4})$$

where the first equality uses eq. (A.3) and the definition of β_{WOLS} , and the second equality uses eq. (8). This expression for t_W implies that conditional on t_1 , the rejection region $\{|t_W| \geq z_{1-\alpha/2}\}$ is quadratic in t_{AR} . Solving this quadratic inequality implies that the rejection region is given by

$$t_{AR} \in \begin{cases} \emptyset & \text{if } t_1^2 \leq (1 - \rho^2)z_{1-\alpha/2}^2, \\ [a_1, a_2] & \text{if } (1 - \rho^2)z_{1-\alpha/2}^2 \leq t_1^2 \leq z_{1-\alpha/2}^2, \\ (-\infty, a_2) \cup (a_1, \infty) & \text{if } t_1^2 \geq z_{1-\alpha/2}^2 \end{cases}$$

where

$$a_1 = \frac{\rho z_{1-\alpha/2}^2 t_1 - |t_1| z_{1-\alpha/2} \sqrt{t_1^2 - (1 - \rho^2)z_{1-\alpha/2}^2}}{z_{1-\alpha/2}^2 - t_1^2},$$

$$a_2 = \frac{\rho z_{1-\alpha/2}^2 t_1 + |t_1| z_{1-\alpha/2} \sqrt{t_1^2 - (1 - \rho^2)z_{1-\alpha/2}^2}}{z_{1-\alpha/2}^2 - t_1^2}.$$

Note that $\text{cor}(t_{AR}, t_1) = \rho$, so that

$$P(t_{AR} \leq x \mid t_1) = \Phi((x - \rho(t_1 - E[t_1]))/\sqrt{1 - \rho^2}). \quad (\text{A.5})$$

Thus, conditional on t_1 , the rejection probability is given by

$$\begin{aligned} P(|t_W| \geq z_{1-\alpha/2} \mid t_1) &= (P(t_{AR} \leq a_2 \mid t_1) - P(t_{AR} \leq a_1 \mid t_1)) \mathbb{I}\{t_1^2 \geq z_{1-\alpha/2}^2(1 - \rho^2)\} \\ &\quad + \mathbb{I}\{t_1^2 \geq z_{1-\alpha/2}^2\} \\ &= f(t_1; E[t_1], \rho) \mathbb{I}\{t_1^2 \geq (1 - \rho^2)z_{1-\alpha/2}^2\} + \mathbb{I}\{t_1^2 \geq z_{1-\alpha/2}^2\}, \end{aligned} \quad (\text{A.6})$$

where

$$f(t_1; E[t_1], \rho) = \Phi\left(\frac{a_2 - \rho(t_1 - E[t_1])}{\sqrt{1 - \rho^2}}\right) - \Phi\left(\frac{a_1 - \rho(t_1 - E[t_1])}{\sqrt{1 - \rho^2}}\right).$$

Since $t_1 \sim \mathcal{N}(E[t_1], 1)$, the rejection probability conditional on $t_1 \geq c$ is therefore given by

$$P(|t_W| \geq z_{1-\alpha/2} \mid t_1 \geq c) = \frac{\int_c^\infty (\mathbb{I}\{\frac{t_1^2}{1-\rho^2} \geq z_{1-\alpha/2}^2\} f(t_1; E[t_1], \rho) + \mathbb{I}\{t_1^2 \geq z_{1-\alpha/2}^2\}) \phi(t_1 - E[t_1]) dt_1}{\Phi(E[t_1] - c)}.$$

The unconditional rejection probability R_W obtains by setting $c = -\infty$. The rejection probability conditional on sign screening, R_W^c , obtains by setting $c = 0$. The rejection contours in Figure 1

evaluate the above expression as a function of $(\rho, E[t_1])$ by numerical integration.

A.3 Proof of Theorem 1

The distribution of $\tilde{\beta}_{IV}$ conditional on t_1 can be written as

$$P_{\omega, E[t_1]}(\tilde{\beta}_{IV} \leq x | t_1) = \begin{cases} P(t_{AR} \leq |\rho|t_1x | t_1) = \Phi(\omega[E[t_1] - (1 - \text{sign}(\omega)x)t_1]) & \text{if } t_1 \geq 0, \\ P(t_{AR} \geq |\rho|t_1x | t_1) = \Phi(-\omega[E[t_1] - (1 - \text{sign}(\omega)x)t_1]) & \text{if } t_1 < 0, \end{cases} \quad (\text{A.7})$$

where the first equality uses eq. (A.3), and the second equality follows from eq. (A.5). Observe that since $P_{-\omega, E[t_1]}(\tilde{\beta}_{IV} \leq x | t_1) = 1 - P_{\omega, E[t_1]}(\tilde{\beta}_{IV} \leq -x | t_1)$, the distribution is symmetric in ω . It therefore suffices to consider $\omega > 0$.

Let $p_c(x) = P_{\omega, E[t_1]}(\tilde{\beta}_{IV} \leq x | t_1 > c)$ denote the distribution of $\tilde{\beta}_{IV}$ conditional on $t_1 > c$ (this shorthand notation ignores the dependence on ω and $E[t_1]$). By eq. (A.7),

$$p_c(x) = \begin{cases} \frac{1}{\Phi(E[t_1] - c)} \int_{c - E[t_1]}^{\infty} \Phi(y(z)) \phi(z) dz & \text{if } c \geq 0, \\ \frac{1}{\Phi(E[t_1] - c)} \left[\int_{c - E[t_1]}^{-E[t_1]} \Phi(-y(z)) \phi(z) dz + \int_{-E[t_1]}^{\infty} \Phi(y(z)) \phi(z) dz \right] & \text{if } c < 0, \end{cases}$$

where $y(z) = \omega[xE[t_1] - (1 - x)z]$. Observe that for all c , the conditional median of $\tilde{\beta}_{IV}$, denoted $m_c = m_c(\omega)$, is smaller than 1. In particular, for $c \geq 0$,

$$p_c(1) = \frac{\Phi(\omega E[t_1])}{\Phi(E[t_1] - c)} \int_{c - E[t_1]}^{\infty} \phi(z) dz = \Phi(\omega E[t_1]) \geq 1/2, \quad (\text{A.8})$$

while if $c < 0$,

$$\begin{aligned} p_c(1) &= \frac{1}{\Phi(E[t_1] - c)} [\Phi(E[t_1] - c) - \Phi(E[t_1]) + \Phi(\omega E[t_1])(-\Phi(E[t_1] - c) + 2\Phi(E[t_1]))] \\ &\geq \frac{1}{\Phi(E[t_1] - c)} \left[\Phi(E[t_1] - c) - \Phi(E[t_1]) + \frac{1}{2}(2\Phi(E[t_1]) - \Phi(E[t_1] - c)) \right] = \frac{1}{2}. \end{aligned}$$

We now show that the m_c is minimized at $c = 0$. If $c \geq 0$, by Leibniz rule,

$$\frac{\partial p_c(x)}{\partial c} = \frac{\phi(E[t_1] - c)}{\Phi(E[t_1] - c)} [p_c(x) - \Phi(\omega[E[t_1] - (1 - x)c])],$$

which is negative for $x \leq 1$, since it follows from eq. (A.7) that $p_c(x) \leq \frac{1}{\Phi(E[t_1] - c)} \int_{c - E[t_1]}^{\infty} \Phi(\omega[E[t_1] - (1 - x)c]) \phi(z) dz = \Phi(\omega[E[t_1] - (1 - x)c])$. Therefore, m_c is decreasing for positive c .

Now consider $c < 0$. From eq. (A.7), it follows that

$$p_c(x) = \frac{1}{\Phi(E[t_1] - c)} \left[\int_{c - E[t_1]}^{-E[t_1]} \Phi(\omega[-E[t_1] + (1 - x)(z + E[t_1])]) \phi(z) dz + \Phi(E[t_1]) p_0(x) \right].$$

Suppose $p_0(m_c) < 1/2$. Then it follows from the preceding display that for $x \leq 1$,

$$\begin{aligned} \frac{1}{2} &< \frac{1}{\Phi(E[t_1] - c)} \left[\int_{c-E[t_1]}^{-E[t_1]} \Phi(\omega[-E[t_1] + (1-x)(z + E[t_1])]) \phi(z) dz + \Phi(E[t_1]) \frac{1}{2} \right] \\ &\leq \frac{1}{\Phi(E[t_1] - c)} \left[\frac{1}{2} \int_{c-E[t_1]}^{-E[t_1]} \phi(z) dz + \Phi(E[t_1]) \frac{1}{2} \right] = \frac{1}{2}, \end{aligned}$$

where the second inequality uses the fact that $\Phi(\omega[-E[t_1] + (1-x)(z + E[t_1])]) \leq \Phi(-E[t_1]\omega) \leq \Phi(0)$ over the range of integration. Hence, $p_0(m_c) \geq 1/2$, which implies that $m_0 \leq m_c$. By the proof of Theorem 2 $m_0 \geq 0$. Thus, setting $c = 0$ minimizes $|m_c|$ for all c , as claimed.

A.4 Proof of Theorem 2

The proof begins by characterizing the distribution of $\tilde{\beta}_{IV}$ conditional on $t_1 > 0$. The previous proof establishes that the conditional median for this distribution is less than 1, so it suffices to consider $p_0(x)$ for $x \leq 1$. By the mean value theorem, for some $\tilde{\omega} = \tilde{\omega}(x, \omega) \in [0, \omega]$,

$$\begin{aligned} p_0(x) &= \Phi(0) + \frac{\omega}{\Phi(E[t_1])} \int_{-E[t_1]}^{\infty} ((x-1)z + xE[t_1]) \phi(\tilde{\omega}((x-1)z + xE[t_1])) \phi(z) dz \\ &= \frac{1}{2} + \frac{\omega}{\tilde{\omega}^2(1-x)\Phi(E[t_1])} \int_{-\infty}^{\tilde{\omega}E[t_1]} y \phi(y) \phi(a + by) dy, \end{aligned}$$

where the second line uses the change of variables $y = \tilde{\omega}xE[t_1] - \tilde{\omega}(1-x)z$, and we let $a = xE[t_1]/(1-x)$, $b = -\frac{1}{\tilde{\omega}(1-x)}$. By line 111 of Table 1 in Owen (1980),

$$\int x \phi(x) \phi(a + bx) = \frac{\phi(a/t)}{t^2} \left[-\phi(tx + ab/t) - \frac{ab}{t} \Phi(tx + ab/t) \right], \quad t = \sqrt{1 + b^2}. \quad (\text{A.9})$$

Applying this result to the preceding display then yields

$$\begin{aligned} p_0(x) &= \frac{1}{2} + \frac{\omega}{\tilde{\omega}^2(1-x)\Phi(E[t_1])} \frac{\phi(a/t)}{1 + b^2} \left[-\phi(t\tilde{\omega}E[t_1] + ab/t) - \frac{ab}{\sqrt{1 + b^2}} \Phi(t\tilde{\omega}E[t_1] + ab/t) \right] \\ &= \frac{1}{2} + \frac{\omega}{\Phi(E[t_1])} \frac{\phi(a/t)(1-x)}{\tilde{\omega}^2(1-x)^2 + 1} \left[\frac{x}{|1-x|} \frac{E[t_1]}{\tilde{g}(x, \tilde{\omega})} \Phi(E[t_1]g(x, \tilde{\omega})) - \phi(E[t_1]g(x, \tilde{\omega})) \right], \end{aligned}$$

where $g(x, \tilde{\omega}) = \frac{\tilde{\omega}^2|1-x| + \text{sign}(1-x)}{\sqrt{\tilde{\omega}^2(1-x)^2 + 1}}$, and $\tilde{g}(x, \tilde{\omega}) = \sqrt{\tilde{\omega}^2(1-x)^2 + 1}$. When evaluated at the conditional median, m_0 , the expression in square brackets must equal zero by definition of the median. Therefore, $m_0 > 0$, and since we also know from eq. (A.8) that $m_0 < 1$, the conditional median must satisfy

$$m_0 = \frac{1}{\frac{E[t_1]}{\tilde{g}(m_0, \tilde{\omega}(m_0, \omega))} \frac{\Phi(E[t_1]g(m_0, \tilde{\omega}(m_0, \omega)))}{\phi(E[t_1]g(m_0, \tilde{\omega}(m_0, \omega)))} + 1}, \quad (\text{A.10})$$

We have

$$\frac{E[t_1]}{\tilde{g}} \frac{\Phi(E[t_1]g)}{\phi(E[t_1]g)} \geq \frac{E[t_1]}{\tilde{g}} \frac{\Phi(E[t_1]\tilde{g})}{\phi(E[t_1]\tilde{g})}, \quad \text{and} \quad \frac{E[t_1]}{\tilde{g}} \frac{\Phi(E[t_1]\tilde{g})}{\phi(E[t_1]\tilde{g})} \geq E[t_1] \frac{\Phi(E[t_1])}{\phi(E[t_1])} \quad \text{if } E[t_1] \geq 0.84.$$

Here the first inequality follows because $\Phi(x)/\phi(x)$ is increasing in x , and $g \geq \tilde{g}$, and the second inequality follows because $\frac{\Phi(x)}{x\phi(x)}$ is increasing for $x \geq 0.84$, and $\tilde{g} \geq 1$. Therefore,

$$m_0 \leq \frac{\phi(E[t_1])}{E[t_1]\Phi(E[t_1]) + \phi(E[t_1])} = \lim_{\omega \downarrow 0} m_0(\omega),$$

where the equality follows since the right-hand side of eq. (A.10) converges to $\frac{\phi(E[t_1])}{E[t_1]\Phi(E[t_1]) + \phi(E[t_1])}$ as $\omega \rightarrow 0$.

We now prove the claims concerning the unconditional distribution of $\tilde{\beta}_{IV}$. By arguments in the proof of Theorem 1, the median is smaller than 1, so it suffices to consider $p_{-\infty}(x)$ for $x \leq 1$. By arguments as in the conditional case,

$$\begin{aligned} p_{-\infty}(x) &= \frac{1}{2} + \frac{\omega}{\tilde{\omega}^2(1-x)} \left[\int_{-\infty}^{\tilde{\omega}E[t_1]} y\phi(y)\phi(a+by)dy - \int_{\tilde{\omega}E[t_1]}^{\infty} y\phi(y)\phi(a+by)dy \right] \\ &= \frac{1}{2} + \frac{\omega}{\tilde{\omega}^2(1-x)} \frac{\phi(a/t)}{1+b^2} \left[-2\phi(t\tilde{\omega}E[t_1] + ab/t) - 2\frac{ab}{t}\Phi(t\tilde{\omega}E[t_1] + ab/t) + \frac{ab}{t} \right] \\ &= \frac{1}{2} + \frac{\omega}{\tilde{\omega}^2(1-x)} \frac{\phi(a/t)}{1+b^2} \left[-2\phi(E[t_1]g(\tilde{\omega}, x)) + 2\frac{x}{1-x} \frac{E[t_1]}{\tilde{g}} \Phi(E[t_1]g(\tilde{\omega}, x)) - \frac{x}{1-x} \frac{E[t_1]}{\tilde{g}} \right]. \end{aligned}$$

Here the first line follows by the mean value theorem, where $\tilde{\omega} = \tilde{\omega}(x, \omega) \in [0, \omega]$, the second line uses eq. (A.9), and the last line follows by algebra. When evaluated at $x = m_{-\infty}$, the expression in square brackets must equal zero by definition of the median. Therefore, $m_{-\infty} > 0$, and it must satisfy

$$m_{-\infty} = \frac{1}{\frac{E[t_1]}{\tilde{g}} \frac{\Phi(E[t_1]g) - 1/2}{\phi(E[t_1]g)} + 1} \tag{A.11}$$

Now,

$$\frac{E[t_1]}{\tilde{g}} \frac{\Phi(E[t_1]g) - 1/2}{\phi(E[t_1]g)} \geq \frac{E[t_1]}{\tilde{g}} \frac{\Phi(E[t_1]\tilde{g}) - 1/2}{\phi(E[t_1]\tilde{g})} \geq E[t_1] \frac{\Phi(E[t_1]) - 1/2}{\phi(E[t_1])}.$$

Here the first inequality follows because $\Phi(x)/\phi(x)$ is increasing in x , and $g \geq \tilde{g}$, and the second inequality follows because $\frac{\Phi(x) - 1/2}{x\phi(x)}$ is increasing for $x > 0$. As a result,

$$m_{-\infty} \leq \frac{\phi(E[t_1])}{E[t_1](\Phi(E[t_1]) - 1/2 + \phi(E[t_1]))} = \lim_{\omega \downarrow 0} m_{-\infty}(\omega),$$

where the equality follows since the right-hand side of eq. (A.11) converges to $\frac{\phi(E[t_1])}{E[t_1](\Phi(E[t_1]) - 1/2 + \phi(E[t_1]))}$ as $\omega \rightarrow 0$.

A.5 Proof of Theorem 3

We may write

$$\tilde{\beta}_U = t_1\mu(t_1)\tilde{\beta}_{IV} + (1 - t_1\mu(t_1))\text{sign}(\rho) = \mu(t_1)\frac{t_{AR}}{|\rho|} + (1 - t_1\mu(t_1))\text{sign}(\rho) \quad (\text{A.12})$$

where the first equality follows from eq. (16), and the fact that $\beta_{WOLS} - \beta$ and ρ have the same sign, and the second equality applies eq. (A.3).

Since $E[t_{AR} | t_1] = \rho(t_1 - E[t_1])$, the relative bias conditional on t_1 is given by

$$E[\tilde{\beta}_U | t_1] = \text{sign}(\rho) [1 - E[t_1]\mu(t_1)].$$

By arguments analogous to those in the proof of Lemma 2.1 in Andrews and Armstrong (2017), we have

$$\begin{aligned} E[t_1]E[\mu(t_1) | t_1 > 0] &= \frac{E[t_1]}{\Phi(E[t_1])} \int_{t=0}^{\infty} \frac{1 - \Phi(t)}{\phi(t)} \phi(t - E[t_1]) dt = \frac{e^{-\frac{E[t_1]^2}{2}}}{\Phi(E[t_1])} \int_{t=0}^{\infty} (1 - \Phi(t)) E[t_1] e^{E[t_1]t} dt \\ &= \frac{e^{-E[t_1]^2/2}}{\Phi(E[t_1])} \left\{ \left[e^{E[t_1]t} (1 - \Phi(E[t_1])) \right]_{t=0}^{\infty} + \int_{t=0}^{\infty} \phi(t) e^{E[t_1]t} dt \right\} \\ &= \frac{1}{\Phi(E[t_1])} \left[-\frac{1}{2} e^{-E[t_1]^2/2} + \int_{t=0}^{\infty} \phi(t - E[t_1]) dt \right] = -\frac{1}{2} \frac{e^{-E[t_1]^2/2}}{\Phi(E[t_1])} + 1, \end{aligned}$$

where the first line uses the definition of the Mills' ratio, the second line uses integration by parts, and the third follows by completing the square. It therefore follows that

$$\frac{E[\hat{\beta}_U - \beta | t_1 > 0]}{\beta_{WOLS} - \beta} = \frac{1}{2} \frac{e^{-E[t_1]^2/2}}{\Phi(E[t_1])}.$$

The second claim follows by an analogous argument.

A.6 Median Bias Comparisons

To evaluate the relative median bias of $\hat{\beta}_{IV}$ as a function of both $E[F]$ and ρ conditional on $t_1 \geq c$, we first evaluate its distribution

$$P(\tilde{\beta}_{IV} \leq x | t_1 \geq c; \rho, E[t_1]) = \frac{1}{\Phi(E[t_1] - c)} \int_{c - E[t_1]}^{\infty} f_{IV}(z; x, E[t_1], \omega) \phi(z) dz \quad (\text{A.13})$$

by numerical integration. Here we use the formula $f_{IV}(z; x, \rho, E[t_1]) = \Phi(\omega[\text{sign}(z + E[t_1])E[t_1] - (1 - x)|z + E[t_1]|])$ from eq. (A.7) for the cdf conditional on $z = t_1 - E[t_1]$. We then numerically solve for the median. The shaded regions in Figure 2 correspond to the range of the absolute value of the relative median bias as ρ varies between -1 and 1 . Similarly, the shaded regions in Figure 3 show how the median bias conditional on $t_1 \geq 0$ relative to the unconditional median bias (that sets $c = -\infty$) varies with ρ .

To compare the relative median bias to that of $\hat{\beta}_U$, it suffices to consider $\rho > 0$, since the distribu-

tions of $\hat{\beta}_U$ and $\hat{\beta}_{IV}$ are symmetric in ρ . By eq. (A.12), it follows that for $t_1 > 0$,

$$P(\tilde{\beta}_U \leq x \mid t_1; \omega) = P\left(\tilde{\beta}_{IV} \leq x - (1-x)\frac{(1-t_1\mu(t_1))}{t_1\mu(t_1)} \mid t_1; \omega\right),$$

which for $x < 1$ is smaller than $P(\tilde{\beta}_{IV} \leq x \mid t_1; \omega)$. Since the median of $\tilde{\beta}_{IV}$ conditional on $t_1 > 0$ is smaller than 1, it follows that the conditional median bias of $\tilde{\beta}_{IV}$ is always smaller than that of $\tilde{\beta}_U$.

To compare the relative magnitudes of the median biases, we compute the relative median bias of $\tilde{\beta}_U$ analogously to that of $\tilde{\beta}_{IV}$, except we replace f_{IV} in eq. (A.13) with $f_U(z; x, E[t_1], \rho) = \Phi(\omega[E[t_1] - (1 - \text{sign}(\omega)x)/\mu(E[t_1] + z)])$ (it follows from eqs. (A.5) and (A.12) that this is the cdf $\tilde{\beta}_U$ conditional on $z = t_1 - E[t_1]$). We then compute the ratio $\text{median}_{E[t_1], \rho}(\tilde{\beta}_U \mid t_1 > 0) / \text{median}_{E[t_1], \rho}(\tilde{\beta}_{IV} \mid t_1 > 0)$ of the median biases on a fine grid of values of $(\rho, E[t_1])$. This ratio is greater than 2 if $E[F] \geq 2$, and greater than 3 if $E[F] \geq 3$, regardless of the value of ρ . Likewise, comparison of the ratio of the conditional and unconditional median IV bias, $\text{median}_{E[t_1], \rho}(\tilde{\beta}_{IV} \mid t_1 > 0) / \text{median}_{E[t_1], \rho}(\tilde{\beta}_{IV}) = \text{median}_{E[t_1], \rho}(\hat{\beta}_{IV} - \beta \mid t_1 > 0) / \text{median}_{E[t_1], \rho}(\hat{\beta}_{IV} - \beta)$ shows that the ratio lies between 0.5 and 0.525 for $E[t_1] \geq 1.5$, regardless of the value of ρ .