ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

https://www.econometricsociety.org/

Econometrica, Vol. 89, No. 3 (May, 2021), 1449–1469

SALVAGING FALSIFIED INSTRUMENTAL VARIABLE MODELS

MATTHEW A. MASTEN
Department of Economics, Duke University

ALEXANDRE POIRIER
Department of Economics, Georgetown University

Econometrica, Vol. 89, No. 3 (May, 2021), 1449–1469

SALVAGING FALSIFIED INSTRUMENTAL VARIABLE MODELS

MATTHEW A. MASTEN
Department of Economics, Duke University

ALEXANDRE POIRIER
Department of Economics, Georgetown University

What should researchers do when their baseline model is falsified? We recommend reporting the set of parameters that are consistent with minimally nonfalsified models. We call this the *falsification adaptive set* (FAS). This set generalizes the standard baseline estimand to account for possible falsification. Importantly, it does not require the researcher to select or calibrate sensitivity parameters. In the classical linear IV model with multiple instruments, we show that the FAS has a simple closed-form expression that only depends on a few 2SLS coefficients. We apply our results to an empirical study of roads and trade. We show how the FAS complements traditional overidentification tests by summarizing the variation in estimates obtained from alternative nonfalsified models.

KEYWORDS: Instrumental variables, nonparametric identification, partial identification, sensitivity analysis.

1. INTRODUCTION

MANY MODELS USED IN EMPIRICAL RESEARCH ARE FALSIFIABLE, in the sense that there exists a population distribution of the observable data which is inconsistent with the model. With finite samples, researchers often use specification tests to check whether their baseline model is falsified. Abstracting from sampling uncertainty, the population versions of such specification tests have a persistent problem: What should researchers do when their baseline model is falsified?

In this paper, we provide a constructive way for researchers to salvage a falsified baseline model. To do this, we consider continuous relaxations of the baseline assumptions of concern. By sufficiently weakening the assumptions, a falsified baseline model becomes nonfalsified. We define the *falsification frontier* as the set of smallest relaxations of the baseline model which are not falsified. Our main recommendation is that researchers report estimates of the identified set for the parameter of interest under the assumption that the true model lies on this frontier. We call this the *falsification adaptive set* (FAS). This set collapses to the baseline identified set or point estimand when the baseline model

Matthew A. Masten: matt.masten@duke.edu
Alexandre Poirier: alexandre.poirier@georgetown.edu
First arXiv draft: December 30, 2018. This paper was presented at Auburn University, UC San Diego, Texas A&M, Duke, Columbia, Cornell, the Yale MacMillan-CSAP workshop, the University of Mannheim, Penn State, the 2018 Incomplete Models conference at Northwestern University, the 2018 and 2019 Southern Economic Association Meetings, the 2019 IAAE conference, the 2019 Georgetown Center for Economic Research Mini-Conference on Non-Standard Methods in Econometrics, the 2019 CeMMAP UCL/Vanderbilt Conference on Advances in Econometrics, the 29th Annual Meeting of the Midwest Econometrics Group, the 2019 Greater New York Area Econometrics Colloquium, the 2020 Winter Meeting of the Econometric Society, and the 2020 World Congress of the Econometric Society. We thank audiences at those seminars and conferences, the referees, as well as Federico Bugni, Tim Christensen, Allan Collard-Wexler, Joachim Freyberger, Chuck Manski, Arnaud Maurel, Francesca Molinari, Adam Rosen, Pedro Sant'Anna, and Alex Torgovitsky for helpful conversations and comments. We thank Paul Diegert and Peiran Xiao for excellent research assistance. Masten thanks the National Science Foundation for research support under Grant 1943138.

 https://doi.org/10.3982/ECTA17969

is not falsified. When the baseline model is falsified, this set expands to include all parameter values consistent with the data and a model which is relaxed just enough to make it nonfalsified. Hence the FAS generalizes the standard baseline estimand to account for possible falsification. Importantly, researchers do not need to select or calibrate sensitivity parameters to compute the falsification adaptive set. We formally define these concepts in Section 2.

To illustrate this method, we study the classical constant coefficients linear model with multiple instruments in Section 3. We relax instrument exclusion by allowing the instruments to have some direct effect on outcomes. We show that the FAS has a particularly simple closed-form expression, depending only on the value of a handful of 2SLS regression coefficients. We then use our results in an empirical study of roads and trade in Section 4. We show that the FAS is an informative complement to traditional overidentification tests: The FAS summarizes the range of estimates obtained from alternative models which are not falsified by the data. Thus the FAS reflects the model uncertainty that arises from a falsified baseline model.

Related Literature

Our paper builds on several large literatures. Manski and Pepper (2018) presented identified sets under relaxations of two assumptions, which can be used to construct a falsification adaptive set in their model; see their Table 2. Ramsahai (2012) studied a heterogeneous treatment effect IV model with continuous relaxations of instrument exogeneity and informally notes that the model is not falsified if exogeneity is sufficiently relaxed. Machado, Shaikh, and Vytlacil (2019) studied a heterogeneous treatment effects IV model and formally define what we call a falsification point. More recently, Andrews and Kwon (2019) introduced a scalar slack parameter to define minimally nonfalsified moment inequality models, which do not nest our results.

Our technical results build on a large literature on sensitivity analysis in linear IV models, including Fisher (1961), Angrist and Krueger (1994), Altonji, Elder, and Taber (2005), Small (2007), Conley, Hansen, and Rossi (2012), Ashley (2009), Kraay (2012), Ashley and Parmeter (2015), and van Kippersluis and Rietveld (2017, 2018). There is also a large literature on falsification and sensitivity analysis in heterogeneous effect IV models; see Flores and Chen (2018) and Swanson et al. (2018) for excellent surveys.

Several recent papers use local asymptotics to study sensitivity to misspecification. For example, see Andrews, Gentzkow, and Shapiro (2017), Bonhomme and Weidner (2018), and Armstrong and Kolesár (2021). This approach assumes the baseline model is approximately correct, in the sense that the magnitude of model misspecification is similar to the magnitude of sampling uncertainty. We focus on clearly falsified models, where it is known that the model is not approximately correct. Hence we use a global approach, which does not rely on linking the size of model misspecification to the size of sampling uncertainty. We compare this local misspecification approach with ours in more detail at the end of Section 4.

2. SALVAGING FALSIFIED MODELS

In this section we consider a general falsifiable model. We use this model to precisely define the falsification frontier and falsification adaptive set. In Section 3 we illustrate these general concepts in the classical linear instrumental variable model.

2.1. *Measuring the Extent of Falsification*

Let W be a vector of observed random variables. Let $\mathcal{F}$ denote the set of all possible cdfs for W. A *model* is a set of underlying parameters which generate the observed distribution F_W and restrictions on those parameters. These parameters could be infinite dimensional. This definition of a model suffices for our purposes. See Section 2 of Matzkin (2007), for example, for a more formal definition. A given model is *falsifiable* if there are some distributions F_W which could not have been generated by the model. When the data follows one of these population distributions, we say the model is falsified (equivalently, refuted). Let $\mathcal{F}_{\mathrm{f}}$ denote the set of cdfs F_W which falsify the model. Let $\mathcal{F}_{\mathrm{nf}}$ denote the set of cdfs F_W which do not falsify the model.

Suppose we begin with a falsifiable baseline model. Suppose this model has L assumptions which we think might be false. For each assumption $\ell \in \{1, \ldots, L\}$, we define a class of assumptions indexed by a parameter δ_ℓ such that the assumption is imposed for $\delta_\ell = 0$, the assumption is not imposed for δ_ℓ equal to its maximum feasible value $\delta_\ell^{\max}$, and the assumption is partially imposed for $\delta_\ell \in (0, \delta_\ell^{\max})$. Two common values of $\delta_\ell^{\max}$ are 1 and $+\infty$. These assumptions must be nested in the sense that for $\delta'_\ell \geq \delta_\ell$, assumption δ'_ℓ is weaker than assumption δ_ℓ.

Consider the model which imposes assumptions $\delta = (\delta_1, \ldots, \delta_L)$. Let $\mathcal{F}_{\mathrm{nf}}(\delta)$ denote the set of joint distributions of the data which are not falsified by this model. In particular, $\mathcal{F}_{\mathrm{nf}}(0_L)$ denotes the set of joint distributions of the data which are not falsified by the baseline model. Since we assumed the baseline model is falsifiable, $\mathcal{F}_{\mathrm{nf}}(0_L)$ is a strict subset of $\mathcal{F}$. Suppose further that the model which does not impose any of the L assumptions is not falsifiable.

Recall that F_W denotes the observed distribution of the data. Suppose $F_W \notin \mathcal{F}_{\mathrm{nf}}(0_L)$, so that the baseline model is falsified. Partition $\mathcal{D} = [0, \delta_1^{\max}] \times \cdots \times [0, \delta_L^{\max}]$ into two sets:

$$\mathcal{D}_{\mathrm{f}} = \{\delta \in \mathcal{D} : F_W \notin \mathcal{F}_{\mathrm{nf}}(\delta)\} \quad \text{and} \quad \mathcal{D}_{\mathrm{nf}} = \{\delta \in \mathcal{D} : F_W \in \mathcal{F}_{\mathrm{nf}}(\delta)\}.$$

$\mathcal{D}_{\mathrm{f}}$ is the set of all assumptions which are falsified. $\mathcal{D}_{\mathrm{nf}}$ is the set of all assumptions which are not falsified. For simplicity assume $\mathcal{D}_{\mathrm{nf}}$ is closed, which holds in our Section 3 analysis.

DEFINITION 1: The *falsification frontier* is the set

$$\mathrm{FF} = \{\delta \in \mathcal{D} : \delta \in \mathcal{D}_{\mathrm{nf}} \text{ and for any other } \delta' < \delta, \text{ we have } \delta' \in \mathcal{D}_{\mathrm{f}}\},$$

where $\delta' < \delta$ means that $\delta'_\ell \leq \delta_\ell$ for all $\ell \in \{1, \ldots, L\}$ and $\delta'_m < \delta_m$ for some $m \in \{1, \ldots, L\}$.

That is, the falsification frontier is the set of assumptions which are not falsified, but if strengthened in any component, leads to a falsified model. When $L = 1$, the falsification frontier is a singleton called the *falsification point*: For all δ below that point, the model is falsified while for all δ above that point the model is not falsified.

2.2. *The Falsification Adaptive Set*

Let $\Theta_I(\delta)$ denote the identified set for a parameter of interest $\theta \in \Theta$, given the model which imposes the assumptions δ. When $\delta \in \mathcal{D}_{\mathrm{f}}$, δ is below the falsification frontier. In this case, the identified set $\Theta_I(\delta)$ is empty. When $\delta \in \mathcal{D}_{\mathrm{nf}}$, δ is on or above the falsification frontier. In this case, the identified set $\Theta_I(\delta)$ is nonempty.

DEFINITION 2: Call

$$\bigcup_{\delta \in \mathrm{FF}} \Theta_I(\delta)$$

the *falsification adaptive set*.

The falsification adaptive set is the identified set for the parameter of interest when the true model satisfies one of the assumptions on the falsification frontier. When the baseline model is not falsified, this set collapses to $\Theta_I(0_L)$, the baseline identified set (which may be a singleton). This is what researchers typically report when their baseline model is not falsified. When the baseline model is falsified, however, the falsification adaptive set expands to account for uncertainty about which assumption along the frontier is true. Hence this set generalizes the standard baseline estimand to account for possible falsification.

3. THE CLASSICAL LINEAR MODEL WITH MULTIPLE INSTRUMENTS

In this section we illustrate our method in the classical linear instrumental variable model. While many kinds of falsifiable assumptions have been considered in the literature, we focus on the classical case where variation from two or more instruments is used to falsify the model.

3.1. *Model and Identification*

Let $Y(x, z)$ denote potential outcomes defined for values $(x, z) \in \mathbb{R}^{K+L}$. Assume

$$Y(x, z) = x'\beta + z'\gamma + U, \tag{1}$$

where β is an unknown constant K-vector, γ is an unknown constant L-vector, and U is an unobserved random variable. Let X be an observed K-vector of endogenous variables. Throughout we suppose X does not contain a constant. Hence U absorbs any nonzero constant intercept. Let Z be an observed L-vector of potentially invalid instruments. We observe the outcome $Y = Y(X, Z)$. For simplicity, we have omitted any additional known exogenous covariates W in equation (1); they can be easily included via partialing out.

Equation (1) imposes homogeneous treatment effects. We also maintain the following relevance and sufficient variation assumptions throughout this section.

ASSUMPTION A1—Relevance: *The $L \times K$ matrix* $\mathrm{cov}(Z, X)$ *has rank K.*

ASSUMPTION A2—Sufficient variation: *The $L \times L$ matrix* $\mathrm{var}(Z)$ *is invertible.*

A1 implies the order condition $L \geq K$. When there is just one endogenous variable ($K = 1$), A1 only requires $\mathrm{cov}(X, Z_\ell) \neq 0$ for at least one instrument. Other instruments may have zero correlation. In this case, these other instruments provide additional falsifying power. We discuss this further below. If one instrument is an affine combination of the others, A2 does not hold. In this case, just remove affinely dependent instruments until $\mathrm{var}(Z)$ is invertible.

The classical model imposes two more assumptions:

ASSUMPTION A3—Exogeneity: $\mathrm{cov}(Z_\ell, U) = 0$ *for all* $\ell \in \{1, \dots, L\}$.

ASSUMPTION A4—Exclusion: $\gamma_\ell = 0$ *for all* $\ell \in \{1, \dots, L\}$.

A1–A4 imply that the coefficient vector β is point identified and equals the two stage least squares (2SLS) estimand. Furthermore, these assumptions imply well-known overidentifying conditions. The following proposition gives these conditions when there is just a single endogenous variable.

PROPOSITION 1: *Suppose* $K = 1$. *Suppose the joint distribution of* (Y, X, Z) *is known and satisfies* A1 *and* A2. *Then the model* (1) *with* A3 *and* A4 *is not falsified if and only if*

$$\operatorname{cov}(Y, Z_m) \operatorname{cov}(X, Z_\ell) = \operatorname{cov}(Y, Z_\ell) \operatorname{cov}(X, Z_m) \tag{2}$$

for all m *and* ℓ *in* $\{1, \dots, L\}$.

When all instruments are relevant, so that $\operatorname{cov}(X, Z_\ell) \neq 0$ for all $\ell \in \{1, \dots, L\}$, equation (2) can be written as

$$\frac{\operatorname{cov}(Y, Z_m)}{\operatorname{cov}(X, Z_m)} = \frac{\operatorname{cov}(Y, Z_\ell)}{\operatorname{cov}(X, Z_\ell)}.$$

That is, the linear IV estimand must be the same for all instruments Z_ℓ. This result is the basis for the classical test of overidentifying restrictions (Anderson and Rubin (1949), Sargan (1958), Hansen (1982)). Suppose the distribution of (Y, X, Z) is such that the model is falsified. This happens when at least one of our model assumptions fails: (a) homogeneous treatment effects, (b) linearity in X, (c) instrument exogeneity, or (d) instrument exclusion.

Here we maintain the homogeneous treatment effects assumption. We consider models with heterogeneous treatment effects in our working paper Masten and Poirier (2020). We also maintain linearity of potential outcomes in x, which could include known functions of covariates like quadratic terms. In principle, our analysis can be extended to allow for relaxations of this functional form assumption, but we leave this to future work.

We thus focus on failure of (c) instrument exogeneity or (d) instrument exclusion as reasons for falsifying the baseline model. These are two different substantive assumptions. Mathematically, however, the same technical analysis can be used to relax both assumptions. For simplicity, here we formally maintain the exogeneity assumption A3 and focus on failure of the exclusion assumption A4.

In general, it is difficult to define a meaningful and tractable class of relaxations of one's baseline assumptions. In the linear model, however, there is a natural way to relax the exclusion restriction. Specifically, we use the following class of assumptions.

ASSUMPTION 4′—Partial exclusion: *There are known constants* $\delta_\ell \geq 0$ *such that* $|\gamma_\ell| \leq \delta_\ell$ *for all* $\ell \in \{1, \dots, L\}$.

A4′ bounds the magnitude of the direct effect of each instrument on the outcome by known constants. This kind of relaxation of the baseline instrumental variable assumptions was previously considered by Small (2007) and Conley, Hansen, and Rossi (2012); also see Angrist and Krueger (1994) and Bound, Jaeger, and Baker (1995). Although the instruments may have a direct causal effect on outcomes, the model may nonetheless continue to be falsified for sufficiently small values of the components in δ. For sufficiently large values, however, the model will not be falsified. To characterize the falsification frontier, we begin by deriving the identified set for β as a function of δ.

THEOREM 1: *Suppose A1–A3 and A4′ hold. Suppose the joint distribution of* (Y, X, Z) *is known. Then*

$$\mathcal{B}(\delta) = \{b \in \mathbb{R}^K : -\delta \leq \text{var}(Z)^{-1}(\text{cov}(Z, Y) - \text{cov}(Z, X)b) \leq \delta\} \tag{3}$$

is the identified set for β. *Here the inequalities are componentwise. The model is falsified if and only if this set is empty.*

The identified set $\mathcal{B}(\delta)$ depends on the data via two terms:

$$\underset{(L\times 1)}{\psi} \equiv \text{var}(Z)^{-1}\,\text{cov}(Z, Y) \quad \text{and} \quad \underset{(L\times K)}{\Pi} \equiv \text{var}(Z)^{-1}\,\text{cov}(Z, X).$$

ψ is the reduced form regression of Y on Z. Π is the first stage of X on Z. If we demeaned (Y, X, Z), then we would have $\psi = \mathbb{E}(ZZ')^{-1}\mathbb{E}(ZY)$ and $\Pi = \mathbb{E}(ZZ')^{-1}\mathbb{E}(ZX')$. Theorem 1 shows that the identified set is the intersection of L pairs of parallel half-spaces in $\mathbb{R}^K$. When $\delta = 0_L$, this identified set becomes the intersection of L hyperplanes in $\mathbb{R}^K$. In this case, β is point identified when $\text{cov}(Z, Y) = \text{cov}(Z, X)b$ for a unique $b \in \mathbb{R}^K$. If $\text{cov}(Z, Y) \neq \text{cov}(Z, X)b$ for all $b \in \mathbb{R}^K$, then the baseline model $\delta = 0_L$ is falsified.

Increasing the components of δ leads to a weakly larger identified set. Furthermore, there always exists a δ with large enough components so that $\mathcal{B}(\delta)$ is nonempty. We characterize the set of such δ below. Before that, we show that the identified set can be written as simple intersection bounds when there is a single endogenous variable.

COROLLARY 1: *Suppose the assumptions of Theorem 1 hold. Suppose* $K = 1$. *Then*

$$\mathcal{B}(\delta) = \bigcap_{\ell=1}^{L} B_\ell(\delta_\ell)$$

is the identified set for β, *where*

$$B_\ell(\delta_\ell) = \begin{cases} \left[\dfrac{\psi_\ell}{\pi_\ell} - \dfrac{\delta_\ell}{|\pi_\ell|}, \dfrac{\psi_\ell}{\pi_\ell} + \dfrac{\delta_\ell}{|\pi_\ell|}\right] & \text{if } \pi_\ell \neq 0, \\ \mathbb{R} & \text{if } \pi_\ell = 0 \text{ and } 0 \in [\psi_\ell - \delta_\ell, \psi_\ell + \delta_\ell], \\ \emptyset & \text{if } \pi_\ell = 0 \text{ and } 0 \notin [\psi_\ell - \delta_\ell, \psi_\ell + \delta_\ell]. \end{cases} \tag{4}$$

Here Π *is an* L*-vector and* π_ℓ *is its* ℓ*th component.*

To interpret this result, first consider an instrument Z_ℓ with a zero first-stage coefficient, $\pi_\ell = 0$. If Z_ℓ has a sufficiently strong relationship with the outcome, so that $\psi_\ell \pm \delta_\ell$ does not contain zero, then the model is falsified. Furthermore, in this case falsification can be solely attributed to the assumption that $|\gamma_\ell| \leq \delta_\ell$ for this specific ℓ. This is similar to what is sometimes called the "zero first-stage test" (e.g., see Slichter (2014) and the references therein). When this relationship with the outcome is sufficiently small, however, Z_ℓ unsurprisingly has no falsifying or identifying power for β.

Next consider a relevant instrument Z_ℓ, so $\pi_\ell \neq 0$. To interpret Corollary 1 in this case, we use the following lemma.

LEMMA 1: *Suppose* $K = 1$. *Let* $\widetilde{X}_\ell = (Z_1, \ldots, Z_{\ell-1}, X, Z_{\ell+1}, \ldots, Z_L)$. *Let* e_ℓ *be the* $L \times 1$ *vector of zeros with a one in the* ℓ*th component*. *Suppose* $\pi_\ell \neq 0$. *Suppose* $\text{cov}(Z, \widetilde{X}_\ell)$ *is invertible*. *Then*

$$\frac{\psi_\ell}{\pi_\ell} = e_\ell' \text{cov}(Z, \widetilde{X}_\ell)^{-1} \text{cov}(Z, Y).$$

This lemma shows that ψ_ℓ / π_ℓ is the population 2SLS coefficient on X using Z_ℓ as the excluded instrument and using the remaining instruments $Z_{-\ell}$ as controls. Thus the identified set $\mathcal{B}(\delta)$ is the intersection of intervals around these 2SLS coefficients using one relevant instrument at a time and controlling for the rest.

Finally, consider the baseline case where $\delta = 0_L$. Corollary 1 implies that $\mathcal{B}(0_L)$ is nonempty if and only if

$$\frac{\psi_m}{\pi_m} = \frac{\psi_\ell}{\pi_\ell}$$

for any $m, \ell \in \{1, \ldots, L\}$ with $\pi_m, \pi_\ell \neq 0$ and $\psi_j = 0$ when $\pi_j = 0$. Moreover, in this case $\mathcal{B}(0_L)$ is a singleton equal to this common value. In this case—when the baseline model is not falsified—we also have

$$\frac{\psi_\ell}{\pi_\ell} = \frac{\text{cov}(Y, Z_\ell)}{\text{cov}(X, Z_\ell)}$$

for all $\ell \in \{1, \ldots, L\}$. That is, ψ_ℓ / π_ℓ equals the population 2SLS coefficient on X using Z_ℓ as an instrument and *not* including $Z_{-\ell}$ as controls. This equality of single instrument 2SLS coefficients with and without controls for the other instruments is an alternative characterization of the classical overidentifying conditions from Proposition 1. Note that, when these overidentifying conditions do not hold, it can be shown that the baseline 2SLS estimand is not necessarily in the identified set $\mathcal{B}(\delta)$. Consequently, it will not necessarily be in the falsification adaptive set that we describe below. Instead, as shown via Corollary 1 and Lemma 1, the identified set depends on the estimands ψ_ℓ / π_ℓ, which use $Z_{-\ell}$ as controls to allow for possible exclusion failures.

3.2. *The Falsification Frontier*

So far we have characterized the identified set for β given a fixed value of δ, the upper bound on the violation of the exclusion restriction. We now consider the possibility that this identified set is empty when $\delta = 0_L$, so that the baseline model is falsified. Our next result characterizes the falsification frontier, the minimal set of δ's which lead to a nonempty identified set. Here we focus on the single endogenous regressor case. We extend this result to multiple endogenous regressors in the Online Supplementary Material (Masten and Poirier (2021)).

PROPOSITION 2: *Suppose* A1–A3 *hold*. *Suppose the joint distribution of* (Y, X, Z) *is known*. *Suppose* $K = 1$. *Then the falsification frontier is the set*

$$\text{FF} = \left\{ \delta \in \mathbb{R}^L_{\geq 0} : \delta_\ell = |\psi_\ell - b\pi_\ell|, \ell = 1, \ldots, L, b \in \left[\min_{\ell=1,\ldots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\ldots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell} \right] \right\}. \quad (5)$$

In the proof we show that this set satisfies our Definition 1 of the falsification frontier. Specifically, any $\delta \in \text{FF}$ maps to a nonempty identified set, and strengthening any of the assumptions for a given $\delta \in \text{FF}$ leads to an empty identified set.

3.3. *The Falsification Adaptive Set*

Next we characterize the falsification adaptive set, which is the identified set for β under the assumption that one of the points on the falsification frontier is true. Here we again focus on the single endogenous regressor case. We generalize our analysis to multiple endogenous regressors in the Online Supplementary Material (Masten and Poirier (2021)).

THEOREM 2: *Suppose A1–A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Suppose $K = 1$. Then*

$$\bigcup_{\delta \in \mathrm{FF}} \mathcal{B}(\delta) = \left[\min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell} \right] \tag{6}$$

is the falsification adaptive set.

We first sketch the proof of this result and then discuss its implications. It follows from two main steps: First, the identified set $\mathcal{B}(\delta)$ is a singleton for any $\delta \in \mathrm{FF}$ (see Lemma 2 in the Appendix). Second, each of these singleton sets corresponds to an element in the interval on the right-hand side of equation (6) (follows using Proposition 2). Thus we obtain the entire interval by taking the union of all these singletons.

Our main recommendation is that researchers report estimates of the falsification adaptive set. Theorem 2 shows that, in the classical linear model we consider here, this set has an exceptionally simple form. Most importantly, no δ's appear on the right-hand side of equation (6). This implies that *we can obtain the falsification adaptive set without precomputing the falsification frontier or selecting any sensitivity parameters*. Furthermore, it is very simple to compute, since it just requires running L different 2SLS regressions.

In this model, we can also immediately see how this set adapts to falsification of the baseline model. When the baseline model is not false, $\psi_m/\pi_m = \psi_\ell/\pi_\ell$ for all $m, \ell \in \{1, \dots, L\}$ with nonzero π_m and π_ℓ. In this case, the falsification adaptive set collapses to the singleton equal to the common value. This is the same point estimand researchers would usually present when their baseline model is not falsified. As the baseline model becomes more falsified, the values of ψ_ℓ/π_ℓ become more different, and the falsification adaptive set expands. Thus the size of this set reflects the severity of baseline falsification.

3.4. *Estimation and Inference*

In finite samples, researchers can present sample analog estimates of the falsification adaptive set, along with corresponding confidence sets. Our characterization of the FAS in equation (6) requires that we first screen for irrelevant instruments. It is not clear how to best do this. We present a first pass approach, but leave a detailed analysis to future work.

Let $\{Y_i, X_i, Z_i\}_{i=1}^n$ be an iid sample from the distribution of (Y, X, Z). Let

$$\mathcal{L}_{\mathrm{rel}} = \{\ell \in \{1, \dots, L\} : \pi_\ell \neq 0\}$$

be the set of indices corresponding to relevant instruments. Estimate this set by

$$\widehat{\mathcal{L}}_{\mathrm{rel}} = \{\ell \in \{1, \dots, L\} : F_\ell \geq C_n\}.$$

F_ℓ is the first-stage F-statistic when considering Z_ℓ as an instrument and $Z_{-\ell}$ as controls. C_n is a cutoff that diverges as the sample size grows. The assumptions in Proposition 3 below ensure that $\widehat{\mathcal{L}}_{\text{rel}}$ is consistent for $\mathcal{L}_{\text{rel}}$.

Let $\widehat{b}_\ell$ be an estimator of ψ_ℓ/π_ℓ, the 2SLS coefficient on X using Z_ℓ as the excluded instrument and $Z_{-\ell}$ as controls. We estimate the falsification adaptive set by

$$\widehat{\text{FAS}} = \left[\min_{\ell \in \widehat{\mathcal{L}}_{\text{rel}}} \widehat{b}_\ell, \max_{\ell \in \widehat{\mathcal{L}}_{\text{rel}}} \widehat{b}_\ell\right].$$

We use this estimator in our empirical analysis of Section 4. The following result gives conditions under which this estimator is consistent for the FAS.

PROPOSITION 3: *For all* $\ell \in \{1, \dots, L\}$ *suppose*:
1. $\widehat{b}_\ell \xrightarrow{p} \psi_\ell/\pi_\ell$ *when* $\pi_\ell \neq 0$.
2. $F_\ell \xrightarrow{d} \chi^2_1$ *when* $\pi_\ell = 0$.
3. $F_\ell/n \xrightarrow{p} \kappa_\ell$ *when* $\pi_\ell \neq 0$, *where* $\kappa_\ell > 0$ *is some positive constant*.
4. *As* $n \to \infty$, $C_n \to \infty$ *and* $C_n = o(n)$.

Let FAS *denote the interval in equation* (6). *Let* d_H *denote the Hausdorff distance. Then* $d_H(\widehat{\text{FAS}}, \text{FAS}) \xrightarrow{p} 0$.

Assumptions 1–3 hold under standard assumptions on random sampling, existence of moments, and the existence of consistent variance estimators used within F_ℓ. Assumption 4 requires that C_n grows slowly enough to ensure that relevant instruments are kept in $\widehat{\mathcal{L}}_{\text{rel}}$ with probability approaching one. In our empirical analysis we choose $C_n = 10$ as our default cutoff, although we sometimes consider other cutoffs, or a sequence of cutoffs. Inference can be done by using a version of the delta method discussed in Fang and Santos (2019), noting that the min and max are directionally differentiable mappings. We leave a detailed analysis of both the choice of the cutoff and procedures for inference to future work.

4. EMPIRICAL APPLICATION: ROADS AND TRADE

In this section we apply our results from Section 3 to the empirical analysis of roads and trade by Duranton, Morrow, and Turner (2014). They consider a dataset of 66 regions ("cities") in the United States. Their treatment variable is the log number of kilometers of interstate highways within a city in 2007. This variable directly affects the cost of leaving a city and, therefore, the cost of exporting from a city: It is easier to export from a city with many kilometers of interstate highways passing through it. Their outcome variable is a measure of how much that city exports. They consider two different ways of measuring exports: Weight (in tons) and value (in dollars). We focus on the weight measure for brevity. They begin by estimating a gravity equation relating the weight of a city's exports to other cities with the highway distance between those cities, both measured in 2007. This equation includes a fixed effect for the exporting city. The estimate of this fixed effect is their main outcome variable. They call this variable the "propensity to export weight." Thus their main goal is to estimate the causal effect of within city highways on the propensity to export weight.

We cannot learn this causal effect by simply regressing the propensity to export weight on within city highways since there is a classic simultaneity problem. We expect that building highways within the city will boost exports. But high export cities may also build more highways to facilitate their existing exports. The authors solve this problem by instrumenting for the number of kilometers of within city highways. They consider three different instruments:

1. *Railroads*: The log number of kilometers of railroads in the city in 1898.
2. *Exploration*: A measure of the quantity of historical exploration routes that passed through the city.
3. *Plan*: The log number of kilometers of highway in the city, according to a planned highway construction map approved by the federal government in 1947. Baum-Snow (2007) had previously used this instrument, and provides a detailed history.

The authors raise concerns about validity of all three instruments. Although they address these concerns with various controls, these controls may still not perfectly fix failures of exogeneity, exclusion, or both. Hence the authors lean on overidentification, stating that "Using different instruments, for which threats to validity differ, allows for informative over-identification tests" (p. 700). With this motivation, we next present the results.

Results

First consider Table I. Panel A reproduces columns 1–4 of Table 5 in Duranton, Morrow, and Turner (2014). These are their main results. In particular, they are interested in the coefficient on log highway km, the log number of highway kilometers within the city. This coefficient represents their estimate of the causal effect of roads on trade. Here it is estimated by 2SLS, using railroads, exploration, and plan as instruments. At the 10% level, the standard test of overidentifying restrictions passes in the two longest specifications, fails in the second specification, and marginally passes in the first specification. Also note that these specifications do not include all of the additional controls the authors consider; they include those in separate analyses, which we discuss later (our Table III).

We add the estimated falsification adaptive set to these baseline results. This is the last row of panel A. There are two things to notice. First, except for the last specification, none of the 2SLS estimates are within the estimated FAS. This is not surprising since it can be shown that the baseline 2SLS estimand does not need to be inside the FAS. Second, the estimated FAS magnitudes are all generally smaller than the 2SLS point estimates.

To better understand how we computed the estimated FAS, and how to interpret it, next consider Table II. Columns 1–3 include the same baseline controls as column 3 in Table I while columns 4–6 include the same baseline controls as column 4 in Table I. The only difference is that we no longer use all three variables (plan, railroads, exploration) as instruments. Instead, in panel A, we use only one of these variables as an instrument and we ignore the other two variables. Panel A reproduces columns 4–6 from Table 6 in Duranton, Morrow, and Turner (2014). The authors used these results as their main robustness check. They argue that the three estimates 0.38, 0.64, and 0.34 from columns 4–6, panel A, Table II are consistent with their baseline estimates of 0.47 and 0.39 from columns 3 and 4, panel A, Table I.

However, omitting an invalid instrument can lead to omitted variable bias. In this application we are concerned that some of the instruments may be invalid. Thus the alternative models of interest are those where one of the instruments is valid but the others are not. When computing results in these alternative models, the invalid instruments should be included as controls (see Lemma 1). Panel B shows these results. Here we use one instrument while controlling for the other two. For example, in column 1 we use plan as an instrument and control for railroads and exploration.

For brevity, here we only describe the results in columns 4–6. These results use the full baseline specification. Consider column 5, panel B. This result uses railroads as an

TABLE I

BASELINE 2SLS RESULTS FOR DURANTON, MORROW, AND TURNER (2014): THE EFFECT OF HIGHWAYS ON EXPORT WEIGHT.[a]

	Dependent Variable: Export Weight			
	(1)	(2)	(3)	(4)
Panel A. Plan, exploration, and railroads used as instruments				
log highway km	1.13 (0.14)	0.57 (0.16)	0.47 (0.14)	0.39 (0.12)
log employment		0.52 (0.11)	0.69 (0.39)	0.47 (0.33)
Market access (export)		−0.45 (0.14)	−0.65 (0.14)	−0.63 (0.11)
log 1920 population			−0.38 (0.25)	−0.29 (0.23)
log 1950 population			1.00 (0.39)	0.65 (0.38)
log 2000 population			−0.74 (0.48)	−0.20 (0.45)
log % manuf. emp.				0.64 (0.12)
First-stage F stat.	97.5	90.3	80.0	84.8
Overid. p-value	0.10	0.043	0.15	0.31
$\widehat{\text{FAS}}$	[0.49, 0.86]	[−0.32, 0.28]	[−0.26, 0.31]	[0.18, 0.42]
Panel B. Plan and exploration used as instruments, controlling for railroads				
log highway km	0.79 (0.24)	0.17 (0.20)	0.21 (0.15)	0.23 (0.14)
log 1898 railroad km	0.33 (0.15)	0.33 (0.12)	0.25 (0.12)	0.16 (0.10)
First-stage F stat.	61.1	65.4	77.8	82.2
Overid. p-value	0.64	0.51	0.48	0.72

[a]*Notes*: 66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. Panel A reproduces columns 1–4 of Table 5 in Duranton, Morrow, and Turner (2014). It also shows the estimated falsification adaptive set. Panel B uses only two of their instruments, controlling for the other.

instrument, controlling for plan and highway. Unlike the uncontrolled result from panel A, railroads is a very weak instrument. Hence we ignore the result using railroads alone, as discussed in Section 3.4. Next consider column 4. Here we use plan as the instrument, controlling for the other two. Despite these controls, plan is still a strong instrument. The estimated effect 0.18 is roughly half as large as the estimate from panel A, 0.38. It is also no longer statistically significant at any conventional level. Next consider column 6. Here we use exploration as the instrument, controlling for the other two. Exploration continues to be a strong instrument with these controls. The estimated effect 0.42 in panel B is similar to the effect from panel A, 0.34. It is no longer statistically significant, however.

Putting these coefficient estimates together gives us the estimated FAS, [0.18, 0.42]. The endpoints of this set correspond to point estimates from alternative models which maintain validity of only one instrument at a time. The interior of this set corresponds to

TABLE II

THE EFFECT OF CONTROLLING FOR UNUSED INSTRUMENTS.[a]

	(1) Plan	(2) Railroads	(3) Exploration	(4) Plan	(5) Railroads	(6) Exploration
Panel A. Without controlling for other instruments						
log highway km	0.49 (0.15)	0.83 (0.25)	0.12 (0.31)	0.38 (0.13)	0.64 (0.22)	0.34 (0.19)
log % manuf. emp.				0.64 (0.12)	0.60 (0.13)	0.65 (0.13)
First-stage F stat.	141	45.2	14.8	130	40.8	23.8
Panel B. Controlling for other instruments						
log highway km	0.31 (0.22)	4.09 (4.09)	−0.26 (0.71)	0.18 (0.21)	3.65 (4.16)	0.42 (0.52)
log % manuf. emp.				0.63 (0.12)	0.36 (0.38)	0.61 (0.12)
log 1898 railroad km	0.21 (0.12)		0.24 (0.11)	0.18 (0.11)		0.16 (0.11)
log 1528–1850 exploration	−0.053 (0.077)	−0.40 (0.36)		0.025 (0.065)	−0.32 (0.40)	
log 1947 highway km		−2.09 (2.50)	0.32 (0.46)		−1.90 (2.48)	−0.13 (0.36)
First-stage F stat.	59.6	1.54	29.1	54.8	1.27	27.0
$\widehat{\text{FAS}}$ for this specification		[−0.26, 0.31]			[0.18, 0.42]	

[a]*Notes*: 66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. All columns have employment, market access, and past populations as controls. Columns 4–6 also have manufacturing share of employment as controls. Panel A reproduces the results from columns 4–6 of Table 6 in Duranton, Morrow, and Turner (2014) while the estimates in panel B are new.

alternative models which relax validity of all instruments at once, but just enough to avoid falsification.

In panel B of Table II, we found that railroads is a weak instrument when controlling for the other two, and hence it yields the largest point estimates. Given this finding, one may also wonder how removing railroads as an instrument affects the baseline analysis. This is shown in panel B of Table I. All of the coefficients on log highway km are smaller, to the point that they are no longer statistically significant for all but the shortest specification. Moreover, the standard overidentification tests now all easily pass. (Note that these tests are only comparing results using plan and exploration as instruments.) However, the coefficients on railroads are statistically significant for all but the fourth column. This suggests that the full baseline model using all three instruments could be rejected, and also explains the source of the relatively small overidentification test p-values in panel A.

Thus far we have focused on the baseline results, which do not include all of the possible control variables that the authors discuss. Table III shows results with their additional controls. We begin with the full set of baseline control variables, as used in column 4 of Table I. We then add just one control. Each row corresponds to a different control. The last row shows the results that add all controls at once. Unlike the main baseline result, column 4 of Table I, here we only use one instrument at a time. Columns 1–3 use a single instrument, without controlling for the other two. These results reproduce Appendix Ta-

TABLE III

THE EFFECT OF CONTROLLING FOR UNUSED INSTRUMENTS, CONTINUED.[a]

Added variable		Without Other IV Controls (1) Plan	(2) Railroads	(3) Exploration	With Other IV Controls (4) Plan	(5) Railroads	(6) Exploration	$\widehat{FAS}$
Water	log highway km	0.34	0.66	0.24	0.13	3.96	0.30	
		(0.16)	(0.26)	(0.29)	(0.21)	(4.56)	(0.49)	[0.13, 0.30]
	F stat.	126	25.3	10.9	67.6	1.17	26.2	
Slope		0.39	0.57	0.46	0.20	3.86	0.66	
		(0.14)	(0.20)	(0.19)	(0.21)	(5.86)	(0.50)	[0.20, 0.66]
		133	44.5	22.6	60.3	0.65	25.6	
Census regions		0.32	0.62	0.36	−0.012	3.68	0.40	
		(0.14)	(0.17)	(0.20)	(0.24)	(3.24)	(0.64)	[−0.012, 0.40]
		122	58.3	22.6	39.4	1.59	10.7	
Percent college		0.29	0.56	0.41	0.013	3.64	0.78	
		(0.13)	(0.23)	(0.18)	(0.19)	(4.13)	(0.49)	[0.013, 0.78]
		116	36.6	29.5	47.9	1.16	25.2	
Income per capita		0.35	0.63	0.35	0.079	3.42	0.54	
		(0.14)	(0.22)	(0.18)	(0.21)	(3.42)	(0.47)	[0.079, 0.54]
		123	36.8	26.4	47.8	1.70	24.9	
Percent wholesale		0.41	0.59	0.49	0.22	2.71	0.75	
		(0.12)	(0.21)	(0.17)	(0.19)	(3.29)	(0.49)	[0.22, 0.75]
		136	39.3	23.2	54.4	1.38	25.4	
Traffic		0.42	1.00	0.34	0.23	5.57	0.40	
		(0.18)	(0.52)	(0.23)	(0.25)	(9.60)	(0.51)	[0.23, 0.40]
		79	13.4	44.6	44.3	0.43	26	
All		0.39	0.73	0.58	0.18	2.43	0.69	
		(0.18)	(0.35)	(0.28)	(0.26)	(3.01)	(0.67)	[0.18, 0.69]
		52.6	19.9	15.5	32.1	0.88	6.53	

[a]*Notes*: 66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. This table extends the analysis of Table II to consider specifications with additional control variables. Columns 1–3 reproduce the results in Appendix Table 6 of Duranton, Morrow, and Turner (2014), while the estimates in Columns 4–6, which add controls for the other instruments, are new.

ble 6 of Duranton, Morrow, and Turner (2014). Based on these results, the authors argue that "None of our main results is affected by these controls, even when we use our instruments individually" (p. 708). They also argue that using one instrument at a time is an "even more demanding exercise" than examining the effect of additional controls when using all three instruments (not shown here; see their Appendix Table 5). As we have discussed, however, omitting the invalid instruments may cause omitted variable bias. So in columns 4–6, we replicate columns 1–3, except now controlling for the other two instruments.

There are three main differences between the results with the instrument controls and those without. First, the railroads instrument is again very weak, leading to large coefficients. This informs our understanding of the results from columns 1–3, since there we observed that the coefficients in column 2 are always larger than those in columns 1 and 3, and are often substantially larger. Second, none of the results are statistically significant at conventional levels. Finally, the coefficients using plan as the instrument all become smaller once the other instruments are controlled for (column 4 versus column 1), while the coefficients using exploration as the instrument all become larger once the other instruments are controlled for (column 6 versus column 3). Thus, ignoring the results using railroads, the overall range of point estimates is larger. This is reflected in the estimated falsification adaptive sets, which are presented in the final column.

Overall, there are two main conclusions from our analysis: First, the evidence suggests that the railroads instrument is the most questionable, and should be used only as a control. Thus the estimates in panel B of Table I are arguably the most appropriate baseline results. Second, there is substantially more uncertainty in the magnitude of the causal effect of roads on trade than suggested by the original results of Duranton, Morrow, and Turner (2014). This is reflected in the various estimated falsification adaptive sets we present. In particular, the estimated FAS for the longest specification is $[0.18, 0.69]$; see Table III. Moreover, accounting for sampling uncertainty would only increase this range. That said, these results do not change the overall qualitative conclusions of the paper: All points in the estimated FAS for the longest specification are still positive, suggesting that the number of within city highways appears to positively affect propensity to export weight.

Comparison With the Andrews, Gentzkow, and Shapiro (2017) Approach

In this subsection we compare our approach with that of Andrews, Gentzkow, and Shapiro (2017). They study general moment equality models, while we focus on the linear IV model. For the linear IV model, in their example 4 they study the sensitivity of the 2SLS estimator to violations of exclusion or exogeneity of the same magnitude as the sampling uncertainty (proportional to $1/\sqrt{n}$). Under such data generating processes, and as in Conley, Hansen, and Rossi (2012), they show that the 2SLS estimator is consistent, but asymptotically biased. The asymptotic bias has the form $A\gamma$ where γ is a vector of sensitivity parameters and A is a matrix that is point identified from the data. Andrews, Gentzkow, and Shapiro (2017) recommended that authors report estimates of A, which allows readers to select a γ and compute their own local asymptotic bias correction for the 2SLS estimator.

Our empirical application has a single endogenous variable, log highway km. We are primarily interested in its coefficient. For a given choice of γ, it can be shown that the local asymptotic bias of the 2SLS estimator for this coefficient is $a_1\gamma_1 + \cdots + a_L\gamma_L$ where

$$a_\ell = \frac{\operatorname{cov}(Z_\ell, X_{\text{pred}}^{\perp W})}{\operatorname{var}(X_{\text{pred}}^{\perp W})}. \tag{7}$$

Here W is the vector of the included exogenous control variables, X_{pred} is the predicted value of the endogenous variable from the first stage regression of X on $(1, Z, W)$, and $X_{\text{pred}}^{\perp W}$ is the residual from the linear regression of X_{pred} on $(1, W)$. Thus the relevant elements a_ℓ of the matrix A are simply the coefficients on X_{pred} in a linear regression of Z_ℓ on $(1, X_{\text{pred}}, W)$ for each $\ell = 1, \dots, L$. In addition to reporting the elements of A,

TABLE IV

ANALYSIS OF THE LOCAL SENSITIVITY OF THE 2SLS ESTIMATOR FOR THE FOUR BASELINE SPECIFICATIONS IN TABLE I[a]

	Dependent Variable: Export Weight			
	(1)	(2)	(3)	(4)
Panel A. Estimates of $a_\ell/\text{stddev}(Z_\ell)$, the standardized elements of the matrix A				
Plan	1.80	2.36	2.46	2.46
Railroads	1.44	1.73	1.49	1.45
Exploration	1.09	1.48	1.55	1.69
Panel B. Point estimates of the coefficient on log highway km				
Baseline 2SLS	1.13	0.57	0.47	0.39
Bias ($\gamma_\ell = 1/\text{stddev}(Z_\ell)$)	0.53	0.69	0.68	0.69
Bias corrected 2SLS ($\gamma_\ell = 1/\text{stddev}(Z_\ell)$)	0.60	−0.11	−0.21	−0.30
Bias corrected 2SLS ($\gamma_\ell = -1/\text{stddev}(Z_\ell)$)	1.67	1.26	1.14	1.08

[a]*Notes*: $n = 66$ observations per column. $1/\sqrt{n} = 0.12$. In panel B, Bias is an estimate of $(a_1\gamma_1 + \cdots + a_L\gamma_L)/\sqrt{n}$.

Andrews, Gentzkow, and Shapiro (2017) also recommended comparing their magnitudes. As they note, this can be difficult because the units of this matrix depend on the units of the moments themselves. In our case, this means that the instrument ℓ with the largest value of $|a_\ell|$ is not necessarily the most important for the local asymptotic bias—it depends on the units of each Z_ℓ. To address this, we standardize the instruments as described below.

Table IV shows the empirical results. Here we focus on the four main baseline specifications, as reported in Table I. Panel A reports estimates of $a_\ell/\text{stddev}(Z_\ell)$ for each of the three instruments. We see that plan has the largest value of the three instruments. Given equation (7) above, this means that plan has the largest correlation with the predicted treatment, after controlling for covariates. This is consistent with what we reported in Table II, where plan had the largest first stage F-statistic conditional on the other two instruments. Note, however, that railroads and exploration both have roughly the same values in panel A. This suggests that they are equally important for the asymptotic bias of the 2SLS estimator. In contrast, our analysis in Table II suggested that, unlike exploration, railroads is a conditionally very weak instrument, and should possibly not be relied on.

In panel B we consider two possible choices of the local sensitivity parameters: $\gamma_\ell = \pm 1/\text{stddev}(Z_\ell)$. These choices can be interpreted as saying that if Z_ℓ increases by one standard deviation, then the direct effect of Z_ℓ on outcomes is $\pm 1/\sqrt{n}$. For these choices, the magnitude of the estimated asymptotic bias is quite large, leading to bias corrected 2SLS estimators which are both positive and negative for the main specification (column 4). This sensitivity analysis thus suggests that even the main qualitative results of Duranton, Morrow, and Turner (2014) are not robust to violations of exclusion of that magnitude. In contrast, the estimated FAS for the main specification (column 4 of Table I) still contains only positive numbers, suggesting that the qualitative results of Duranton, Morrow, and Turner (2014) *are* robust to exclusion violations that are sufficiently large to explain falsification of the baseline model.

Overall, this discussion highlights two key practical differences between our analysis and that of Andrews, Gentzkow, and Shapiro (2017). First, their approach is estimator specific: Different estimators of the same parameter can lead to different conclusions about sensitivity. Thus our comparisons of the relative sensitivity of exclusion violations for each of the three instruments above could change if we used an estimator other than 2SLS. In contrast, our approach is a population level identification analysis that does not depend on a specific choice of baseline estimator. Second, their approach ultimately requires researchers—either authors or readers—to choose the sensitivity parameter γ. In contrast, our approach can be thought of as leveraging falsification of the baseline model to automatically calibrate the parameter γ, by considering the minimal relaxations that make the instruments consistent with each other.

5. CONCLUSION

In this paper we suggested a constructive answer to the question "What should researchers do when their baseline model is falsified?" We recommend reporting estimates of the set of parameters that are consistent with minimally nonfalsified models. We call this the *falsification adaptive set* (FAS) because it generalizes the standard baseline estimand to account for possible falsification. We illustrated this recommendation in the classical linear instrumental variable model with multiple instruments. We showed that the FAS has a particularly simple closed-form expression, depending only on the value of a handful of 2SLS regression coefficients. Finally, we showed how to use our results in practice. There we discussed the importance of controlling for the possibly invalid instruments when considering alternative models. Overall, we showed that the FAS is an informative complement to traditional overidentification test p-values: It directly summarizes the range of estimates corresponding to nonfalsified alternative models.

APPENDIX: PROOFS FOR SECTION 3

PROOF OF PROPOSITION 1: Note that we assumed A1–A2 hold since they depend on observables only. Suppose equation (2) holds for all $m, \ell \in \{1, \dots, L\}$. We will construct a joint distribution $(Y, X, Z, \widetilde{U})$ and a parameter $\widetilde{\beta}$ consistent with the data, equation (1), and assumptions A3–A4.

By the relevance assumption A1, there exists an ℓ such that $\operatorname{cov}(X, Z_\ell) \neq 0$. Let $\widetilde{\beta} = \operatorname{cov}(Y, Z_\ell)/\operatorname{cov}(X, Z_\ell)$. Let $\widetilde{U} = Y - X\widetilde{\beta}$. For every $m \in \{1, \dots, L\}$,

$$\begin{aligned}
\operatorname{cov}(\widetilde{U}, Z_m) &= \operatorname{cov}(Y, Z_m) - \operatorname{cov}(X, Z_m)\widetilde{\beta} \\
&= \operatorname{cov}(Y, Z_m) - \operatorname{cov}(X, Z_m)\frac{\operatorname{cov}(Y, Z_\ell)}{\operatorname{cov}(X, Z_\ell)} \\
&= \frac{\operatorname{cov}(Y, Z_m)\operatorname{cov}(X, Z_\ell) - \operatorname{cov}(Y, Z_\ell)\operatorname{cov}(X, Z_m)}{\operatorname{cov}(X, Z_\ell)} \\
&= 0.
\end{aligned}$$

Thus A3 holds. A4 holds by definition of $\widetilde{U}$. Thus the model is not falsified.

Next suppose the model is not falsified. Then there exists a joint distribution of (Y, X, Z, U) and a value β consistent with the model assumptions, equation (1), and the data. By A3–A4, we have

$$\operatorname{cov}(Y, Z_\ell) = \beta \operatorname{cov}(X, Z_\ell) \tag{8}$$

for all $\ell \in \{1, \dots, L\}$. Suppose $\beta = 0$. Then $\operatorname{cov}(Y, Z_\ell) = 0$ for all ℓ, and hence equation (2) holds for all $m, \ell \in \{1, \dots, L\}$.

Suppose $\beta \neq 0$. Then multiplying equation (8) for ℓ by equation (8) for m gives

$$\operatorname{cov}(Y, Z_\ell) \times \big(\beta \operatorname{cov}(X, Z_m)\big) = \big(\beta \operatorname{cov}(X, Z_\ell)\big) \times \operatorname{cov}(Y, Z_m).$$

Divide by β to see that equation (2) holds for all $m, \ell \in \{1, \dots, L\}$. *Q.E.D.*

PROOF OF THEOREM 1: First we show that any value of β consistent with the model must lie in $\mathcal{B}(\delta)$. By the outcome equation (1) and instrument exogeneity (A3),

$$\operatorname{cov}(Z, Y) = \operatorname{cov}(Z, X)\beta + \operatorname{var}(Z)\gamma.$$

By A2,

$$\gamma = \operatorname{var}(Z)^{-1}\big(\operatorname{cov}(Z, Y) - \operatorname{cov}(Z, X)\beta\big).$$

Since $-\delta \leq \gamma \leq \delta$ (componentwise) by A4′, we have $\beta \in \mathcal{B}(\delta)$.

Next we show that $\mathcal{B}(\delta)$ is sharp. Let $b \in \mathcal{B}(\delta)$. Define

$$\gamma = \operatorname{var}(Z)^{-1}\big(\operatorname{cov}(Z, Y) - \operatorname{cov}(Z, X)b\big).$$

Then γ satisfies A4′ by definition of $\mathcal{B}(\delta)$. Next, define $\widetilde{U} \equiv Y - X'b - Z'\gamma$. Then

$$\begin{aligned}
\operatorname{cov}(Z, \widetilde{U}) &= \operatorname{cov}(Z, Y) - \operatorname{cov}(Z, X)b - \operatorname{var}(Z)\gamma \\
&= \operatorname{cov}(Z, Y) - \operatorname{cov}(Z, X)b - \operatorname{var}(Z)\operatorname{var}(Z)^{-1}\big(\operatorname{cov}(Z, Y) - \operatorname{cov}(Z, X)b\big) \\
&= 0.
\end{aligned}$$

Thus A3 holds. Hence $\mathcal{B}(\delta)$ is sharp. That the model is falsified if and only if this set is empty follows by the definition of the (sharp) identified set. *Q.E.D.*

PROOF OF COROLLARY 1: Write the identified set from Theorem 1 as

$$\begin{aligned}
\mathcal{B}(\delta) &= \{b \in \mathbb{R} : -\delta \leq \psi - b\pi \leq \delta\} \\
&= \{b \in \mathbb{R} : \psi_\ell - \delta_\ell \leq b\pi_\ell \leq \psi_\ell + \delta_\ell, \ell = 1, \dots, L\}.
\end{aligned}$$

Equation (4) follows immediately by considering the cases $\pi_\ell = 0$, $\pi_\ell < 0$, and $\pi_\ell > 0$ separately. *Q.E.D.*

PROOF OF LEMMA 1: Without loss of generality, let $\ell = 1$. The result for $\ell \neq 1$ can be obtained by permuting the components of the vector Z. Then $\widetilde{X}_1 = (X, Z_2, \dots, Z_L)$. Hence

$$\operatorname{cov}(Z, \widetilde{X}_1) = \begin{pmatrix} \operatorname{cov}(Z_1, X) & \operatorname{cov}(Z_1, Z_{-1}) \\ \operatorname{cov}(Z_{-1}, X) & \operatorname{var}(Z_{-1}) \end{pmatrix}.$$

By block matrix inversion, the first row of $\operatorname{cov}(Z, \widetilde{X}_1)^{-1}$ is

$$\begin{aligned}
&e_1' \operatorname{cov}(Z, \widetilde{X}_1)^{-1} \\
&= \begin{pmatrix} \big(\operatorname{cov}(Z_1, X) - \operatorname{cov}(Z_1, Z_{-1})\operatorname{var}(Z_{-1})^{-1}\operatorname{cov}(Z_{-1}, X)\big)^{-1} \\ -\big(\operatorname{cov}(Z_1, X) - \operatorname{cov}(Z_1, Z_{-1})\operatorname{var}(Z_{-1})^{-1}\operatorname{cov}(Z_{-1}, X)\big)^{-1}\operatorname{cov}(Z_1, Z_{-1})\operatorname{var}(Z_{-1})^{-1} \end{pmatrix}'.
\end{aligned}$$

Let $\widetilde{Z}_1 = Z_1 - \text{cov}(Z_1, Z_{-1})\,\text{var}(Z_{-1})^{-1} Z_{-1}$ be the population residual from a regression of Z_1 on Z_{-1}. Then

$$\begin{aligned}
e_1' \,\text{cov}(Z, \widetilde{X}_1)^{-1} \,\text{cov}(Z, Y) &= \frac{\text{cov}(Z_1, Y) - \text{cov}(Z_1, Z_{-1})\,\text{var}(Z_{-1})^{-1}\,\text{cov}(Z_{-1}, Y)}{\text{cov}(Z_1, X) - \text{cov}(Z_1, Z_{-1})\,\text{var}(Z_{-1})^{-1}\,\text{cov}(Z_{-1}, X)} \\
&= \frac{\text{cov}(\widetilde{Z}_1, Y)}{\text{var}(\widetilde{Z}_1)} \Big/ \frac{\text{cov}(\widetilde{Z}_1, X)}{\text{var}(\widetilde{Z}_1)} \\
&= \frac{\psi_1}{\pi_1}.
\end{aligned}$$

The last line follows by the partitioned regression formula. *Q.E.D.*

We use the following lemma in the proofs of Proposition 2 and Theorem 2. It says that the identified set for β is a singleton at any point δ in the set FF defined in equation (5).

LEMMA 2: *Suppose* A1–A3 *hold. Suppose* $K = 1$. *Let*

$$b \in \left[\min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}\right].$$

Define $\delta(b) = (|\psi_1 - b\pi_1|, \dots, |\psi_L - b\pi_L|)$. *Then* $\mathcal{B}(\delta(b)) = \{b\}$.

PROOF OF LEMMA 2: We have

$$\begin{aligned}
\mathcal{B}\big(\delta(b)\big) &= \bigcap_{\ell=1,\dots,L:\pi_\ell \neq 0} \left[\frac{\psi_\ell}{\pi_\ell} - \frac{|\psi_\ell - b\pi_\ell|}{|\pi_\ell|}, \frac{\psi_\ell}{\pi_\ell} + \frac{|\psi_\ell - b\pi_\ell|}{|\pi_\ell|}\right] \\
&= \left(\bigcap_{\ell=1,\dots,L:\psi_\ell \geq b\pi_\ell, \pi_\ell \neq 0} \left[\frac{\psi_\ell}{\pi_\ell} - \left|\frac{\psi_\ell}{\pi_\ell} - b\right|, \frac{\psi_\ell}{\pi_\ell} + \left|\frac{\psi_\ell}{\pi_\ell} - b\right|\right]\right) \\
&\quad \cap \left(\bigcap_{\ell=1,\dots,L:\psi_\ell < b\pi_\ell, \pi_\ell \neq 0} \left[\frac{\psi_\ell}{\pi_\ell} - \left|\frac{\psi_\ell}{\pi_\ell} - b\right|, \frac{\psi_\ell}{\pi_\ell} + \left|\frac{\psi_\ell}{\pi_\ell} - b\right|\right]\right) \\
&= \left(\bigcap_{\ell=1,\dots,L:\psi_\ell \geq b\pi_\ell, \pi_\ell \neq 0} \left[b, 2\frac{\psi_\ell}{\pi_\ell} - b\right]\right) \cap \left(\bigcap_{\ell=1,\dots,L:\psi_\ell < b\pi_\ell, \pi_\ell \neq 0} \left[2\frac{\psi_\ell}{\pi_\ell} - b, b\right]\right) \\
&= \{b\}.
\end{aligned}$$

The first line follows by equation (4) and the definition of $\delta(b)$. The remaining lines follow by considering two cases so that we can eliminate the absolute values. *Q.E.D.*

PROOF OF PROPOSITION 2: Let FF denote the true falsification frontier from Definition 1. Let

$$\text{FF}^{\text{guess}} = \left\{\delta \in \mathbb{R}^L_{\geq 0} : \delta_\ell = |\psi_\ell - b\pi_\ell|, \ell = 1, \dots, L, b \in \left[\min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}\right]\right\}.$$

We will show $\text{FF} = \text{FF}^{\text{guess}}$. We split the proof in three parts. The first two parts together show that $\text{FF}^{\text{guess}} \subseteq \text{FF}$. The third part shows that $\text{FF}^{\text{guess}} \supseteq \text{FF}$.

1. We first show that if $\delta \in \text{FF}^{\text{guess}}$, then the identified set $\mathcal{B}(\delta)$ is not empty. This follows immediately from Lemma 2.
2. We next show that $\delta' < \delta$ for $\delta \in \text{FF}^{\text{guess}}$ implies that $\mathcal{B}(\delta')$ is empty. So let $\delta' < \delta$ where $\delta \in \text{FF}^{\text{guess}}$ and $\delta' \geq 0$. Consider two cases:
   (a) First suppose $\delta'_\ell < \delta_\ell$ for some ℓ such that $\pi_\ell = 0$. By the definition of FF^{guess}, $\delta_\ell = |\psi_\ell|$. Note that $\delta_\ell > \delta'_\ell \geq 0$ implies $\psi_\ell \neq 0$. If $\psi_\ell > 0$ then $\psi_\ell - \delta'_\ell > 0$. So $0 \notin [\psi_\ell - \delta'_\ell, \psi_\ell + \delta'_\ell]$. Hence $B_\ell(\delta'_\ell) = \emptyset$ by equation (4). The case for $\psi_\ell < 0$ is similar. Thus in this case we must have $\mathcal{B}(\delta') = \emptyset$.
   (b) Next suppose $\delta'_\ell < \delta_\ell$ for some ℓ such that $\pi_\ell \neq 0$. $\delta' < \delta$ implies that $\mathcal{B}(\delta') \subseteq \mathcal{B}(\delta)$. By Lemma 2, $\mathcal{B}(\delta) = \{b^*\}$ for some value b^*. Thus it suffices to show that $b^* \notin \mathcal{B}(\delta')$. That will imply that $\mathcal{B}(\delta') = \emptyset$.

   To show that $b^* \notin \mathcal{B}(\delta')$ it suffices to show that $b^* \notin B_\ell(\delta')$ for some ℓ, since $\mathcal{B}(\delta')$ is the intersection of these sets over all ℓ's, by Corollary 1. From that corollary we have

$$B_\ell(\delta') = \left[\frac{\psi_\ell}{\pi_\ell} - \frac{\delta'_\ell}{|\pi_\ell|}, \frac{\psi_\ell}{\pi_\ell} + \frac{\delta'_\ell}{|\pi_\ell|}\right].$$

   If $b^* \leq \psi_\ell/\pi_\ell$, then $b^* = \psi_\ell/\pi_\ell - \delta_\ell/|\pi_\ell| < \psi_\ell/\pi_\ell - \delta'_\ell/|\pi_\ell|$. Therefore, $b^* \notin B_\ell(\delta')$. The case where $b^* > \psi_\ell/\pi_\ell$ is analogous. Hence $b^* \notin \mathcal{B}(\delta')$.

   Steps 1 and 2 together imply that $\text{FF}^{\text{guess}} \subseteq \text{FF}$.
3. Finally, we show that $\text{FF}^{\text{guess}} \supseteq \text{FF}$. We show the contrapositive: $\delta \notin \text{FF}^{\text{guess}}$ implies that $\delta \notin \text{FF}$. So let $\delta \notin \text{FF}^{\text{guess}}$. Denote

$$b_{\min} = \min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell} \quad \text{and} \quad b_{\max} = \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}.$$

   There are two cases to consider.
   (a) Suppose $\mathcal{B}(\delta) \subseteq [b_{\min}, b_{\max}]$. If $\mathcal{B}(\delta) = \emptyset$ then $\delta \notin \text{FF}$. So we can assume $\mathcal{B}(\delta) \neq \emptyset$. We will show that we can find a $\delta' < \delta$ such that $\mathcal{B}(\delta') \neq \emptyset$, and hence $\delta \notin \text{FF}$. Let $b' \in \mathcal{B}(\delta) \subseteq [b_{\min}, b_{\max}]$. First, $\delta(b') \in \text{FF}^{\text{guess}}$ and $\delta \notin \text{FF}^{\text{guess}}$ imply that $\delta \neq \delta(b')$. Since $b' \in \mathcal{B}(\delta)$, $\delta_\ell(b') = |\psi_\ell - b'\pi_\ell| \leq \delta_\ell$ for all ℓ. Thus $\delta(b') < \delta$. Next, $\mathcal{B}(\delta(b')) = \{b'\} \neq \emptyset$ by Lemma 2. So if we let $\delta' = \delta(b')$ then we have $\delta' < \delta$ and $\mathcal{B}(\delta') \neq \emptyset$. So $\delta \notin \text{FF}$ by definition of the falsification frontier.
   (b) Suppose $\mathcal{B}(\delta)$ contains an element $b \notin [b_{\min}, b_{\max}]$. Suppose $b > b_{\max}$. Let $\delta' = \delta(b_{\max}) \in \text{FF}^{\text{guess}}$. By $\delta \notin \text{FF}^{\text{guess}}$, $\delta \neq \delta'$. If ℓ is such that $\pi_\ell = 0$, then $\delta'_\ell = |\psi_\ell| = |\psi_\ell - b\pi_\ell| \leq \delta_\ell$ by definition of $b \in \mathcal{B}_\ell(\delta_\ell)$. If ℓ is such that $\pi_\ell \neq 0$, then

$$\delta'_\ell = |\psi_\ell - b_{\max}\pi_\ell| = |\pi_\ell|(b_{\max} - \psi_\ell/\pi_\ell) < |\pi_\ell|(b - \psi_\ell/\pi_\ell) \leq \delta_\ell$$

   by $b > b_{\max}$. So $\delta' \leq \delta$. Together with $\delta \neq \delta'$, we have $\delta' < \delta$. Also, $\mathcal{B}(\delta') = \{b_{\max}\} \neq \emptyset$ by Lemma 2. So $\delta \notin \text{FF}$ by definition of the falsification frontier. A similar argument applies if instead we have $b < b_{\min}$. *Q.E.D.*

PROOF OF THEOREM 2: We have

$$\begin{aligned}
\bigcup_{\delta \in \text{FF}} \mathcal{B}(\delta) &= \bigcup_{b \in \left[\min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}\right]} \mathcal{B}\big(\delta(b)\big) \\
&= \bigcup_{b \in \left[\min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}\right]} \{b\}
\end{aligned}$$

$$= \left[\min_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,\dots,L:\pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell} \right].$$

The first equality follows by the characterization of the falsification frontier in Proposition 2. The second equality follows by Lemma 2. *Q.E.D.*

PROOF OF PROPOSITION 3: It suffices to show that

$$\begin{pmatrix} \min\limits_{\ell \in \widehat{\mathcal{L}}_{\text{rel}}} \widehat{b}_\ell \\ \max\limits_{\ell \in \widehat{\mathcal{L}}_{\text{rel}}} \widehat{b}_\ell \end{pmatrix} \xrightarrow{p} \begin{pmatrix} \min\limits_{\ell \in \mathcal{L}_{\text{rel}}} \psi_\ell / \pi_\ell \\ \max\limits_{\ell \in \mathcal{L}_{\text{rel}}} \psi_\ell / \pi_\ell \end{pmatrix}.$$

We have

$$\begin{aligned} \mathbb{P}\left(\min_{\ell \in \widehat{\mathcal{L}}_{\text{rel}}} \widehat{b}_\ell = \min_{\ell \in \mathcal{L}_{\text{rel}}} \widehat{b}_\ell \right) &\geq \mathbb{P}(\widehat{\mathcal{L}}_{\text{rel}} = \mathcal{L}_{\text{rel}}) \\ &= \mathbb{P}\left(\bigcap_{\ell:\pi_\ell = 0} \{F_\ell < C_n\} \cap \bigcap_{\ell:\pi_\ell \neq 0} \{F_\ell \geq C_n\} \right) \\ &= \mathbb{P}\left(\bigcap_{\ell:\pi_\ell = 0} \{C_n^{-1} F_\ell < 1\} \cap \bigcap_{\ell:\pi_\ell \neq 0} \{n^{-1} F_\ell - n^{-1} C_n \geq 0\} \right). \end{aligned}$$

This probability converges to 1 as $n \to \infty$. To see that, use assumptions 2–4 to get $C_n^{-1} F_\ell \mathbb{1}(\pi_\ell = 0) = C_n^{-1} O_p(1) = o_p(1)$, which is strictly less than 1 with probability approaching 1. Similarly, $n^{-1} F_\ell \mathbb{1}(\pi_\ell \neq 0) - n^{-1} C_n \mathbb{1}(\pi_\ell \neq 0) = \kappa_\ell + o_p(1) - o(1)$, which is greater than or equal to 0 with probability approaching 1. Thus $\min_{\ell \in \widehat{\mathcal{L}}_{\text{rel}}} \widehat{b}_\ell = \min_{\ell \in \mathcal{L}_{\text{rel}}} \widehat{b}_\ell + o_p(1)$. By consistency of $\widehat{b}_\ell$ for $\ell \in \mathcal{L}_{\text{rel}}$ (assumption 1) and continuity of the minimum function, $\min_{\ell \in \mathcal{L}_{\text{rel}}} \widehat{b}_\ell \xrightarrow{p} \min_{\ell \in \mathcal{L}_{\text{rel}}} \psi_\ell / \pi_\ell$. The same analysis applies to the maximum. *Q.E.D.*

REFERENCES

ALTONJI, J. G., T. E. ELDER, AND C. R. TABER (2005): "An Evaluation of Instrumental Variable Strategies for Estimating the Effects of Catholic Schooling," *Journal of Human Resources*, 40, 791–821. [1450]

ANDERSON, T. W., AND H. RUBIN (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *The Annals of Mathematical Statistics*, 20, 46–63. [1453]

ANDREWS, D. W., AND S. KWON (2019): "Inference in Moment Inequality Models That Is Robust to Spurious Precision Under Model Misspecification," Cowles Foundation Discussion Paper No. 2184. [1450]

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 132, 1553–1592. [1450,1462,1463]

ANGRIST, J., AND A. B. KRUEGER (1994): "Why Do World War II Veterans Earn More Than Nonveterans?" *Journal of Labor Economics*, 12, 74–97. [1450,1453]

ARMSTRONG, T., AND M. KOLESÁR (2021): "Sensitivity Analysis Using Approximate Moment Condition Models," *Quantitative Economics*, 12, 77–108. [1450]

ASHLEY, R. (2009): "Assessing the Credibility of Instrumental Variables Inference With Imperfect Instruments via Sensitivity Analysis," *Journal of Applied Econometrics*, 24, 325–337. [1450]

ASHLEY, R. A., AND C. F. PARMETER (2015): "Sensitivity Analysis for Inference in 2SLS/GMM Estimation With Possibly Flawed Instruments," *Empirical Economics*, 49, 1153–1171. [1450]

BAUM-SNOW, N. (2007): "Did Highways Cause Suburbanization?" *The Quarterly Journal of Economics*, 122, 775–805. [1458]

BONHOMME, S., AND M. WEIDNER (2018): "Minimizing Sensitivity to Model Misspecification," arXiv preprint, arXiv:1807.02161. [1450]

BOUND, J., D. A. JAEGER, AND R. M. BAKER (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, 443–450. [1453]

CONLEY, T. G., C. B. HANSEN, AND P. E. ROSSI (2012): "Plausibly Exogenous," *The Review of Economics and Statistics*, 94, 260–272. [1450,1453,1462]

DURANTON, G., P. M. MORROW, AND M. A. TURNER (2014): "Roads and Trade: Evidence From the US," *The Review of Economic Studies*, 81, 681–724. [1457-1463]

FANG, Z., AND A. SANTOS (2019): "Inference on Directionally Differentiable Functions," *The Review of Economic Studies*, 86, 377–412. [1457]

FISHER, F. M. (1961): "On the Cost of Approximate Specification in Simultaneous Equation Estimation," *Econometrica*, 29, 139–170. [1450]

FLORES, C., AND X. CHEN (2018): *Average Treatment Effect Bounds With an Instrumental Variable: Theory and Practice*. Springer. [1450]

HANSEN, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054. [1453]

KRAAY, A. (2012): "Instrumental Variables Regressions With Uncertain Exclusion Restrictions: A Bayesian Approach," *Journal of Applied Econometrics*, 27, 108–128. [1450]

MACHADO, C., A. SHAIKH, AND E. VYTLACIL (2019): "Instrumental Variables and the Sign of the Average Treatment Effect," *Journal of Econometrics*, 212, 522–555. [1450]

MANSKI, C. F., AND J. V. PEPPER (2018): "How Do Right-to-Carry Laws Affect Crime Rates? Coping With Ambiguity Using Bounded-Variation Assumptions," *The Review of Economics and Statistics*, 100, 232–244. [1450]

MASTEN, M. A., AND A. POIRIER (2020): "Salvaging Falsified Instrumental Variable Models," arXiv preprint, arXiv:1812.11598v3. [1453]

——— (2021): "Supplement to 'Salvaging Falsified Instrumental Variable Models'," *Econometrica Supplemental Material*, 89, https://doi.org/10.3982/ECTA17969. [1455,1456]

MATZKIN, R. L. (2007): "Nonparametric Identification," *Handbook of Econometrics*, 6, 5307–5368. [1451]

RAMSAHAI, R. R. (2012): "Causal Bounds and Observable Constraints for Non-Deterministic Models," *Journal of Machine Learning Research*, 13, 829–848. [1450]

SARGAN, J. D. (1958): "The Estimation of Economic Relationships Using Instrumental Variables," *Econometrica*, 26, 393–415. [1453]

SLICHTER, D. (2014): "Testing Instrument Validity and Identification With Invalid Instruments," Working paper. [1454]

SMALL, D. S. (2007): "Sensitivity Analysis for Instrumental Variables Regression With Overidentifying Restrictions," *Journal of the American Statistical Association*, 102, 1049–1058. [1450,1453]

SWANSON, S. A., M. A. HERNÁN, M. MILLER, J. M. ROBINS, AND T. RICHARDSON (2018): "Partial Identification of the Average Treatment Effect Using Instrumental Variables: Review of Methods for Binary Instruments, Treatments, and Outcomes," *Journal of the American Statistical Association*, 113, 933–947. [1450]

VAN KIPPERSLUIS, H., AND C. A. RIETVELD (2017): "Pleiotropy-Robust Mendelian Randomization," *International Journal of Epidemiology*, 47, 1279–1288. [1450]

——— (2018): "Beyond Plausibly Exogenous," *The Econometrics Journal*, 21, 316–331. [1450]

Co-editor Guido Imbens handled this manuscript.

Manuscript received 27 December, 2019; final version accepted 24 November, 2020; available online 22 January, 2021.