

Salvaging Falsified Instrumental Variable Models*

Matthew A. Masten[†] Alexandre Poirier[‡]

January 7, 2020

Abstract

What should researchers do when their baseline model is refuted? We provide four constructive answers. First, researchers can measure the extent of falsification. To do this, we consider continuous relaxations of the baseline assumptions of concern. We then define the falsification frontier: The smallest relaxations of the baseline model which are not refuted. This frontier provides a quantitative measure of the extent of falsification. Second, researchers can present the identified set for the parameter of interest under the assumption that the true model lies somewhere on this frontier. We call this the *falsification adaptive set*. 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. Third, researchers can present the identified set for a specific point on this frontier. Finally, as a sensitivity analysis, researchers can present identified sets for points beyond the frontier. To illustrate these four ways of salvaging falsified models, we study overidentifying restrictions in two instrumental variable models: a homogeneous effects linear model, and heterogeneous effect models with either binary or continuous outcomes. In the linear model, we consider the classical overidentifying restrictions implied when multiple instruments are observed. We generalize these conditions by considering continuous relaxations of the exclusion restrictions. By sufficiently weakening the assumptions, a falsified baseline model becomes non-falsified. This lets us derive the falsification adaptive set, which has a simple closed form expression that depends only on a few 2SLS regression coefficients. We obtain similar results in the heterogeneous effect models, where we derive identified sets for marginal distributions of potential outcomes, falsification frontiers, and falsification adaptive sets under continuous relaxations of the instrument exogeneity assumptions. We apply our results to four different empirical studies: Duranton, Morrow, and Turner (2014), Alesina, Giuliano, and Nunn (2013), Acemoglu, Johnson, and Robinson (2001), and Nevo (2001). We show how our results, especially the falsification adaptive set, can help researchers go beyond the binary pass/fail outcome of a specification test and instead summarize the variation in estimates obtained from alternative non-falsified models.

JEL classification: C14; C18; C21; C26; C51

Keywords: Instrumental Variables, Nonparametric Identification, Partial Identification, Sensitivity Analysis

*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, 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, and the 2019 Greater New York Area Econometrics Colloquium. We thank audiences at those seminars and conferences 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.

Data Acknowledgements: Researchers own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the authors and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

[†]Department of Economics, Duke University, matt.masten@duke.edu

[‡]Department of Economics, Georgetown University, alexandre.poirier@georgetown.edu

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. We equivalently call these overidentified or refutable models.¹ With finite samples, researchers often use specification tests to check whether their baseline model is refuted. A well known example is the overidentifying restrictions test in linear instrumental variable models (Anderson and Rubin 1949, Sargan 1958, Hansen 1982). Abstracting from sampling uncertainty², the population versions of such specification tests have a persistent problem: What should researchers do when their baseline model is refuted?

One option is to ignore the finding and report some point estimand, like the 2SLS estimand in an overidentified linear instrumental variable model. There are two problems with this approach. First, it is often justified by appealing to large sample sizes. For example, Nevo (2001) refers to this justification on page 325: “It is well known that with a large enough sample a chi-squared test will reject essentially any model.” If we are willing to relax assumptions, however, this common justification is not true: There always exist assumptions which are sufficiently weak that they are not refuted. Second, such point estimands can be highly misleading. For example, Heckman and Hotz (1989) compare estimates from observational data with those from experimental data. In their data, they show that the observational estimates from models which fail specification tests are much farther from the experimental estimates than observational estimates from models which pass specification tests. Essentially, parameters like the 2SLS estimand are motivated by arguments that they deliver a causal effect of interest under a set of baseline assumptions. When those assumptions are known to be false, there is no reason for the corresponding estimand to be close to the causal effect of interest. In appendix B we discuss five other responses to baseline falsification from the literature. We review the related literature in appendix A.

Instead of ignoring findings from specification tests, we provide four constructive ways for researchers to salvage a falsified baseline model. First, researchers can measure the extent of falsification. To do this, we consider continuous relaxations of the baseline assumptions of concern. We then define the falsification frontier: The smallest relaxations of the baseline model which are not refuted. This frontier provides a quantitative measure of the extent of falsification. Second, researchers can present the identified set for the parameter of interest under the assumption that the true model lies somewhere 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 is not refuted. When the baseline model is refuted, this set expands to include all parameter values consistent with the data and a model which is relaxed just enough to make it non-refuted. This set generalizes the

¹Keep in mind that falsifiability is a property of a model, not a parameter. It is nonetheless common for researchers to discuss “overidentified parameters.” These are parameters whose identified sets are either empty (when the model is falsified) or a singleton (when the model is not falsified). Thus such parameters are point identified when the model is not falsified. To avoid confusion between properties of a parameter and properties of a model, we only use the terms falsifiable and refutable when describing models.

²We study population level identification and falsification in this paper. We briefly discuss finite sample estimation and inference in sections 2.5 and 3.4, but this is not the focus of the paper.

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*. While this set is agnostic about the direction of relaxation, our third suggestion is that researchers present the identified set for a specific minimally relaxed model. Finally, as a further sensitivity analysis, researchers can present identified sets for points beyond the falsification frontier. We discuss a complementary Bayesian approach in appendix D.

To illustrate these four constructive ways to salvage a falsified baseline model, we study a classic source of overidentification: observation of several instrumental variables. We do this in two different models. The first is the classical constant coefficients linear model with multiple instruments. This model imposes homogeneous treatment effects, but allows for continuous treatments. We generalize the usual overidentifying restrictions by allowing the instruments to have some direct effect on outcomes. We use this result to characterize the falsification frontier, which trades off bounds on the magnitudes of these direct instrument effects. We then characterize the identified set along the falsification frontier. This leads to a particularly simple closed form expression for the falsification adaptive set, depending only on the value of a handful of 2SLS regression coefficients. We also show how to use these identified sets to do sensitivity analysis for non-falsified models beyond the frontier.

It is well known that the classical overidentifying conditions may not hold when treatment effects are heterogeneous, even if all instrument exogeneity and exclusion restrictions hold. We therefore also study a second model, which allows for heterogeneous treatment effects and multiple instruments. In this model we focus on binary treatments. We consider both binary and continuous outcomes. Although this model allows for heterogeneous treatment effects, it is nonetheless falsifiable. We relax statistical independence between each instrument and potential outcomes using a latent propensity score distance from our previous work, Masten and Poirier (2018a). Under these relaxations, we derive the identified set for the marginal distributions of potential outcomes. We show how to quickly compute the falsification frontier and the falsification adaptive set using convex optimization. We then discuss how to use these identified sets to do sensitivity analysis regardless of whether the baseline model is refuted.

We show how to use our results in four previously published empirical studies: Duranton, Morrow, and Turner (2014, *The Review of Economic Studies*), Alesina, Giuliano, and Nunn (2013, *The Quarterly Journal of Economics*), Acemoglu, Johnson, and Robinson (2001, *The American Economic Review*), and Nevo (2001, *Econometrica*). Each paper reports 2SLS estimates using multiple instrumental variables and each paper discusses concerns about instrument validity. All four papers run overidentification tests, which sometimes fail. Even when they do not fail, the authors sometimes express a concern that this could simply be due to low sample size. We show that the falsification adaptive set is an informative complement to these traditional tests: Rather than focusing on null hypothesis significance testing, 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: Relying on different instruments

to different degrees yields different results. The FAS gives this range of results.

Finally, it is important to distinguish between two kinds of falsifiable models. The first kind is falsified because of an assumption made directly on observed random variables. For example, suppose we observe a scalar random variable. Consider the model which assumes that this variable is normally distributed. At the population level, we can simply check whether the observed variable is actually normally distributed. If not, the model is refuted. When the model is refuted, it can be salvaged by removing the normality assumption. A massive literature in statistics and econometrics on semi- and non-parametrics is largely concerned with salvaging refuted models by relaxing parametric assumptions on observed random variables. For example, see the survey by Spanos (2018). The second kind of model is falsified because of an assumption made on unobserved random variables, or unobserved structural parameters. Salvaging these models is delicate because, even at the population level, the data themselves do not tell us the correct alternative assumptions. We study this second kind of falsifiable model in this paper.

2 Salvaging Falsified Models

In this section we consider a general falsifiable model. We use this model to precisely define our four recommended responses to baseline falsification. In particular, we formally define the falsification frontier and falsification adaptive set. In sections 3 and 5 we illustrate these general concepts in two specific instrumental variable models. We discuss some interpretation issues in section 2.2. In section 2.4 we discuss the relationship between the falsification frontier and the breakdown frontier we studied in our previous work (Masten and Poirier 2019). We discuss estimation and inference in section 2.5.

2.1 Measuring the Extent of Falsification

Let W be a vector of observed random variables. Let \mathcal{F} denote the set of all cdfs on the support of W , $\text{supp}(W)$. A *model* is a set of underlying structural parameters which generate the observed distribution F_W and restrictions on those structural parameters. 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 \mathcal{M} is *falsifiable* if there are some observed distributions F_W which could not have been generated by the model. If such a cdf is observed, we say the model is falsified (equivalently, refuted). For a given model, let \mathcal{F}_f denote the set of cdfs F_W which falsify the model. Let \mathcal{F}_{nf} denote the set of cdfs F_W which do not falsify the model. Falsifiable models are often said to have *testable implications*. We prefer the terms ‘falsifiable’ or ‘refutable’, to distinguish this population level feature from issues arising in statistical hypothesis testing.³

³Informally, it is possible that for every non-falsified distribution of observables there exists a falsified distribution of observables which is arbitrarily close. Thus hypothesis tests in finite samples have a difficult time distinguishing the two cases, even though we can distinguish them at the population level. In this case we say the model is falsifiable but not testable. Canay, Santos, and Shaikh (2013) study an example of this; also see Freyberger (2017).

Suppose we begin with a falsifiable baseline model $\mathcal{M}(0_L)$. Let $\mathcal{F}_{\text{nf}}(0_L)$ denote the set of joint distributions of the data which are not falsified by this model. Hence $\mathcal{F}_{\text{nf}}(0_L)$ is a strict subset of \mathcal{F} . L denotes the number of assumptions which we believe may be the reason the model is falsifiable. For each assumption $\ell \in \{1, \dots, L\}$, 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 $\delta_\ell = 1$, and the assumption is partially imposed for $\delta_\ell \in (0, 1)$.⁴ These assumptions must be nested in the sense that for $\delta'_\ell \geq \delta_\ell$, assumption δ'_ℓ is weaker than assumption δ_ℓ . Let $\mathcal{M}(\delta)$ denote the model which imposes assumptions $\delta = (\delta_1, \dots, \delta_L)$. Let $\mathcal{F}_{\text{nf}}(\delta)$ denote the set of joint distributions of the data which are not falsified by this model. Suppose further that the model $\mathcal{M}(1_L)$ which does not impose any of the L assumptions is not falsifiable.

Recall that F_W denotes the observed distribution of the data. Suppose $F_W \notin \mathcal{F}_{\text{nf}}(0_L)$, so that the baseline model is falsified. We first consider the case where we only relax a single assumption.

Definition 1. Suppose $L = 1$. The *falsification point* is

$$\delta^* = \inf\{\delta \in [0, 1] : F_W \in \mathcal{F}_{\text{nf}}(\delta)\}.$$

That is, the falsification point is the smallest relaxation of the baseline assumption such that the model is not falsified by the observed data F_W . For any $\delta < \delta^*$, the model is falsified. For any $\delta > \delta^*$, the model is not falsified. Under mild conditions, we can show that the set $\{\delta \in [0, 1] : F_W \in \mathcal{F}_{\text{nf}}(\delta)\}$ is closed, and hence the model is not falsified at δ^* .

The falsification frontier is the multi-dimensional version of the falsification point. To formally define it, partition $[0, 1]^L$ into the set of all assumptions which are falsified,

$$\mathcal{D}_f = \{\delta \in [0, 1]^L : F_W \notin \mathcal{F}_{\text{nf}}(\delta)\}$$

and the set of all assumptions which are not falsified,

$$\mathcal{D}_{\text{nf}} = \{\delta \in [0, 1]^L : F_W \in \mathcal{F}_{\text{nf}}(\delta)\}.$$

Definition 2. The *falsification frontier* is the set

$$\text{FF} = \{\delta \in [0, 1]^L : \delta \in \mathcal{D}_{\text{nf}} \text{ and for any other } \delta' < \delta, \text{ we have } \delta' \in \mathcal{D}_f\}$$

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

That is, the falsification frontier is the set of assumptions which are not falsified, but if strengthened in any component, leads to a refuted model. In the one dimensional case, this corresponds to the smallest relaxation δ such that the model is not refuted. For example, suppose $L = 2$. For each $\delta_1 \in [0, 1]$ define

$$\delta_2^*(\delta_1) = \inf\{\delta_2 \in [0, 1] : F_W \in \mathcal{F}_{\text{nf}}(\delta_1, \delta_2)\}.$$

⁴More generally, the range of δ could be $[0, \delta_{\text{max}}]$ for some value $\delta_{\text{max}} \geq 0$, possibly $+\infty$.

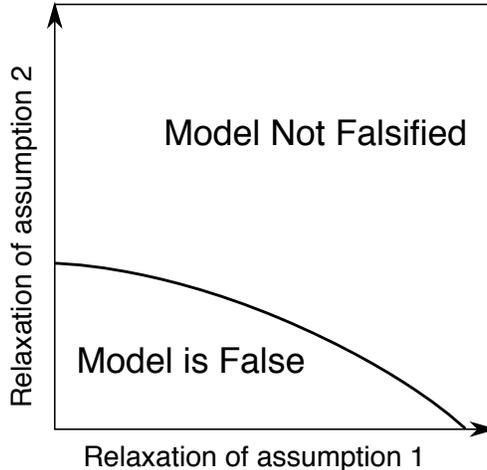


Figure 1: An example falsification frontier, partitioning the space of assumptions into the set for which the model is falsified and the set for which the model is not falsified.

This function $\delta_2^*(\cdot)$ plots the falsification frontier in the two-dimensional case. Figure 1 shows an example of such a function. Our baseline model is refuted and we have picked two different assumptions to focus on as possible explanations. The horizontal axis measures the relaxation of the first assumption. The vertical axis measures the relaxation of the second assumption. The origin at the lower left represents the baseline model. The point at the top right of the box represents the model where neither of the two baseline assumptions are imposed. First, look at the top left and lower right corners. We see that if we completely drop one assumption while maintaining the other, the model is no longer falsified. However, we can learn more by studying the points at which the model becomes non-refutable—that is, by examining the falsification frontier. This is shown as the boundary between the two regions. Consider its vertical and horizontal intercepts. Supposing the units are comparable, then the model continues to be refuted even for large relaxations of assumption 1, while maintaining assumption 2. Conversely, the model is not refuted if we allow moderate relaxations of assumption 2, while maintaining assumption 1. Looking along the frontier, strengthening assumption 2 by some amount requires relaxing assumption 1 by even more, if we want to avoid falsification. Thus the falsification frontier gives the trade-off between different assumptions’ ability to falsify the model.

2.2 Interpreting Falsified and Non-Falsified Models

In this section we discuss two well-known points that are nevertheless important to keep in mind when discussing falsification. First, although a falsified model cannot be true, a non-falsified model is not necessarily true. This asymmetry underlies Popper’s (1934) philosophy of falsification. In the classical linear instrumental variables model of section 3, this is discussed in Kadane and Anderson (1977), Newey (1985), Breusch (1986), Small (2007), Parente and Santos Silva (2012), and Guggenberger (2012). They noted that the classical overidentifying conditions are really constraints that

the instruments are consistent with each other. Thus it is possible that the instrument exogeneity assumptions fail in such a way that all instruments yield the same biased estimand for the true parameter. In this case, the baseline model is not refuted even though it is false. This is often explained by saying that the classical overidentification test is “not consistent against all alternatives.” This is why the falsification frontier is not a partition of true and false assumptions, but rather of falsified and non-falsified assumptions. We discuss this point further in appendix I.

Second, when the baseline model is falsified, it could be due to any combination of our baseline assumptions. Importantly, this is true *regardless* of the shape of the falsification frontier. The falsification frontier does *not* tell us why the baseline assumptions were refuted. For example, consider figure 1. This figure does not tell us that ‘assumption 1 is probably the reason the baseline model was refuted’. Instead, it tells us about the *robustness of falsification* to certain relaxations from the baseline assumptions. It tells us that, maintaining assumption 2, we must substantially relax assumption 1 in order to prevent falsification.

2.3 The Falsification Adaptive Set

In section 2.1, we formalized our first recommendation: Measure the extent of falsification. When the falsification frontier is far from the origin in all directions, researchers may wish to stop there. In other cases, researchers may want to proceed and present identified sets for non-falsified relaxations of the baseline model.

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 $\delta \in \mathcal{D}_f$, δ is below the falsification frontier. In this case, the identified set $\Theta_I(\delta)$ is empty. When $\delta \in \mathcal{D}_{nf}$, δ is on or above the falsification frontier. In this case, the identified set $\Theta_I(\delta)$ is nonempty.

The set of δ 's on the falsification frontier are minimally non-falsified, in the sense that they are the closest to the baseline model while still not leading to a falsified model. These δ 's, along with the δ 's beyond the frontier, are consistent with the data. As discussed in section 2.2, the data alone cannot tell us which of these models are true. It only tells us that δ 's below the falsification frontier are not true. Indeed, the same remark is true when the baseline model is not falsified. Yet in that case, it is common to present identified sets or point estimands under the baseline model. To generalize this practice to the case where the baseline model is falsified, consider the following definition.

Definition 3. Call

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

the *falsification adaptive set*.

The falsification adaptive set is the identified set for the parameter of interest under the assumption that the true model lies somewhere on the frontier. When the baseline model is not refuted, this set collapses to the $\Theta_I(0_L)$, the baseline identified set (which may be a singleton). This is what

researchers typically report when their baseline model is not refuted. When the baseline model is refuted, 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. We recommend that researchers report this set.

In some cases, researchers may have a prior belief over the relative role of the various assumptions in falsification. In this case, researchers may also want to present $\Theta_I(\delta)$ for a specific δ on the falsification frontier. For example, consider figure 1. Two natural relaxations of the baseline model are the horizontal intercept $(\delta_1^*, 0)$ and the vertical intercept $(0, \delta_2^*)$. These correspond to fully maintaining one assumption while sufficiently relaxing the other. The corresponding identified sets are $\Theta_I(\delta_1^*, 0)$ and $\Theta_I(0, \delta_2^*)$. Presenting sets like these for specific points on the frontier is our third recommendation.

Cautious Optimism

Although the points on the falsification frontier are consistent with the data, this does not mean one of them is true. For this reason, the falsification adaptive set—although it represents uncertainty due to the nature of the misspecification—is nonetheless an optimistic set. It assumes that, although the model is misspecified, it is only minimally misspecified. Since this may not be true, our fourth and final recommendation is that researchers present identified sets $\Theta_I(\delta)$ for points δ beyond the falsification frontier, as a sensitivity analysis. We discuss this recommendation in the next subsection.

The widespread practice in existing empirical work is to first present results from a usually optimistic baseline model—which typically point identifies a parameter of interest—and then to present results from various sensitivity analyses. Our recommendations directly generalize this standard practice to falsified baseline models: First present the falsification adaptive set and then present identified sets for further relaxations.

2.4 Breakdown Frontier Analysis

Our fourth recommendation is to present identified sets $\Theta_I(\delta)$ for δ 's beyond the falsification frontier. To guide this analysis, researchers can use the breakdown frontier concept which we previously studied in Masten and Poirier (2019). In this subsection we briefly summarize this concept and relate it to the falsification frontier. Consider the model $\mathcal{M}(\delta)$ from section 2.1. As above, let $\Theta_I(\delta) \subseteq \Theta$ be the identified set for some parameter of interest $\theta \in \Theta$. Let $\mathcal{C} \subseteq \Theta$. Suppose we are interested in the conclusion that $\theta \in \mathcal{C}$. For example, if our parameter is the average treatment effect, then we may be interested in the conclusion that the average treatment effect is positive, $\text{ATE} \geq 0$. Then $\mathcal{C} = [0, \infty)$. Define the *robust region*

$$\text{RR} = \{\delta \in [0, 1]^L : \Theta_I(\delta) \subseteq \mathcal{C}, \Theta_I(\delta) \neq \emptyset\}.$$

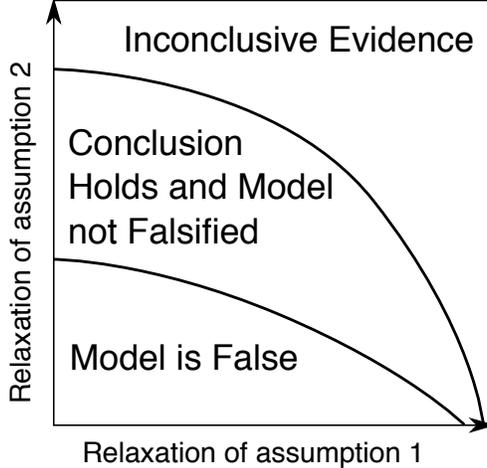


Figure 2: An example of using the falsification frontier jointly with a breakdown frontier for a specific conclusion of interest. Here there is a band of assumptions which are sufficiently weak that the model is not falsified, but are sufficiently strong that the conclusion of interest holds. This band is the robust region.

The conclusion of interest holds for all assumptions in this set. That is, for all assumptions such that all elements of the identified set are also elements of \mathcal{C} . We also restrict the set to assumptions such that the model is not refuted. The breakdown frontier is the boundary between the robust region and its complement:

$$\text{BF} = \overline{\text{RR}} \cap \overline{\text{RR}^c}.$$

In the two dimensional case, we can write the breakdown frontier as the function

$$\text{BF}(\delta_1) = \sup\{\delta_2 \in [0, 1] : \Theta_I(\delta_1, \delta_2) \subseteq \mathcal{C}, \Theta_I(\delta_1, \delta_2) \neq \emptyset\}.$$

It is possible for the robust region to be empty. This happens when the baseline model is falsified, but as soon as the model is not falsified, the identified set $\Theta_I(\delta)$ contains values of θ outside of \mathcal{C} . When the robust region is not empty, so that the breakdown frontier is not empty, the breakdown frontier will always be weakly larger than the falsification frontier. This follows since the identified set is empty for all points below the falsification frontier.

Figure 2 depicts an analysis combining both the falsification frontier and the breakdown frontier. Here the robust region is a band between the falsification frontier and the breakdown frontier. It is the region of assumptions which are weak enough that the model is not falsified, but strong enough that our conclusion of interest holds.

A key difference between the breakdown and falsification frontiers is that the falsification frontier does not depend on (a) the parameter of interest or (b) the conclusion of interest. For a given class of relaxations from the baseline assumptions, there is only one falsification frontier. In contrast, there are many different breakdown frontiers, depending on the parameter and conclusion of interest. This difference is important since there have been extensive debates in the literature regarding

what parameters researchers should study (for example, see Imbens 2010 and Heckman and Urzua 2010). Thus one’s opinions on this issue do not affect the problem of measuring the extent of falsification.

2.5 Estimation and Inference

In this paper we focus on population level analysis. In this section we briefly discuss how to implement our recommendations with finite sample data. First, to measure the extent of falsification, one can estimate and do inference on the falsification frontier. The basic idea is similar to our analysis of inference on breakdown frontiers (Masten and Poirier 2019). For the breakdown frontier, we recommended constructing lower confidence bands. For the falsification frontier, we recommend constructing *upper* confidence bands, since this provides an outer confidence set for the area under the falsification frontier—the set of models which are refuted. When doing a combined analysis, we recommend constructing an inner confidence set for the robust region; that is, for the area above the falsification frontier but below the breakdown frontier. This will require constructing two-sided confidence bands, whereas one only needs to construct one-sided confidence bands when considering each frontier separately.

When there are closed form expressions for the identified sets $\Theta_I(\delta)$ —like in our analysis below—the falsification adaptive set can be estimated by sample analog:

$$\widehat{\text{FAS}} = \bigcup_{\delta \in \widehat{\text{FF}}} \widehat{\Theta}_I(\delta).$$

In the linear instrumental variable model of section 3, we provide a particularly simple characterization of the falsification adaptive set which does not require pre-estimation of the falsification frontier. In that case, estimation and inference is straightforward. Recall that, by definition, the falsification adaptive set simplifies to the standard estimand when the baseline model is not refuted. Thus there is no need to do a specification pre-test—researchers can just always present the estimated falsification adaptive set with accompanying confidence intervals. However, note that in finite samples the estimated falsification adaptive set will generally always be an interval. This situation is similar to Haile and Tamer (2003, page 10), who note that their population bounds collapse to a single point in a special case, although in finite samples their estimator still generally gives an interval.

3 The Classical Linear Model with Multiple Instruments

In this section and section 5, we illustrate our recommendations using two variations of instrumental variable models. While many kinds of refutable assumptions have been considered in the literature, we focus on the classical case where variation from two or more instruments is used to falsify a model. In this section we consider the classical linear model with homogeneous treatment effects and continuous outcomes. In section 5 we consider a model with heterogeneous treatment effects.

3.1 Model and Identification

We thus begin with the classical constant coefficient model

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

where $Y(x, z)$ are potential outcomes defined for values $(x, z) \in \mathbb{R}^{K+L}$ and U is an unobserved random variable. β is an unknown constant K -vector. γ is an unknown constant L -vector. 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. For simplicity we have omitted any additional known exogenous covariates W in equation (1). In appendix F we show how to easily include them via partialling out.

We observe the outcome $Y = Y(X, Z)$. Thus our equation for observed outcomes is

$$Y = X'\beta + Z'\gamma + U.$$

Equation (1) imposes homogeneous treatment effects since the difference between two potential outcomes

$$Y(x_1, z) - Y(x_0, z) = (x_1 - x_0)'\beta$$

is constant across the population. We consider a model with heterogeneous treatment effects in section 5.

Remark 1. In this section we present the linear model with homogeneous treatment effects. It is well known, however, that linear model analysis can be generalized to allow for heterogeneous treatment effects under some assumptions. For example, suppose β is a random coefficient. Heckman and Vytlacil (1998) and Wooldridge (1997, 2003, 2008) showed that if the causal effect of Z on X is homogeneous, then the 2SLS estimand equals $\mathbb{E}(\beta)$. Alternatively, we can allow for heterogeneous instrument effects if we assume these are independent of β ; for example, see the discussion following equation (11) in Card (2001, page 1142). Under these kinds of assumptions, all of our results in this section extend to allow for heterogeneous treatment effects. Our analysis of heterogeneous treatment effects in section 5 applies when we are not willing to impose assumptions that restrict the nature of the heterogeneous instrument effects. Conley, Hansen, and Rossi (2012, page 261) make a similar observation.

In addition to this homogeneous treatment effect assumption, we maintain the following relevance and sufficient variation assumptions throughout this section.

Assumption A1 (Relevance). The $L \times K$ matrix $\text{cov}(Z, X)$ has rank K .

Assumption A2 (Sufficient variation). The $L \times L$ matrix $\text{var}(Z)$ is invertible.

A1 implies the order condition $L \geq K$. When there is just one endogenous variable ($K = 1$), A1 only requires $\text{cov}(X, Z_\ell) \neq 0$ 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. A2 simply requires that the instruments are not degenerate and are not affine combinations of each other.

The classical model imposes two more assumptions:

Assumption A3 (Exogeneity). $\text{cov}(Z_\ell, U) = 0$ 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. Consider model (1). Suppose A1–A4 hold. Suppose $K = 1$. Suppose the joint distribution of (Y, X, Z) is known. Then the model is not refuted if and only if

$$\text{cov}(Y, Z_m) \text{cov}(X, Z_\ell) = \text{cov}(Y, Z_\ell) \text{cov}(X, Z_m) \quad (2)$$

for all m and ℓ in $\{1, \dots, L\}$.

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

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

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 refuted. 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.

In this section we maintain the homogeneous treatment effects assumption and focus on instrument exclusion or exogeneity as an explanation for refutation. We consider heterogeneous treatment effects in section 5. We also maintain linearity of potential outcomes in X . In principle our analysis can be extended to allow for relaxations of the linearity assumption, but we leave this to future work. Note that linearity holds automatically when X is binary.

We thus focus on failure of (c) instrument exogeneity or (d) instrument exclusion as reasons for refuting the baseline model. These are two different substantive assumptions. Mathematically, however, we can use the same analysis to relax both assumptions. For simplicity, here we formally maintain the exogeneity assumption A3 and focus on failure of the exclusion assumption A4. In appendix G we explain how to use our results to relax either or both assumptions.

We relax exclusion as follows.

Assumption A4' (Partial exclusion). There are known constants $\delta_\ell \geq 0$ such that

$$|\gamma_\ell| \leq \delta_\ell$$

for all $\ell \in \{1, \dots, L\}$.

This kind of relaxation of the baseline instrumental variable assumptions was previously considered by Small (2007) and Conley et al. (2012); also see Angrist and Krueger (1994) and Bound, Jaeger, and Baker (1995). For continuous instruments one may want to standardize the instruments to have variance one, to make the magnitudes of the components in $\delta = (\delta_1, \dots, \delta_L)$ comparable. We discuss interpretation of δ_ℓ further in section 3.6 of Masten and Poirier (2018b); our empirical analysis in section 4 does not require us to select or calibrate this parameter.

Under partial exclusion, the instruments may have a direct causal effect on outcomes. For sufficiently small values of the components in δ , the model may nonetheless continue to be refuted. For sufficiently large values, however, the model will not be refuted. 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

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

is the identified set for β . Here the inequalities are component-wise. The model is refuted if and only if this set is empty.

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

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

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 $\delta = 0$, 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 $\delta = 0$ is refuted.

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

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

Here π_ℓ is the ℓ th component of Π , which is an L -vector.

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 large covariance with the outcome, so that $\psi_\ell \pm \delta_\ell$ does not contain zero, then the model is refuted. Furthermore, in this case falsification can be solely attributed to the assumption that $|\gamma_\ell| \leq \delta_\ell$. This is similar to what is sometimes called the ‘zero first-stage test’ (for example, see Slichter 2014 and the references therein). When this covariance 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 $\tilde{X}_\ell = (Z_1, \dots, Z_{\ell-1}, X, Z_{\ell+1}, \dots, Z_L)$. 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 $\text{cov}(Z, \tilde{X}_\ell)$ is invertible. Then

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

This lemma shows that ψ_ℓ/π_ℓ is the population 2SLS coefficient on X using Z_ℓ as the excluded instrument and $Z_{-\ell}$ 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 $\delta = 0$. Corollary 1 implies that $\mathcal{B}(0)$ is nonempty if and only if

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

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

$$\frac{\psi_\ell}{\pi_\ell} = \frac{\text{cov}(Y, Z_\ell)}{\text{cov}(X, Z_\ell)} \quad (5)$$

for all $\ell \in \{1, \dots, L\}$. 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 classic overidentifying conditions from proposition 1.

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 $\delta = 0$, so that the baseline model is falsified. Our next result characterizes the falsification frontier, the minimal set of δ 's which lead to a non-empty identified set. Here we focus on the single endogenous regressor case. We extend this result to multiple endogenous regressors in section 3.6.

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

$$\text{FF} = \left\{ \delta \in \mathbb{R}_{\geq 0}^L : \delta_\ell = |\psi_\ell - b\pi_\ell|, \ell = 1, \dots, L, b \in \left[\min_{\ell=1, \dots, L: \pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1, \dots, L: \pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell} \right] \right\}. \quad (6)$$

In the proof we show that this set satisfies our definition 2 of the falsification frontier. Specifically: Any $\delta \in \text{FF}$ maps to a non-empty identified set, and strengthening any of the assumptions for a given $\delta \in \text{FF}$ leads to an empty identified set.

If we knew a priori that some instruments are valid, then we could obtain the falsification frontier for relaxing the potentially invalid instrument assumptions by simply setting $\delta_\ell = 0$ for the valid instruments ℓ .

To better understand proposition 2, consider the two instrument case. In this case, the following corollary characterizes the falsification frontier.

Corollary 2. Suppose A1–A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Suppose $K = 1$ and $L = 2$. Suppose $\pi_1 \neq 0$ and $\pi_2 \neq 0$. Then the falsification frontier is the set

$$\text{FF} = \left\{ (\delta_1, \delta_2) \in \mathbb{R}_{\geq 0}^2 : \delta_2 = \left| \frac{\psi_1}{\pi_1} - \frac{\psi_2}{\pi_2} \right| |\pi_2| - \left| \frac{\pi_2}{\pi_1} \right| \delta_1 \right\}.$$

Thus, in the two instrument case, the falsification frontier is simply the line with horizontal intercept $(\delta_1^*, 0)$ and vertical intercept $(0, \delta_2^*)$, where

$$\delta_1^* = \left| \frac{\psi_1}{\pi_1} - \frac{\psi_2}{\pi_2} \right| |\pi_1| \quad \text{and} \quad \delta_2^* = \left| \frac{\psi_1}{\pi_1} - \frac{\psi_2}{\pi_2} \right| |\pi_2|.$$

The classical overidentifying restrictions (proposition 1) state that the baseline model is refuted if and only if $|\psi_1/\pi_1 - \psi_2/\pi_2|$ is nonzero. A key aspect of our analysis is that not all refuted models are equivalent. Specifically, two dgps can have the same value of this difference but have different falsification frontiers. Figure 3 illustrates this point. It shows three example falsification frontiers, corresponding to three different distributions of (Y, X, Z_1, Z_2) . All three dgps have

$$\left| \frac{\psi_1}{\pi_1} - \frac{\psi_2}{\pi_2} \right| = 4.$$

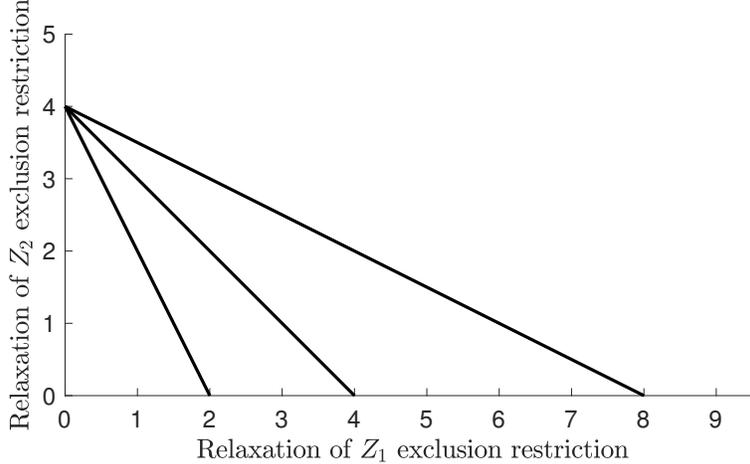


Figure 3: Three example falsification frontiers. Middle: Instruments are equally strong. Right: Z_1 is two times stronger than Z_2 . Left: Z_1 is half as strong as Z_2 .

The relative strength of the two instruments differs across the dgps, however. Specifically, we set $|\pi_1/\pi_2|$ to be either 0.5, 1, or 2. That is, the middle dgp has two instruments with exactly the same first stage coefficient. Hence there is a one-to-one trade off in the falsification frontier. If we weaken Z_1 relative to Z_2 , however, the falsification frontier rotates inwards as the horizontal intercept δ_1^* gets smaller. Conversely, if we strengthen Z_1 relative to Z_2 , the falsification frontier rotates outwards, as the horizontal intercept δ_1^* gets larger. That is, a false baseline model is harder to save by relaxing instrument exclusion for a strong instrument, relative to a weak one.

3.3 The Falsification Adaptive Set

Thus far we have characterized the identified set for β given any assumption $\delta \in \mathbb{R}_{\geq 0}^L$ as well as the falsification frontier. 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 to multiple endogenous regressors in section 3.6.

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

$$\bigcup_{\delta \in \text{FF}} \mathcal{B}(\delta) = \left[\min_{\ell=1, \dots, L: \pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1, \dots, L: \pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell} \right] \quad (7)$$

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 $\delta \in \text{FF}$ (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 (7) (follows using proposition 2). Thus we obtain the entire interval by taking the union over all of these singletons.

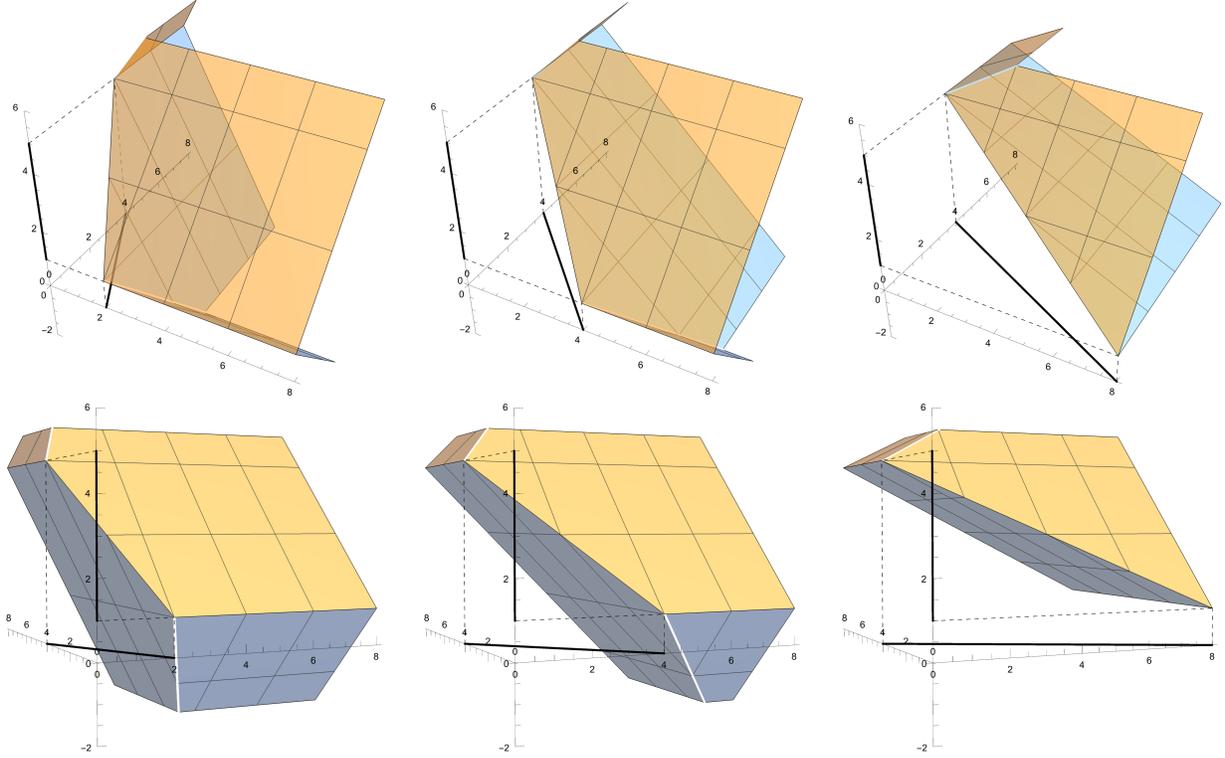


Figure 4: These plots simultaneously show the three key objects in our analysis: The identified set $\mathcal{B}(\delta)$, the falsification frontier, and the falsification adaptive set. See the text for a full description.

One of our main recommendations is that researchers report 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 (7). This implies that *we can obtain the falsification adaptive set without pre-computing 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 appendix H we show that the baseline 2SLS estimand using all instruments does not have to be inside the falsification adaptive set. In fact, the baseline 2SLS estimand can be arbitrarily far from the FAS.

In finite samples, researchers can present sample analogs or bias-corrected versions of equation (7), along with corresponding confidence sets (e.g., Chernozhukov, Lee, and Rosen 2013). We further discuss finite sample practice in section 3.4.

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 refuted. As the baseline model becomes more refuted, 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.

To show how all of the concepts we have discussed fit together, we return to the earlier two instrument numerical illustration discussed earlier. Figure 4 extends figure 3 by adding a third dimension: Values of β . Specifically, each column in figure 4 is a different dgp. Within a column, the top and bottom row are just different viewing angles. For each plot, the horizontal plane shows different values of (δ_1, δ_2) . The solid line in that plane shows the falsification frontier, as in figure 3. The vertical axis shows the identified set $\mathcal{B}(\delta)$ as a function of δ . For δ values sufficiently close to zero, this identified set is empty. But for δ 's beyond the falsification frontier, this set is nonempty. For any point δ on the falsification frontier, this set is a singleton. But the set grows as both δ_1 and δ_2 increase. Finally, the falsification adaptive set is shown as the solid line segment on the vertical axis. It is the union over over the identified sets $\mathcal{B}(\delta)$ as δ varies over the falsification frontier. Since $\mathcal{B}(\delta_1^*, 0) = \{\psi_2/\pi_2\}$, $\mathcal{B}(0, \delta_2^*) = \{\psi_1/\pi_1\}$, and $\mathcal{B}(\delta)$ varies monotonically as we traverse the falsification frontier, the falsification adaptive set is simply the interval between the points ψ_1/π_1 and ψ_2/π_2 . In this illustration, these points are 1 and 5 for all three dgps and hence the falsification adaptive set is the interval $[1, 5]$.

With more than two instruments, we cannot draw the space of δ 's and β at the same time. Nonetheless, theorem 2 shows that we can compute the falsification adaptive set by simply taking the convex hull of the L different 2SLS estimands ψ_ℓ/π_ℓ .

3.4 Estimation

In section 2.5 we discussed estimation and inference in general. Here we discuss the specific case given in theorem 2. This characterization of the falsification adaptive set requires that we first screen for weak 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

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

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

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

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 converges to zero as the sample size grows. Specifying the cutoff to shrink ensures that asymptotically we only discard instruments whose coefficients are exactly zero, so that $\widehat{\mathcal{L}}$ is consistent for \mathcal{L} .

We then estimate the falsification adaptive set by

$$\widehat{\text{FAS}} = \left[\min_{\ell \in \widehat{\mathcal{L}}} \frac{\widehat{\psi}_\ell}{\pi_\ell}, \max_{\ell \in \widehat{\mathcal{L}}} \frac{\widehat{\psi}_\ell}{\pi_\ell} \right]$$

where $\widehat{\psi}_\ell/\pi_\ell$ is the estimated 2SLS coefficient on X using Z_ℓ as the excluded instrument and $Z_{-\ell}$

as controls. We use this estimator in our empirical analysis of section 4. There we use $C_n = 10$ as our default, although we sometimes consider other cutoffs, or a sequence of cutoffs.

3.5 Directional Falsification Points

Beyond presenting the falsification frontier and the falsification adaptive set, researchers may want to also present $\mathcal{B}(\delta)$ for specific choice of model relaxation. We consider the case where the researcher specifies the *direction* of δ a priori and then chooses the smallest magnitude $\|\delta\|$ such that $\mathcal{B}(\delta) \neq \emptyset$.

Let $\delta = m \cdot d$ for $m \in \mathbb{R}$ and known $d = (d_1, \dots, d_L)' \in (0, \infty)^L$. For example, when all the instruments are continuous, one reasonable direction may be $d = (\sqrt{\text{var}(Z_1)}, \dots, \sqrt{\text{var}(Z_L)})'$. Given this fixed direction d , our assumptions are now parameterized by a single scalar, m . The following proposition shows that there is a unique falsification point m^* .

Proposition 3. Suppose A1–A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Suppose $\delta = m \cdot d$ for known $d = (d_1, \dots, d_L)' \in (0, \infty)^L$. Suppose $\pi_\ell \neq 0$ for all $\ell \in \{1, \dots, L\}$. Then

$$m^* = \max_{\ell, \ell' \in \{1, \dots, L\}} \frac{\frac{\psi_\ell}{\pi_\ell} - \frac{\psi_{\ell'}}{\pi_{\ell'}}}{\frac{d_\ell}{|\pi_\ell|} + \frac{d_{\ell'}}{|\pi_{\ell'}|}} \quad (8)$$

is the falsification point. That is, $\mathcal{B}(m^* \cdot d) \neq \emptyset$ and $\mathcal{B}(m \cdot d) = \emptyset$ for all $m < m^*$.

The following corollary characterizes the identified set at the point on the falsification frontier in the direction d .

Corollary 3. Suppose the assumptions of proposition 3 hold. Define m^* by equation (8) and let (ℓ^*, ℓ'^*) be the argmax in that equation. Let $\delta^* = m^* \cdot d$. Then $\mathcal{B}(\delta^*)$ is the singleton equal to

$$\begin{aligned} \mathcal{B}(\delta^*) &= \left\{ \frac{\psi_{\ell^*}}{\pi_{\ell^*}} - \frac{m^* \cdot d_{\ell^*}}{|\pi_{\ell^*}|} \right\} \\ &= \left\{ \frac{\psi_{\ell'^*}}{\pi_{\ell'^*}} + \frac{m^* \cdot d_{\ell'^*}}{|\pi_{\ell'^*}|} \right\}. \end{aligned}$$

3.6 The FF and FAS for K Endogenous Variables

Theorem 1 characterizes the identified set for the vector of coefficients on the endogenous variables, as a function of the exclusion restriction relaxation. Our subsequent characterizations of the falsification frontier and the falsification adaptive set, however, restricted attention to the case with just one endogenous variable—see proposition 2 and theorem 2. In this section, we extend these two results to the general case with $K \geq 1$ endogenous variables. These results can be used in at least three different cases: (1) If there are multiple ‘basic’ endogenous variables. For example, in our empirical application in section 4.4, one could allow both price and advertising spending to be endogenous. (2) If the outcome equation has interactions of one basic endogenous variable with

covariates. (3) If the outcome equation is nonlinear in one basic endogenous variable. For example, it could be a quadratic function.

Recall our notation for the reduced form and first stage regressions:

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

L denotes the number of instruments while K denotes the number of endogenous variables. Let π'_ℓ denote the ℓ th row of the matrix Π . When $K = 1$, this is a scalar and $\pi'_\ell = \pi_\ell$.

When there is a single endogenous variable, theorem 2 shows that the falsification adaptive set is the interval

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

This interval can be interpreted as follows: Pick any set of K instruments. Impose the exclusion restriction for those instruments. For the remaining instruments, completely drop the exclusion restriction. This yields a non-refutable model that is weaker than the original model which imposed exclusion for all of the instruments. But this weaker model is still strong enough that β is point identified (ignoring the possibility of irrelevant instruments, which we discuss below). Moreover, in this weaker model β equals the population 2SLS coefficient on X using the chosen K instruments as the excluded instruments and the remaining instruments as controls. Collect all of these different 2SLS estimands, as we vary which set of K instruments we pick. Then take their convex hull.

When $K = 1$, this convex hull is precisely the interval above: We pick a single relevant instrument to exclude, and include the others as controls. This gives us a single just-identified 2SLS point estimand. By cycling through which instrument we exclude, we obtain a variety of different point estimands. Since $K = 1$, these estimands are scalars. So their convex hull is simply the interval between the smallest and largest values.

When $K > 1$ and $L = K + 1$, the falsification adaptive set is precisely the convex hull of just-identified 2SLS estimands that we've just described. For $L > K + 1$, it is somewhat more complicated. We discuss these general results next.

In the $K = 1$ case we allowed for irrelevant instruments. For $K > 1$ we focus on the case where all instruments are relevant for simplicity. Specifically, we impose assumption A1' below. To state this assumption, we consider submatrices of Π . Let $\mathcal{L} \subseteq \{1, \dots, L\}$. Let $\Pi_{\mathcal{L}}$ be the $|\mathcal{L}| \times K$ submatrix of Π formed by removing all rows $\ell \notin \mathcal{L}$.

Assumption A1' (Relevance). The following hold:

1. For all $\mathcal{L} \subseteq \{1, \dots, L\}$ with $|\mathcal{L}| = K$, $\Pi_{\mathcal{L}}$ has full rank.
2. For all $\mathcal{L} \subseteq \{1, \dots, L\}$ such that $|\mathcal{L}| = K + 1$, $\{\pi_\ell : \ell \in \mathcal{L}\}$ are affinely independent. That is,

$$\begin{pmatrix} 1 & 1 & \dots & 1 \\ \pi_{\ell_1} & \pi_{\ell_2} & \dots & \pi_{\ell_{K+1}} \end{pmatrix}$$

has full rank, where $\mathcal{L} = \{\ell_1, \dots, \ell_{K+1}\}$.

A1.1' is equivalent to the following:

- Pick any K components of the instrument vector $Z = (Z_1, \dots, Z_L)$. Let $Z_{\mathcal{L}}$ denote this subvector of instruments. Let $Z_{-\mathcal{L}}$ denote the remaining instruments. Then the population 2SLS coefficient on X using $Z_{\mathcal{L}}$ as excluded instruments and $Z_{-\mathcal{L}}$ as controls is well-defined and unique. Moreover, it equals

$$\beta_{\mathcal{L}}^{2\text{SLS}} = \Pi_{\mathcal{L}}^{-1} \psi_{\mathcal{L}}$$

where $\psi_{\mathcal{L}}$ equals the subvector of ψ after removing all components $\ell \notin \mathcal{L}$.

For example, suppose $L = K + 1$. Then dropping the exclusion restriction for any single instrument returns us to a non-refutable model. Here we simply assume that β is still point identified regardless of which exclusion restriction we choose to drop. Finally, note that A1.1' implies our earlier relevance assumption A1. A1.2' means that there does not exist a hyperplane that passes through all of the π_{ℓ} vectors. It is equivalent to linear independence of $(\pi_L - \pi_1, \dots, \pi_2 - \pi_1)$.

Let

$$\text{FAS}^* = \text{conv}(\{\beta_{\mathcal{L}}^{2\text{SLS}} : \mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}| = K\}). \quad (9)$$

denote the polytope defined by the convex hull of the set of just-identified 2SLS estimands.

Proposition 4. Suppose A1', A2, and A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Suppose $L = K + 1$. Then the falsification frontier is the set

$$\text{FF} = \{\delta \in \mathbb{R}_{\geq 0}^L : \delta_{\ell} = |\psi_{\ell} - \pi'_{\ell} b|, \ell = 1, \dots, L, b \in \text{FAS}^*\}. \quad (10)$$

Theorem 3. Suppose A1', A2, and A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Suppose $L = K + 1$. Then FAS^* is the falsification adaptive set.

To illustrate theorem 3, consider the two endogenous variables ($K = 2$) and three instruments ($L = 3$) case. Since there are two endogenous variables, there are two coefficients we're interested in. Consider the left plot in figure 5. This plot shows possible values (b_1, b_2) for these two coefficients. The exclusion restriction from instrument ℓ imposes a single linear constraint $\psi_{\ell} = \pi'_{\ell} b$. These constraints are simply lines in \mathbb{R}^2 . Since there are three instruments, there are three constraints. When these three lines do not intersect at a common point, the baseline model is refuted. This case is shown in the figure. Suppose we drop the exclusion restriction for instrument 1. Then two linear constraints remain, β is point identified, and it equals the intersection point $\beta_{\{2,3\}}^{2\text{SLS}}$ in the lower right of the figure. A similar interpretation applies to the other two intersection points $\beta_{\{1,3\}}^{2\text{SLS}}$ and $\beta_{\{1,2\}}^{2\text{SLS}}$. The falsification adaptive set is then simply the convex hull of these three points in \mathbb{R}^2 , which is shown as the shaded triangular region.

Next consider the case with $K > 1$ and $L > K + 1$. We handle this case by reducing it to the case we just studied where $L = K + 1$. Let

$$\mathcal{P}_{\mathcal{L}} = \text{conv}(\{\beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} : \ell \in \mathcal{L}\}).$$

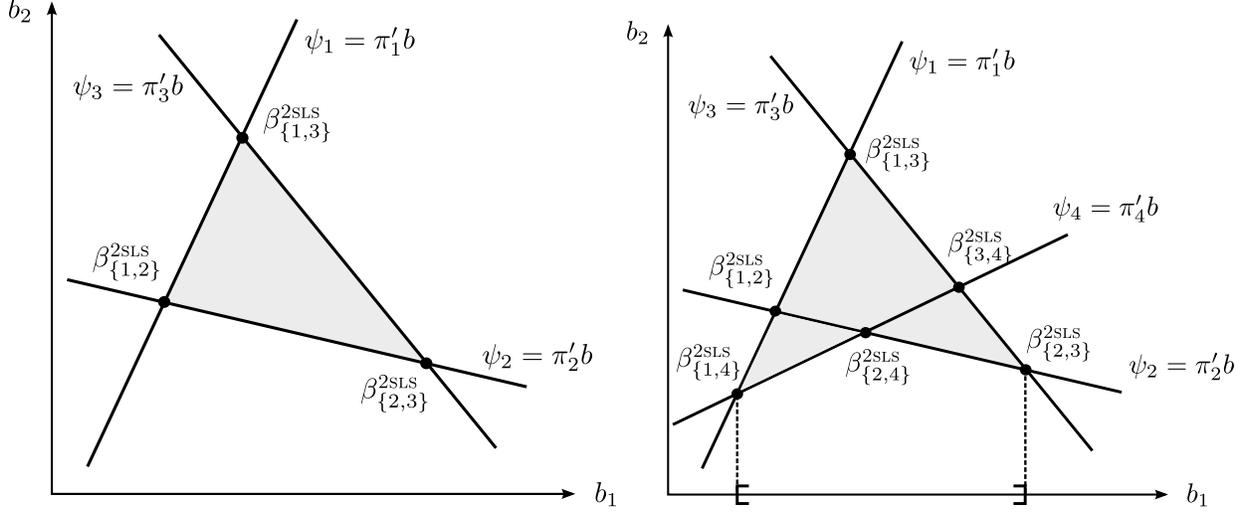


Figure 5: Example with $K = 2$ endogenous variables. Left: $L = 3$ instruments. Right: $L = 4$ instruments. In both plots, the falsification adaptive set for (β_1, β_2) is the shaded region. In the right plot, the falsification adaptive set for β_1 is shown as the projection onto the first component. See text for additional explanation.

In the following proposition, we consider subsets of indices $\mathcal{L} \subseteq \{1, \dots, L\}$ such that $|\mathcal{L}| = K + 1$. Thus $|\mathcal{L} \setminus \{\ell\}| = K$. So $\beta_{\mathcal{L} \setminus \{\ell\}}^{2SLS}$ is a just-identified 2SLS estimand. We form $K + 1$ different estimands by cycling through which exclusion restriction to drop. We then take their convex hull. This set is similar to the set defined in equation (9). The key difference is that here we only look at the estimands which use all but one index from a reference set \mathcal{L} .

Proposition 5. Suppose A1', A2, and A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Then the falsification frontier is the set

$$\text{FF} = \{\delta \in \mathbb{R}_{\geq 0}^L : \delta_\ell = |\psi_\ell - \pi'_\ell b|, b \in \mathcal{P}_\mathcal{L}, \mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}| = K + 1\}.$$

Theorem 4. Suppose A1', A2, and A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Let

$$\mathcal{P} = \bigcup_{\mathcal{L} \subseteq \{1, \dots, L\}; |\mathcal{L}| = K + 1} \mathcal{P}_\mathcal{L}.$$

Then \mathcal{P} is the falsification adaptive set.

We use Farkas' lemma and Caratheodory's theorem from convex analysis to prove these results. As in the previous cases, \mathcal{P} can be computed by just running a variety of 2SLS regressions. Next note that, although each $\mathcal{P}_\mathcal{L}$ is convex, their union generally is not. Nonetheless, we are often only interested in linear functionals of the coefficient vector β . For example, we often care about just one component of β . The following corollary shows that the falsification adaptive set for a linear functional of β again has a simple form. For this result, recall the definition of FAS* from equation (9).

Corollary 4. Suppose A1', A2, and A3 hold. Suppose the joint distribution of (Y, X, Z) is known. Then FAS* contains the falsification adaptive set for β . Moreover, for any $\alpha \in \mathbb{R}^K$ the falsification adaptive set for $\alpha'\beta$ is

$$\left[\min_{\mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}|=K} \alpha' \beta_{\mathcal{L}}^{2\text{SLS}}, \max_{\mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}|=K} \alpha' \beta_{\mathcal{L}}^{2\text{SLS}} \right].$$

This result shows that we can simply cycle through all possible just identified models, compute the corresponding 2SLS estimand, take the convex hull, and project it onto one component to get the FAS for that component.

The right plot in figure 5 illustrates the $L > K + 1$ case. Here we have $K = 2$ and $L = 4$. There are 6 different just-identified 2SLS estimands. The falsification adaptive set is no longer a convex set. Nonetheless, the projection of the convex hull of *all* just-identified 2SLS estimands onto the first component still gives the falsification adaptive set for β_1 . Moreover, this projection can be computed by simply taking the largest and smallest estimated values of β_1 , among the just-identified 2SLS estimands. In this sense, corollary 4 directly generalizes theorem 2 to the case $K > 1$.

We conclude this subsection by briefly discussing estimation and inference. Corollary 4 shows that the FAS for a linear functional of β has an intersection bounds form. As in the $K = 1$ case, researchers can present sample analog estimators or bias-corrected versions, along with corresponding confidence sets (see Chernozhukov et al. 2013). Our discussion of weak instruments in section 3.4 applies here as well. More generally, researchers may want to do estimation and inference on the FAS for the entire vector β . Theorem 4 shows that this set is defined by a finite set of unconditional linear moment inequalities. Note that here we have omitted covariates. Depending on how covariates enter the model, the FAS may continue to depend only on unconditional linear moment inequalities, or it may also depend on conditional linear moment inequalities. See appendix F for details. While there is a large literature on inference in general moment inequality models—see Canay and Shaikh (2017) and Molinari (2019) for surveys—there is now a growing literature on the linear case. This includes Hsieh, Shi, and Shum (2017), Andrews, Roth, and Pakes (2019), Cho and Russell (2019), and Gafarov (2019). We conjecture that some of these results can be applied to do inference on the FAS for β , the FAS for subvectors of β with two or more components, and the FAS for nonlinear functionals of β . We leave a full analysis for future work, however.

4 Empirical Applications

In this section we apply our results from section 3 to four different empirical studies: Duranton, Morrow, and Turner (2014, *The Review of Economic Studies*), Alesina, Giuliano, and Nunn (2013, *The Quarterly Journal of Economics*), Acemoglu, Johnson, and Robinson (2001, *The American Economic Review*), and Nevo (2001, *Econometrica*). Each paper reports 2SLS estimates using multiple instrumental variables and each paper discusses concerns about instrument validity. In

particular, all four papers run overidentification tests, which sometimes fail. Even when they do not fail, the authors sometimes express a concern that this could simply be due to low sample size. To address these concerns, we estimate falsification adaptive sets and compare them with the results reported in the original papers. We argue that the FAS is an informative complement to traditional overidentification test p -values: Rather than focusing on the null hypothesis that the instruments are consistent with each other, the FAS summarizes the range of estimates obtained from alternative models which are not refuted by the data.

4.1 Roads and Trade: Duranton, Morrow, and Turner (2014)

Duranton et al. (2014) study the relationship between roads and trade. Specifically, 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. We discuss the value measure in appendix J. 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. Specifically, they digitized five maps of exploration routes between 1528 and 1850. For each map, they counted the number of 1×1 km pixels crossed by an exploration route in the city, summed over all five maps, and took logs.
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. This map was written following the 1944 Federal Aid Highway Act. Baum-Snow (2007) had previously used this instrument, and provides a detailed history.

We first summarize their arguments for validity of these instruments. We then review their analysis and present our new results.

Arguments for Instrument Validity

Instrument relevance can be directly assessed from the first stage results in the data. We discuss this later. Still, there are a priori reasons to believe the instruments will be correlated with treatment. Building railroad tracks requires leveling ground. So does building roads. Thus old unused railroad tracks can be easily converted to highways. Exploration routes were easy to travel on foot, horseback, or wagon. Such routes are also likely to be good for cars. Finally, most of the highways on the 1947 plan were ultimately built, although additional highways not on the plan were built as well.

Next consider instrument exogeneity and exclusion. Exogeneity concerns the presence of omitted variables that affect the instrument and the outcome, or of simultaneity between the instrument and the outcome. Exclusion concerns the direct causal effect of the instrument on outcomes. We consider each instrument in turn:

1. *Railroads*. Consider exogeneity. First they argue that simultaneity is not a concern, as railroads were “constructed by private companies expecting to make a profit from railroad operations in a not too distant future,” rather than over 100 years later. Moreover, in 1898 the U.S. was primarily agricultural, and railroads were built “to transport grain, livestock and lumber as well as passengers over long distances”. This suggests that omitted variables may not be a concern, since modern manufactured goods trade in 2007 is quite different from agricultural and commodities trade in 1898. There are, however, three specific omitted variables they are worried about:

- (a) City population in 1898 affects the present of railroads in 1898. It also affects 2007 population, which affects trade in 2007. To address this, they control for population in 1920, the closest data they have to 1898.
- (b) Certain geographical features may make some cities more attractive for both railroad construction in 1898, as well as for trade in 2007. Thus they control for census region fixed effects, slope, and distance to the nearest body of water.
- (c) Time-invariant city productivity may affect 1898 railroads and trade in 2007. They argue that this can be controlled for by including population in 1950 and 2000 as controls.

Next consider exclusion. Here they are worried about two direct paths from 1898 railroads to trade in 2007:

- (a) Railroads may cause a city to specialize in manufacturing, which may then persist through the present. They address this by controlling for the log share of manufacturing employment in 1956 and in 2003. They also control for the weight of a city’s exports in 1956.
- (b) Railroads may affect the socioeconomic characteristics of cities, which then affects patterns of trade in 2007. They address this by controlling for log income per capita and log share of the population with at least a college degree.

2. *Exploration*. Consider exogeneity. First they argue that simultaneity is not a concern, as exploration routes were forged for a variety of reasons, including the search for the fountain of youth, expanding U.S. territory to the Pacific Ocean, and finding emigration routes to Oregon. Moreover, this also suggests that omitted variables may not be a concern, since the variables affecting choice of exploration route are related to these motives like the search for the fountain of youth, rather than modern manufacturing. That said, there is one set of omitted variables they are concerned about: Geographical features. These could affect both exploration routes and trade patterns in 2007. They address this concern via the same set of geographic controls mentioned above in our discussion of the railroad instrument. Finally, they do not discuss any specific concerns regarding the exclusion restriction.
3. *Plan*. Consider exogeneity. As with the other two instruments, they first argue that simultaneity is not a concern: The plan was explicitly drawn to “connect population centers [in 1947], not to anticipate future population levels and trade patterns.” However, because population is persistent, this could be an omitted variable. Hence they control for population in 1950 and 2007 as before. As with the exploration instrument, they do not discuss any specific concerns regarding the exclusion restriction.

Overall, 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.” (page 700)

With this motivation, we next present the results.

Results

First consider table 1. In this and all other tables, the non-highlighted parts reproduce results from the original paper. The highlighted parts are new computations which we have added. Panel A reproduces columns 1–4 of table 5 in Duranton et al. (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. As mentioned above, we are concerned about possible invalidity of the instruments. Thus they implement the standard test of overidentifying restrictions. At conventional sizes, it easily passes in the longest specification, fails in the second specification, and marginally passes in the first specification. Also note that these specifications do not include all of the controls discussed above; the authors include those in separate analyses, which we discuss later (our table 3).

We add the 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

Table 1: Baseline 2SLS results for Duranton et al. (2014): The effect of highways on export weight. Non-highlighted parts reproduce results from their paper. Highlighted parts are new. Panel A reproduces columns 1–4 of their table 5. It also shows the estimated falsification adaptive set. Panel B uses only two of their instruments, controlling for the other. See text for discussion.

	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	84.8
Overid. p -value	0.10	0.043	0.15	0.31
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

Notes: 66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. ***, **, *: statistically significant at 1%, 5%, 10%.

are within the FAS. We knew this was possible, given our theoretical results in appendix H (which relate the 2SLS estimand to the FAS). Second, the FAS magnitudes are all generally smaller than the 2SLS point estimates.

To better understand how we computed the FAS, and how to interpret it, next consider table 2. Columns 1–3 include the same baseline controls as column 3 in table 1 while columns 4–6 include the same baseline controls as column 4 in table 1. The only difference is that we no longer use all three variables (plan, railroad, 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 et al. (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 2 are consistent with their baseline estimates of 0.47 and 0.39 from columns 3 and 4, panel

Table 2: The effect of controlling for unused instruments. Non-highlighted parts reproduce results from their paper. Highlighted parts are new. All columns have employment, market access, and past populations as controls. Columns 4-6 also have manufacturing share of employment as controls. Compare to table 6 of Duranton et al. (2014).

	(1) Plan	(2) Railroad	(3) Exploration	(4) Plan	(5) Railroad	(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
FAS for this specification		[-0.26, 0.31]			[0.18, 0.42]	

Notes: 66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. ***, **, *: statistically significant at 1%, 5%, 10%.

A, table 1.

However, as we have discussed, 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. 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 railroad 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 railroad as an instrument, controlling for plan and highway. Unlike the uncontrolled result from panel A, railroad is a very weak instrument. Hence we ignore the result using railroad 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 alternative models which relax validity of all instruments at once, but just enough to avoid falsification. Thus the FAS reflects model uncertainty: Relying on different instruments to different degrees yields different results. This range of results is given by the FAS.

In panel B of table 2 we found that railroad 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 1. 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 are now all easily passed. (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 discussed earlier. Table 3 shows results with these controls. We begin with the full set of baseline control variables, as used in column 4 of table 1. 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 1, 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 table 6 of Duranton et al. (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.” (page 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’ve 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 railroad 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

Table 3: The effect of controlling for unused instruments, continued. Non-highlighted parts reproduce results from their paper. Highlighted parts are new. This table extends the analysis of table 2 to consider specifications with additional control variables. Columns 1–3 reproduce the results in appendix table 6 of Duranton et al. (2014). Columns 4–6 just add controls for the other instruments.

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

66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. ***, **, *: statistically significant at 1%, 5%, 10%.

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 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 railroad instrument is the most questionable, and should be used only as a control. Thus the estimates in panel B of table 1 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 et al. (2014). This is reflected in the various falsification adaptive

sets we present. In particular, the FAS for the longest specification is [0.18, 0.69]; see table 3. Accounting for sampling uncertainty only increases 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.

4.2 The Origin of Gender Roles: Alesina, Giuliano, and Nunn (2013)

Alesina et al. (2013) study the relationship between traditional agricultural practices and modern gender norms. Specifically, following Boserup (1970) they distinguish between two kinds of traditional agriculture: shifting cultivation, which “is labor intensive and uses handheld tools like the hoe and the digging stick,” and plough cultivation, which “is much more capital intensive, using the plough to prepare the soil.” They note that “unlike the hoe or digging stick, the plough requires significant upper body strength, grip strength, and bursts of power.” Hence “when plough agriculture is practiced, men have an advantage in farming relative to women.” Consequently, in societies that used plough agriculture, “men tended to work outside the home in the fields, while women specialized in activities within the home.” Boserup hypothesized that, as Alesina et al. (2013) put it, “this division of labor then generated norms about the appropriate role of women in society” which persist today.

Alesina et al. (2013) consider a variety of datasets and methods to test this hypothesis. We focus on their cross country instrumental variable analysis. Their dataset includes 160 countries. Their treatment variable is the “estimated proportion of citizens with ancestors that traditionally used the plough in pre-industrial agriculture.” They explain how they construct this variable in section 3. For this analysis they consider three different outcome variables: The female labor force participation rate in 2000, the share of a country’s firms with some female ownership (measured using data between 2005–2011), and the share of political positions held by women in 2000. Their appendix A2 gives further details on the definition of these variables. They also consider a composite outcome variable, called the average effect size (AES). This is constructed by first dividing each outcome variable by its standard deviation and then averaging the three variables.

Any correlation between treatment and outcomes is not necessarily causal for a few reasons. First, there may be reverse causality: Countries with less equal gender norms may have actively adopted innovations like the plough which conformed with these norms. Second, there may be omitted variables: Areas that were economically more developed may have been more likely to both adopt the plough and to have more equal gender norms today. This would tend to mask any causal effect of the plough on gender norms. To address these concerns, the authors use an instrumental variable approach.

The authors use properties of countries’ ancestral land geography as instruments. Specifically, Pryor (1985) classifies crops as either *plough positive* or *plough negative*. Plough positive crops benefit from the use of the plough while plough negative crops generally do not. The authors focus on cereal crops only, since they are similar except for how much they benefit from the plough.

Table 4: Country level 2SLS results from Alesina et al. (2013). Non-highlighted parts reproduce results from their paper. Highlighted parts are new. This table reproduces panels A and C of their table 8. It also shows the estimated falsification adaptive set, which here is a singleton since one of the two instruments is weak.

	Female labor force participation in 2000		Share of firms with female ownership, 2005–2011		Share of political positions held by women in 2000		Average effect size (AES)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Second stage 2SLS estimates								
Traditional plough use	-21.6*** (5.25)	-25.0*** (7.51)	-17.5*** (5.53)	-22.7*** (7.62)	-6.46*** (2.33)	-9.73*** (3.75)	-0.92*** (0.23)	-1.31*** (0.39)
Continent fixed effects		Y		Y		Y		Y
Overid. <i>p</i> -value	0.00091	0.00012	0.41	0.25	0.72	0.86	0.048	0.027
FAS	-14.3*** (5.41)	-18.0*** (6.78)	-15.3** (6.18)	-18.9** (8.11)	-7.08** (3.00)	-9.99** (4.11)	-0.73*** (0.23)	-1.08*** (0.38)
Observations	160	160	122	122	140	140	104	104
Panel B. First stage estimates. Dependent variable: Traditional plough use								
Plough-pos. environment	0.74*** (0.084)	0.63*** (0.089)	0.86*** (0.078)	0.67*** (0.10)	0.82*** (0.082)	0.68*** (0.10)	0.87*** (0.089)	0.72*** (0.12)
Plough-neg. environment	0.12 (0.12)	0.18 (0.13)	0.100 (0.17)	0.12 (0.17)	0.13 (0.13)	0.19 (0.14)	0.13 (0.18)	0.14 (0.19)
<i>F</i> -stat (both)	40.2	25.1	66.8	21.9	52	21.9	49.5	18.5
<i>F</i> -stat (pos. only)	80	57.7	123	42.8	101	43.6	95.6	36.6
<i>F</i> -stat (neg. only)	0.58	1.55	0.36	0.45	1.03	1.76	0.51	0.57

Notes: These results include a variety of historical and contemporary controls. See the table 8 notes in Alesina et al. (2013) along with their discussion in the text for details. Heteroskedasticity robust standard errors in parentheses. ***, **, *: statistically significant at 1%, 5%, 10%.

Given this classification, they construct instruments as follows: Consider a large piece of land. First measure the area of land that is suitable for growing crops in general. Pick a single plough positive cereal. Within this area of overall suitability, compute the area of land suitable for growing this specific plough positive cereal. Do this for each plough positive cereal. Note that some land is suitable for growing multiple crops, so there will be overlap in these areas. So take the average area across all plough positive cereals. Divide by the area that is suitable for growing crops overall. That is their measure of plough positive share. Next do the same for plough negative crops. Thus we have a procedure for defining suitability of a single piece of land. Using this procedure, the authors define the country level variable *plough positive environment* as the “average fraction of ancestral land that was suitable for growing [plough positive cereals] divided by the fraction that was suitable for any crops.” They use the same procedure as in their definition of the treatment variable, described in their section 3. They define the country level variable *plough negative environment* analogously.

Their main concern with validity of these instruments is failure of exogeneity due to geographical variables which affect the instruments and also modern gender norms. In their baseline results (table 8) they include a variety of geographical controls to account for this. They consider more extensive geographical controls in appendix table A14. They consider many other controls in appendix table

A15. Finally, they also use the presence of two instruments to conduct overidentification tests.

Table 4 reproduces their main instrumental variable results. As in our analysis of Duranton et al. (2014), the non-highlighted parts reproduce results from the original paper. The highlighted parts are new computations which we have added. First focus on their results. The overidentification test fails at the conventional 5% level for columns 1 and 2, as well as for the average effect size, columns 7 and 8.⁵ The authors do not comment on this finding. To provide a constructive response to this rejection, we compute the falsification adaptive set. To do this, we must first check instrument relevance. As shown in panel B, plough positive environment satisfies the relevance assumption. Plough negative environment, however, does not. The authors commented on this, since they did present the first stage coefficients on both instruments, although they did not present the separate F -test statistics as we do. Despite this relevance failure, the authors used both instruments in their main results.

In panel A, we present the estimated falsification adaptive set. Since there are only two instruments, one of which fails relevance, this set is simply a singleton. Specifically, this is the point estimate from the model using plough positive environment as an instrument, controlling for plough negative environment. Overall, the authors' qualitative findings all continue to hold using the FAS: All point estimates are still negative and statistically significant at conventional levels. Quantitatively, however, the FAS point estimates all have smaller magnitudes than the original baseline findings, except for columns 5 and 6. Excluding those columns, the point estimates are an average of 22% smaller. Finally, note that the columns where the estimates changed the least—columns 5 and 6—are those with the largest overidentification test p -values. Likewise, the columns where the estimates changed the most—columns 1 and 2—are the ones with the smallest overidentification test p -values. This is to be expected, given how the overidentification test is constructed. Our point is that the FAS is an informative complement to these p -values, since it directly summarizes the range of estimates corresponding to non-refuted alternative models. As in our analysis of Duranton et al. (2014), here it also highlighted the weakness of the plough negative instrument, suggesting that the FAS point estimates we present in panel A of table 4 are more appropriate baselines than those originally presented by the authors.

4.3 Colonial Origins of Development: Acemoglu, Johnson, and Robinson (2001)

Acemoglu et al. (2001) study the relationship between institutions and economic development. Since this paper is exceptionally well known, here we only briefly summarize it and then present our results. They consider a cross section of about 60 countries. Their treatment variable is the average protection against expropriation risk between 1985 and 1995, a measure of the strength of property rights in the country. Their outcome variable is log GDP per capita in 1995. Their primary instrument measures mortality rates faced by early settlers in the country. Their main results (table

⁵The authors report overidentification test p -values of 0.31 in column 4 and 0.06 in column 8. These differ slightly from the p -values we report in table 4. The reason is that we use heteroskedasticity robust standard errors for all calculations, including columns 4 and 8. The authors used homoskedastic standard errors for their columns 4 and 8 overidentification test (but not for the standard errors reported under their point estimates).

Table 5: Results from Acemoglu et al. (2001). Non-highlighted parts reproduce results from their paper. Highlighted parts are new. Panel B reproduces panel D of their table 8. Panel A is analogous to their panel A of table 8, except that we also use log settler mortality as an instrument. Columns 2, 3, 7, and 8 include a control for the number of years since independence. All other columns do not include controls, except as indicated in the table.

	Dependent variable: Log GDP per capita in 1995									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A. 2SLS estimates using two instruments at a time: log settler mortality paired with each of the other five instruments										
Average protection against expropriation risk, 1985–1995	0.89*** (0.13)	0.81*** (0.13)	0.80*** (0.12)	0.67*** (0.11)	0.63*** (0.11)	0.95*** (0.17)	0.83*** (0.17)	0.82*** (0.16)	0.70*** (0.14)	0.65*** (0.12)
Latitude						-0.60 (1.16)	-0.30 (1.08)	-0.21 (1.04)	-0.75 (0.88)	-0.51 (0.81)
Observations	63	60	59	60	59	63	60	59	60	59
First-stage F stat.	17.4	12.1	13.3	9.97	11.4	10.5	7.26	7.98	7.45	8.79
Overid. p -value	0.70	0.25	0.28	0.45	0.20	0.79	0.27	0.30	0.42	0.18
FAS	[0.81,0.99]	[0.45,1.03]	[0.51,1.03]	[0.48,0.77]	[0.40,0.85]	[0.88,1.02]	[0.42,1.06]	[0.48,1.04]	[0.49,0.84]	[0.41,0.93]
Panel B. 2SLS estimates using one of the other five instruments, controlling for log settler mortality										
Average protection against expropriation risk, 1985–1995	0.81*** (0.22)	0.45* (0.24)	0.51** (0.22)	0.48** (0.22)	0.40** (0.17)	0.88*** (0.28)	0.42 (0.29)	0.48* (0.27)	0.49** (0.24)	0.41** (0.18)
Log settler mortality	-0.067 (0.16)	-0.25 (0.16)	-0.21 (0.15)	-0.14 (0.15)	-0.19 (0.12)	-0.050 (0.18)	-0.26 (0.16)	-0.22 (0.16)	-0.14 (0.14)	-0.19 (0.12)
Latitude						-0.52 (1.11)	0.38 (0.86)	0.28 (0.83)	-0.38 (0.81)	-0.17 (0.70)
First-stage F stat.	9.34	3.66	4.87	3.56	5.88	6.58	2.50	3.21	3.25	5.36
Panel C. 2SLS estimates using log settler mortality, controlling for one of the other five instruments										
Average protection against expropriation risk, 1985–1995	0.99*** (0.32)	1.03*** (0.29)	1.03*** (0.31)	0.77*** (0.19)	0.85*** (0.25)	1.02*** (0.34)	1.06*** (0.34)	1.04*** (0.33)	0.84*** (0.25)	0.93*** (0.32)
European settlements in 1900	-0.0039 (0.011)					-0.0029 (0.012)				
Constraint on executive in 1900		-0.10 (0.11)					-0.099 (0.11)			
Democracy in 1900			-0.074 (0.086)					-0.071 (0.085)		
Constraint on executive in first year of independence				-0.043 (0.064)					-0.050 (0.071)	
Democracy in first year of independence					-0.058 (0.057)					-0.064 (0.064)
Latitude						-0.55 (1.25)	-0.57 (1.39)	-0.30 (1.30)	-1.07 (1.12)	-1.08 (1.22)
First-stage F stat.	6.51	7.69	6.91	8.60	6.38	6.15	5.86	5.88	6.22	4.71

Notes: Homoskedastic standard errors in parentheses. ***, **, *: statistically significant at 1%, 5%, 10%.

4) use just this one instrument. In table 8, however, they consider five more instruments. We focus on that analysis.

They use the additional instruments as follows: First, they only ever consider pairs of instruments. As they say in footnote 29, they do this to avoid the many instruments problem. We follow them and also only consider two instruments at a time. For each of the five additional instruments, they estimate the coefficient on the treatment variable using either (a) that instrument alone (their panel A) or (b) that instrument alone but also including settler mortality as a control variable (their panel D). They test overidentification by using a Hausman test to compare these two coefficient estimates (their panel C). They also present the estimated coefficient on settler mortality when it is included as a control, noting that it is not statistically indistinguishable from zero at conventional levels (their panel D).

In our table 5 we replicate and extend those results. Whereas Acemoglu et al. (2001) focused on using these additional instruments to perform overidentification tests, we instead focus on the point estimates themselves. Thus in our panel A we present a new set of baseline results, which use two instruments at a time. Acemoglu et al. (2001) did not present these results. Here we include the classical overidentification test p -value; the authors had instead presented a Hausman test result. Columns 1–5 present results with no control variables (except for 3 and 4, which include a single control: years since independence) while columns 6–10 add a control for latitude.

The overidentification tests do not reject at conventional levels in each column. Nonetheless, it is informative to ask: What estimates would we obtain from alternative non-refutable models? The last line of panel A answers this question by presenting the FAS. This set is computed from the estimated coefficients on the treatment variable in panels B and C. The authors had originally computed one of the endpoints of this set, although they focused on the coefficient on the controlled instrument. Here we compute the other endpoint of the FAS. First note that none of the instruments are particularly strong. This is a well known problem in this specific empirical application (for example, see pages 3072–3073 of Albouy 2012). Since we have not formally studied inference on the FAS, we leave the question of how to accurately measure sampling uncertainty in this application to future work. Here we will simply focus on the point estimates. For these estimates, we use a small F -statistic cut-off for inclusion in the FAS, so that the FAS is an interval for all specifications. For every specification, the estimated FAS always lies above zero. Hence the FAS supports the qualitative conclusion of this paper: Increasing property rights causes an increase in GDP. That said, for some of the specifications the length of the FAS is quite large, indicating substantial variation in the estimated magnitude of this effect.

As in our previous examples, the authors used multiple instruments to assess validity of the exogeneity and exclusion restrictions. They did this by focusing on a variety of statistical tests of overidentification. Here we illustrated the value of using the falsification adaptive set to go beyond the binary pass/fail outcome of a specification test and to instead summarize the substantive variation in estimates obtained from alternative non-refutable models.

4.4 Discrete Choice Demand Estimation: Nevo (2001)

In our final example, we study a classical application of instrumental variables: estimation of consumer demand. In a series of papers Aviv Nevo estimated the demand for ready-to-eat cereal as a first step to answering several important economic questions: How can we predict the welfare effects of a merger (Nevo 2000)? How can we test for collusion (Nevo 2001)? How can we construct price indices that account for new products and quality changes (Nevo 2003)? Answering these questions requires reliable estimates of price elasticities of demand. If one’s baseline model of consumer demand is refuted, however, it is not clear whether the estimated elasticities from this known-to-be misspecified model are relevant for these questions. For this reason, we recommend computing falsification adaptive sets. These sets show the range of estimates consistent with non-refuted models.

Since all three of the papers Nevo (2000, 2001, 2003) use essentially the same data and demand estimation, we focus on Nevo (2001). First we briefly summarize his analysis, focusing on falsification. We then discuss our results. Here we assume familiarity with standard discrete choice demand models. See Nevo (2011) for a survey.

Analysis from Nevo (2001)

Nevo begins by estimating a variety of logit demand models; see his table 5. In 9 out of the 10 specifications, the overidentification test rejects the model. Commenting on this, he says “it is unclear whether the large number of observations is the reason for the rejection or that the IV’s are not valid” (page 325). He notes that the most flexible specification, column 10, passes the test. This specification added city fixed effects, and hence he concludes that it is important to flexibly control for demographics. In our results below, however, the specification with city fixed effects is strongly rejected.

In logit demand models estimated using overidentified 2SLS, the model can be refuted for two distinct reasons: (1) The instruments are inconsistent with each other or (2) the logit functional form is incorrect. The second issue is well known and is commonly addressed by using more flexible models of demand, like the random coefficients logit model. Nevo presents estimates from this model in table 6. Using these estimates he notes that the logit model “is easily rejected” (footnote 28 on page 332).

The random coefficients logit model itself, however, is falsifiable. Nevo does not report the overidentification test result for his full model. He does report the GMM objective function value in table 6, but it is not clear if this objective function uses the optimal weighting matrix, as required to compute the J -test statistic. If it does, and assuming the statistic has not already been scaled by the sample size, then the overidentification test based on the reported value easily rejects the random coefficients logit model as well. Regardless, we also estimated the random coefficients logit model using our data (not reported here), and it is easily rejected.

As with the baseline logit model, this could be either because (1) the instruments are invalid

Table 6: Brands used for estimating demand

All Family / Basic Segment	Taste Enhanced Wholesome Segment	Simple Health Nutrition Segment	Kids Segment
K Rice Krispies K Corn Flakes GM Cheerios	P Raisin Bran P Honey Bunches of Oats K Raisin Bran Crunch K Raisin Bran K Frosted Mini-Wheats	P Grape Nuts K Special K Red Berry K Special K GM Multigrain Cheerios	P Fruity Pebbles K Frosted Flakes K Froot Loops K Corn Pops K Apple Jacks Q Life Q Cinnamon Life Q Cap'n Crunch Crunch Berries Q Cap'n Crunch GM Reese's Puffs GM Lucky Charms GM Honey Nut Cheerios GM Cinnamon Toast Crunch

Table 7: Volume market shares

Firm	2012 Q4	2013 Q4	2014 Q4	2015 Q4	2016 Q4
Kellogg	28.23	28.18	27.76	29.00	29.89
General Mills	31.03	31.03	30.48	30.74	28.92
Post	10.46	11.26	11.73	11.56	12.21
Quaker Oats	6.59	6.50	7.08	6.82	7.25
Malt-O-Meal	3.99	4.62	4.49	4.00	4.51
Kashi	2.69	2.65	2.09	2.12	1.92
C3	69.72	70.47	69.97	71.30	71.01
C6	83.00	84.24	83.64	84.24	84.70
Private Label	15.19	13.74	13.78	13.10	12.35

or (2) the random coefficients logit functional form is incorrect. One response is to use even more flexible models of demand. For recent work along these lines, see Compiani (2018) and Tebaldi, Torgovitsky, and Yang (2019). An alternative approach is to maintain the random coefficients logit functional form, but relax the exogeneity and exclusion restrictions on the instruments. Or one could relax both assumptions simultaneously. Here we focus only on the instrument assumptions. Moreover, we also focus only on the baseline logit demand model. Our characterization of the falsification adaptive set (theorem 2) can be extended to the random coefficients logit model, but since this is a substantial extension we leave it to separate work.

Finally, recall that Berry (1994, pages 252–254) shows how the nested logit model can be written as a linear instrumental variable model with multiple endogenous variables. This result combined with our multiple endogenous variable results in section 3.6 allows us to immediately study the nested logit model with endogenous price. Moreover, nested logit has been used as the baseline model in previous studies of the demand for ready-to-eat cereal—see Cotterill and Haller (1997). We omit a full empirical study of the nested logit baseline for brevity, however.

Table 8: Prices and market shares of brands in sample

Description	Mean	Median	Std	Min	Max	Variation (percentage)		
						Brand	City	Quarter
Prices (¢ per serving)	10.38	9.56	2.45	5.01	18.02	85	6.7	0.8
Advertising (M\$ per quarter)	5.92	4.98	5.15	0	23.76	53.6	–	3.5
Share within Cereal Market (%)	2.8	2.47	1.73	0.22	10.95	77.6	4.5	1.2

Table 9: Sample statistics

Description	Mean	Median	Std	Min	Max
Total Calories	136	120	35.59	100	200
Fat Calories (/100)	.11	.1	.08	0	.3
Sodium (% RDA/100)	.07	.07	.03	0	.12
Fiber (% RDA/100)	.1	.08	.09	0	.31
Sugar (g/100)	.09	.09	.05	.01	.2
Mushy (=1 if cereal gets soggy in milk)	.56	1	.51	0	1
Serving weight (g)	35.96	30	11.62	26	61
Income (\$)	12,659	9,052	14,453	0	791,777
Age (years)	34.92	34	21.96	0	85
Child (=1 if age < 16)	.25	0	.43	0	1

Data

Unlike our previous three applications, here the original data is not publicly available. We instead construct a new dataset, following the details of Nevo (2001) as closely as possible. We consider a panel dataset of 53 cities between 2012 and 2016. Nevo’s data covered 1988 to 1992. We obtained quarterly data on price and quantity of each cereal sold in these cities from Nielsen scanner data. We focus on the top 25 cereal brands, shown in table 6. Table 7 shows the volume market shares (pounds sold) for the top 6 firms, for the last quarter of each year. Comparing this to table 1 from Nevo (2001), we see that firm market shares are relatively stable over time. The industry continues to be highly concentrated, with the top 3 firms having around 70% of the market and the top 6 having around 84%. The market share of private label cereals has increased somewhat, ranging from 12 to 15% in our data, whereas it ranged from 3 to 8% between 1988 and 1992.

We obtained quarterly advertising spending from Kantar Media. Table 8 shows summary statistics for advertising, along with price and market shares. Compared to Nevo’s table 4, we see that prices are generally substantially lower, with less variation over time. Advertising spending has increased substantially.

We gathered nutrition data for each cereal by examining cereal boxes. For each city we obtained demographic information by sampling individuals from the March Current Population Survey. Table 9 provides summary statistics for nutrition and demographics.

Table 10: Logit demand parameter estimates. The specifications here are similar to those in table 5 of Nevo (2001).

	(1)	(2)	(3)	(4)
Panel A. OLS				
Price	-2.45 (0.62)	-1.49 (0.59)	-2.41 (0.53)	-14.97 (0.34)
Adjusted R^2	0.34	0.48	0.56	0.92
Panel B. 2SLS with per-quarter Hausman instruments				
Price	-6.67 (0.84)	-8.85 (0.93)	-6.89 (0.79)	-15.15 (1.12)
First-stage F stat.	730.13	391.52	378.66	65.49
Overid. test stat.	2467.31	1038.85	874.62	1085.28
Overid. p -value	0.00	0.00	0.00	0.00
Panel C. 2SLS with aggregated Hausman instruments				
Price	-7.50 (0.76)	-11.21 (0.82)	-10.60 (0.73)	-16.23 (0.58)
First-stage F stat.	6574.61	3715.17	3578.36	2134.16
Overid. test stat.	83.34	143.45	78.08	30.13
Overid. p -value	0.00	0.00	0.00	0.00
FAS ($F \geq 10$)	-21.27 (1.71)	-15.99 (1.59)	-16.93 (1.48)	-17.97 (0.66)
FAS ($F \geq 5$)	-21.27	[-332.10, -15.99]	[-182.92, -16.93]	[-58.96, -17.97]
Advertising	Y	Y	Y	Y
Time fixed effects	Y	Y	Y	Y
Product characteristics	Y			
Brand dummies		Y	Y	Y
Demographics			Y	
City fixed effects				Y

Heteroskedasticity robust standard errors in parentheses.

Results

Nevo considered two kinds of instruments:

1. Cost shifter instruments: (1) annual average wages paid in the supermarket sector in each city and (2) city density, which is viewed as a proxy for the cost of rent. In our data, these instruments are very weak. We do not use them in our main analysis.
2. Hausman instruments: “Regional quarterly average prices (excluding the city being instrumented) in all twenty quarters.” Taken together, these instruments are very strong. Our falsification adaptive set analysis, however, requires us to consider the strength of each instrument conditional on the others. In this case, most of these instruments become very weak. For this reason we also consider aggregated versions of the Hausman instruments: (1) The regional average *lagged* price, (2) the regional average *contemporaneous* price, and (3) the regional average *lead* price. For all three instruments we exclude the city being instrumented.

Table 10 shows our main results. We consider four specifications. All specifications include advertising and time fixed effects. Column 1 includes product characteristics as controls. Columns

2–4 replace these with brand dummies. Column 3 adds demographic controls: log of median income, log of median age, and median household size. Column 4 replaces this with city fixed effects. Panel A shows the OLS results. Panel B shows 2SLS results using the per-quarter Hausman instruments while panel C uses the aggregated Hausman instruments.⁶

First consider panel A. In the first three columns, the OLS point estimates are much smaller in magnitude than the 2SLS point estimates. In column 4 the OLS and 2SLS estimates are quite similar. However, notice that the overidentification test easily rejects *all* specifications (columns 1–4, panels B and C). This suggests that the instruments may not be valid. There has in fact been substantial debate over the validity of these instruments. See Bresnahan’s comment on Hausman (1996).

To address this concern, we next estimate the falsification adaptive set. As discussed above, we only do this for panel C. We consider two different cut-offs for the definition of a weak instrument: $F \geq 10$ and $F \geq 5$ (see section 3.4). With the stricter cut-off, only one of the three 2SLS estimates is not weak, and thus the estimated FAS is a singleton. All the point estimates are larger in magnitude than the corresponding baseline estimates. The coefficients are also more stable across specifications. With the weaker cut-off, the FAS is an interval for columns 2–4. The length of this interval shrinks substantially as we move from column 2 to 3 to 4. Consider the FAS in column 4: $[-58.96, -17.97]$. All elements of this interval are larger in magnitude than the estimate -16.23 from the refuted baseline model. They are also larger than the OLS estimate of -14.97 . In the logit model the price elasticities are proportional to the coefficient on price. So the lower bound on the FAS in column 4 suggests that the price elasticities could be as much as about three times larger than those suggested by the refuted baseline estimate. Thus ignoring the result of the overidentification test and using the baseline estimates anyway could lead to a large under-estimate of these elasticities. The falsification adaptive set is a constructive and informative tool researchers can use when their baseline model is rejected.

We conclude this section by briefly comparing our analysis with that of Nevo and Rosen (2012), who also study discrete choice demand models. They derive identified sets under a discrete relaxation of the classical instrument exogeneity and exclusion restrictions. They do not discuss the general problem of salvaging a falsified model. In particular, their identified sets can be empty. They do not discuss what researchers should do in this case. Furthermore, we relax instrument validity in a different and non-nested way. Hence our analysis is complementary to theirs. Using our data, we estimated the Nevo and Rosen identified set for the coefficient on price. Under their assumptions 3 and 4, and for the specification in column 4 of table 10, this set is $(-\infty, -16.727]$. This set is unbounded from below and strictly contains the estimated FAS $[-58.96, -17.97]$. In their proposition 5 they discuss an additional assumption which can sometimes be used to obtain a bounded identified set. They only consider the two instrument case, and hence we cannot apply

⁶In panel A: Column 1 corresponds to column 1 of Nevo’s table 5, Column 2 to his column 2, and column 3 to his column 3. He did not present the specification in our panel A, column 4. In panel B: Column 3 corresponds to Nevo’s column 9, and column 4 to his column 10. He did not present the specifications in our panel B, columns 1 and 2. He also did not present any of the specifications in panel C.

their result here. Moreover, this result requires choosing a sensitivity parameter γ a priori. By definition, the FAS does not require choosing any sensitivity parameters.

5 Heterogeneous Treatment Effect Models

Even when both instrument exogeneity and exclusion assumptions hold, the classical overidentifying restrictions (2) may fail when the homogeneous treatment effects assumption fails. Hence we next consider models with heterogeneous treatment effects. In section 5.1 we consider binary outcomes while in section 5.2 we consider continuous outcomes. As earlier, we omit additional covariates for simplicity.

5.1 Binary Outcomes

Unlike the baseline model in section 3 using zero-correlation conditions, the baseline model we study here using statistical independence assumptions is falsifiable even with a single instrument. Thus we begin with the case where a single binary instrument is available. This allows us to explain the main ideas and results while keeping the notation simple. After this, we consider the general case where multiple discrete instruments are available.

Let $X \in \{0, 1\}$ denote the observed treatment variable. Let $Y_1, Y_0 \in \{0, 1\}$ denote binary potential outcomes. We observe

$$Y = Y_1X + Y_0(1 - X). \tag{11}$$

Let $Z \in \{0, 1\}$ denote an observed instrument. As mentioned above, we consider multiple discretely supported instruments later. Thus we observe the joint distribution of (Y, X, Z) . Consider the following assumption.

Assumption B1 (Instrument exogeneity). $Z \perp\!\!\!\perp Y_1$ and $Z \perp\!\!\!\perp Y_0$.

In this model, Manski (1990) derived the identified set for $\mathbb{P}(Y_x = 1)$ for $x \in \{0, 1\}$ as well as the identified set for the average treatment effect

$$\text{ATE} = \mathbb{P}(Y_1 = 1) - \mathbb{P}(Y_0 = 1)$$

under B1.⁷ Balke and Pearl (1997) showed these identified sets can be empty, and hence that this model is falsifiable. Consequently, researchers whose model is falsified face the same question we discussed earlier: What should they do? We answer this question using the four recommendations provided in section 2. To do this, we must first define continuous relaxations of the instrument exogeneity assumption B1. We do this using the following concept from Masten and Poirier (2018a).

⁷Manski's (1990) analysis considered a general case which does not require outcomes, treatment, or instruments to be binary. In this general setting, he used a mean independence assumption. When outcomes are binary, mean independence of Y_x from Z is equivalent to statistical independence of Y_x and Z .

Definition 4. Let $x \in \{0, 1\}$. Let c be a scalar between 0 and 1. Say Z is c -dependent with Y_x if

$$\sup_{y_x \in \text{supp}(Y_x)} |\mathbb{P}(Z = 1 | Y_x = y_x) - \mathbb{P}(Z = 1)| \leq c. \quad (12)$$

When $c = 0$, c -dependence is equivalent to $Z \perp\!\!\!\perp Y_x$. When $c \geq \max\{\mathbb{P}(Z = 1), \mathbb{P}(Z = 0)\}$, c -dependence does not constrain the stochastic relationship between Z and Y_x . $c \in (0, 1)$ partially constrains the stochastic relationship between Z and Y_x . Specifically, while it allows the probability of $Z = 1$ given $Y_x = y_x$ to vary with y_x , it must be within a band around the unconditional probability:

$$\mathbb{P}(Z = 1 | Y_x = y_x) \in [\mathbb{P}(Z = 1) - c, \mathbb{P}(Z = 1) + c] \cap [0, 1]$$

for all $y_x \in \text{supp}(Y_x)$. Masten and Poirier (2018a) give additional discussion of how to interpret c -dependence.

Thus we relax B1 as follows.

Assumption B1' (Instrument partial exogeneity). Z is c -dependent with Y_1 and with Y_0 .

Under this relaxation, we will derive identified sets for various parameters of interest. We use these identified sets to characterize the falsification point and the falsification adaptive set. We can also use these identified sets to do sensitivity analysis beyond the falsification point.

Let $p_Z = \mathbb{P}(Z = 1)$. The following assumption rules out trivial cases.

Assumption B2. Suppose $p_Z \in (0, 1)$. Suppose $\mathbb{P}(X = 1 | Z = z) \in (0, 1)$ for $z \in \{0, 1\}$.

First we set up some notation. Then we state and discuss the main theorem. For $z \in \{0, 1\}$, define

$$k_z(c) = \frac{\mathbb{P}(Z = z) \max\{\mathbb{P}(Z = 1 - z) - c, 0\}}{\mathbb{P}(Z = 1 - z) \min\{\mathbb{P}(Z = z) + c, 1\}}. \quad (13)$$

Note that $k_z(0) = 1$, $k_z(1) = 0$, and $k_z(\cdot)$ is weakly decreasing on $c \in [0, 1]$. Using this function, define the set

$$\mathcal{D}(c) = \{(a_0, a_1) \in [0, 1]^2 : a_1 \in [a_0 k_0(c), 1 - k_0(c) + a_0 k_0(c)], a_0 \in [a_1 k_1(c), 1 - k_1(c) + a_1 k_1(c)]\}, \quad (14)$$

which depends only on c and the marginal distribution of Z . For $x \in \{0, 1\}$, define

$$\begin{aligned} \mathcal{H}_x = & [\mathbb{P}(Y = 1, X = x | Z = 0), \mathbb{P}(Y = 1, X = x | Z = 0) + \mathbb{P}(X = 1 - x | Z = 0)] \\ & \times [\mathbb{P}(Y = 1, X = x | Z = 1), \mathbb{P}(Y = 1, X = x | Z = 1) + \mathbb{P}(X = 1 - x | Z = 1)], \end{aligned} \quad (15)$$

which depends on the joint distribution of (Y, X, Z) . Let

$$\Theta_x(c) = \mathcal{D}(c) \cap \mathcal{H}_x$$

denote the intersection of these two sets.

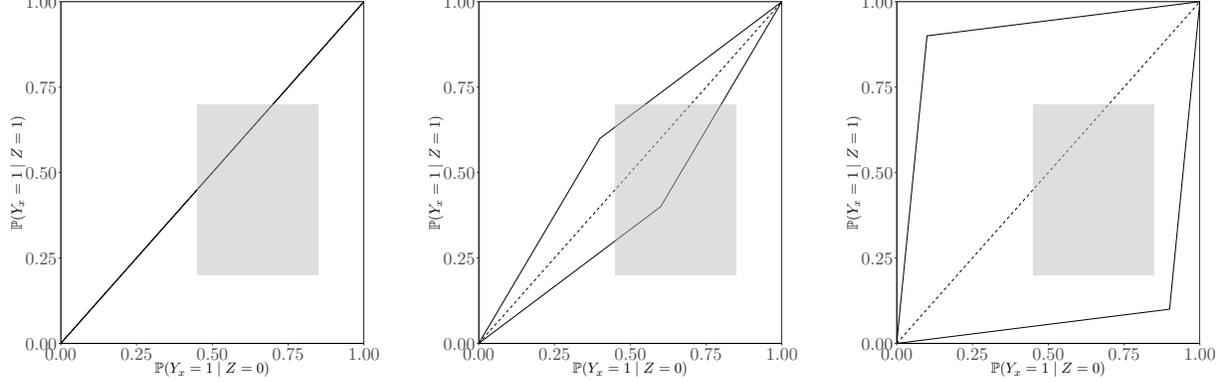


Figure 6: Example identified sets for $(\mathbb{P}(Y_x = 1 | Z = 1), \mathbb{P}(Y_x = 1 | Z = 0))$. Here $p_Z = 0.5$. Left: $c = 0$. Middle: $c = 0.1$. Right: $c = 0.4$.

Theorem 5. Consider the binary outcome, binary treatment model (11). Suppose the joint distribution of (Y, X, Z) is known, where Z is binary. Suppose B1' and B2 hold. Then the identified set for

$$(\mathbb{P}(Y_0 = 1 | Z = 0), \mathbb{P}(Y_0 = 1 | Z = 1), \mathbb{P}(Y_1 = 1 | Z = 0), \mathbb{P}(Y_1 = 1 | Z = 1))$$

is $\Theta_0(c) \times \Theta_1(c)$. Consequently, the model is falsified if and only if this set is empty.

The identified sets in theorem 5 are defined via two kinds of constraints: the set $\mathcal{D}(c)$ and the sets \mathcal{H}_x for $x \in \{0, 1\}$. The set \mathcal{H}_x is simply the identified set for

$$(\mathbb{P}(Y_x = 1 | Z = 0), \mathbb{P}(Y_x = 1 | Z = 1))$$

without imposing any assumptions on the stochastic relationship between Y_x and Z , as derived by Manski (1990). This set can be thought of as the intersection of two pairs of parallel half-spaces. The shaded boxes in figure 6 show one example of this set.

The set $\mathcal{D}(c)$ imposes the c -dependence constraint. For example, if $c = 0$, it imposes that

$$\mathbb{P}(Y_x = 1 | Z = 0) = \mathbb{P}(Y_x = 1 | Z = 1)$$

for $x \in \{0, 1\}$. This is simply the statistical independence assumption B1 studied by Manski (1990). Thus the identified sets in theorem 5 simplify to those derived by Manski when $c = 0$. As c gets larger, we no longer require the values $\mathbb{P}(Y_x = 1 | Z = z)$ for $z \in \{0, 1\}$ to exactly equal each other, but we do require them to be sufficiently close, as defined by the set $\mathcal{D}(c)$. Specifically, it constrains the pairs $(\mathbb{P}(Y_x = 1 | Z = 0), \mathbb{P}(Y_x = 1 | Z = 1))$ to lie inside a diamond-shaped parallelogram, as shown in figure 6. As c gets larger, this parallelogram gets larger, eventually approaching the entire square $[0, 1]^2$ as $c \rightarrow \max\{p_Z, 1 - p_Z\}$. Thus, for $c \geq \max\{p_Z, 1 - p_Z\}$, theorem 5 simplifies to the no assumption bounds derived by Manski (1990). Thus theorem 5 continuously connects the no assumption bounds with the bounds under the statistical independence assumption B1.

While the no assumption bounds ($c \geq \max\{p_Z, 1 - p_Z\}$) are never empty, the bounds under

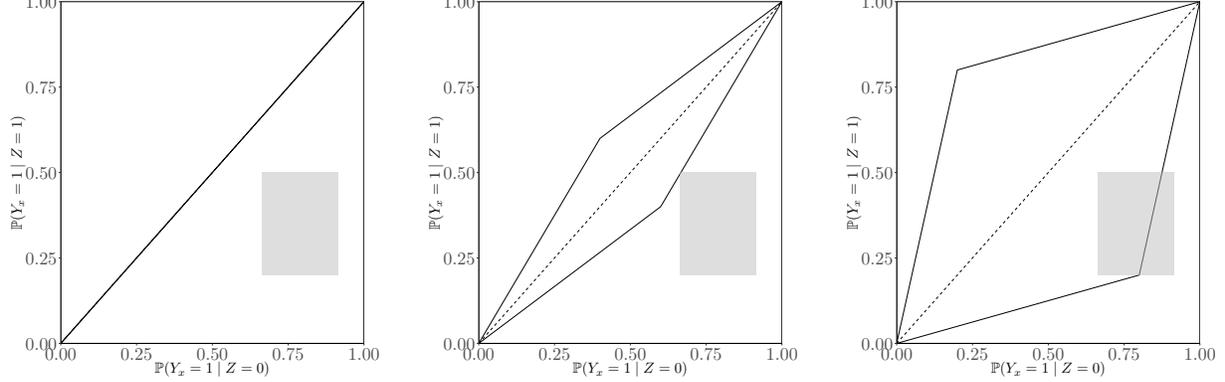


Figure 7: Example identified set for $(\mathbb{P}(Y_x = 1 | Z = 1), \mathbb{P}(Y_x = 1 | Z = 0))$. Here $p_Z = 0.5$. Left: $c = 0$. We see that the baseline model is falsified. Middle: $c = c^* = 0.1$, the falsification point. Here we have relaxed the assumption just enough that the model is no longer falsified. Right: $c = 0.3$. Relaxing independence even further increases the size of the identified set.

statistical independence ($c = 0$) can be empty, and hence the baseline statistical independence assumption can be falsified. This happens when, for some $x \in \{0, 1\}$, the no assumption bounds \mathcal{H}_x have an empty intersection with the statistical independence constraint set $\mathcal{D}(0)$. Graphically, this happens when the box defined by the no assumption bounds does not intersect the 45-degree line. This is shown in figure 7. The falsification point is simply the smallest c such that the parallelogram defined by $\mathcal{D}(c)$ has a nonempty intersection with the no assumption bounds \mathcal{H}_x for each $x \in \{0, 1\}$. This intersection is illustrated in the middle plot of figure 7.

We next state some useful properties of the identified set given in theorem 5.

Proposition 6. Suppose the assumptions of theorem 5 hold. Then:

1. There is a value $c^* \in [0, 1]$ such that $\Theta_0(c) \times \Theta_1(c)$ is empty for all $c \in [0, c^*)$ and is non-empty for all $c \in [c^*, 1]$. The value c^* is point identified.
2. For each $c \in [c^*, 1]$, the set $\Theta_0(c) \times \Theta_1(c)$ is a closed, convex polytope in $[0, 1]^4$.
3. The correspondence $\phi : [c^*, 1] \rightrightarrows [0, 1]^2$ defined by $\phi(c) = \Theta_0(c) \times \Theta_1(c)$ is continuous.

The value c^* is the falsification point. In the simple model with a single binary instrument it is possible to give a closed form expression for the falsification point, but we omit it for brevity. For part 2, recall that a polytope in $[0, 1]^4$ is a set that can be written as the intersection of finitely many half-spaces. We use continuity of the correspondence ϕ below to show that bounds on functionals like ATE are continuous in c .

We next show how to use theorem 5 to get identified sets for the counterfactual probabilities $\mathbb{P}(Y_x = 1)$ and for the average treatment effect. By the law of total probability,

$$\mathbb{P}(Y_x = 1) = \mathbb{P}(Y_x = 1 | Z = 0)\mathbb{P}(Z = 0) + \mathbb{P}(Y_x = 1 | Z = 1)\mathbb{P}(Z = 1).$$

The weights $\mathbb{P}(Z = 0)$ and $\mathbb{P}(Z = 1)$ are point identified. The joint identified set for $\mathbb{P}(Y_x = 1 \mid Z = 0)$ and $\mathbb{P}(Y_x = 1 \mid Z = 1)$ is given by $\Theta_x(c)$. Thus we can simply minimize and maximize the above convex combination over this set to obtain the identified set for $\mathbb{P}(Y_x = 1)$. Hence we define

$$\bar{P}_x(c) = \sup_{(a_0, a_1) \in \Theta_x(c)} (a_0 \mathbb{P}(Z = 0) + a_1 \mathbb{P}(Z = 1))$$

and

$$\underline{P}_x(c) = \inf_{(a_0, a_1) \in \Theta_x(c)} (a_0 \mathbb{P}(Z = 0) + a_1 \mathbb{P}(Z = 1)).$$

These are both finite dimensional linear programs and hence can be computed easily.

Corollary 5. Suppose the assumptions of theorem 5 hold. Then:

1. Let $c \in [c^*, 1]$. The identified set for $(\mathbb{P}(Y_0 = 1), \mathbb{P}(Y_1 = 1))$ is $I_0(c) \times I_1(c)$ where $I_x = [\underline{P}_x(c), \bar{P}_x(c)]$ for $x \in \{0, 1\}$.
2. For $x \in \{0, 1\}$, the functions $\underline{P}_x(c)$ and $\bar{P}_x(c)$ are continuous and monotonic over $c \in [c^*, 1]$.
3. Let $c \in [c^*, 1]$. The identified set for ATE is $[\underline{\text{ATE}}(c), \overline{\text{ATE}}(c)]$ where

$$\underline{\text{ATE}}(c) = \underline{P}_1(c) - \bar{P}_0(c) \quad \text{and} \quad \overline{\text{ATE}}(c) = \bar{P}_1(c) - \underline{P}_0(c).$$

The falsification adaptive set for $\mathbb{P}(Y_x = 1)$ is $[\underline{P}_x(c^*), \bar{P}_x(c^*)]$. This will be a singleton for one $x \in \{0, 1\}$ and will typically be an interval with a nonempty interior for the other value of x . This can be seen in figure 7. In the middle plot, $\mathbb{P}(Y_x = 1)$ is point identified since $\Theta_x(c^*)$ is a singleton. There is an analogous plot, however, for $\mathbb{P}(Y_{1-x} = 1 \mid Z = z)$. In this plot, the box \mathcal{H}_x will typically not also have a singleton intersection with the diamond $\mathcal{D}(c^*)$. Thus $\mathbb{P}(Y_{1-x} = 1)$ will typically be partially identified. This discussion implies that ATE will typically be partially identified at the falsification point. That is: The falsification adaptive set for ATE, $[\underline{\text{ATE}}(c^*), \overline{\text{ATE}}(c^*)]$, will generally be an interval with a nonempty interior.

Generalization to Multiple Discrete Instruments

We next consider the case where there are multiple discrete instruments. We relax each instrument exogeneity condition separately. We derive identified sets for a given relaxation of the instrument exogeneity conditions. As before, this allows us to derive the falsification frontier and the falsification adaptive set.

Suppose we observe $L \geq 1$ instruments, $Z = (Z_1, \dots, Z_L) \in \mathbb{R}^L$. Let $\text{supp}(Z_\ell) = \{z_\ell^1, \dots, z_\ell^{J_\ell}\}$ for some constant J_ℓ . For notational simplicity we only consider the case $J_\ell = J$ for all $\ell \in \{1, \dots, L\}$. As in theorem 5, we will derive the joint identified set for the probabilities

$$\mathbb{P}(Y_x = 1 \mid Z_\ell = z_\ell^j)$$

for $x \in \{0, 1\}$, $\ell \in \{1, \dots, L\}$, and $j \in \{1, \dots, J\}$. For each $x \in \{0, 1\}$, there are LJ conditional probabilities. As in the single binary instrument case, we make the following assumption to rule out trivial cases.

Assumption B2' For all $\ell \in \{1, \dots, L\}$ and $j \in \{1, \dots, J\}$,

1. $\mathbb{P}(Z_\ell = z_\ell^j) \in (0, 1)$.
2. $\mathbb{P}(X = 1 \mid Z_\ell = z_\ell^j) \in (0, 1)$.

Let $Z_{-\ell}$ denote the vector of instruments without Z_ℓ . Z has J^L support points while $Z_{-\ell}$ has J^{L-1} support points. Let $\text{supp}(Z_{-\ell}) = \{z_{-\ell}^1, \dots, z_{-\ell}^{J^{L-1}}\}$. By the law of total probability,

$$\begin{aligned}
& \mathbb{P}(Y_x = 1 \mid Z_\ell = z_\ell^j) \\
&= \mathbb{P}(Y = 1 \mid X = x, Z_\ell = z_\ell^j) \mathbb{P}(X = x \mid Z_\ell = z_\ell^j) \\
&+ \sum_{k=1}^{J^{L-1}} \mathbb{P}(Y_x = 1 \mid X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k) \mathbb{P}(X = 1 - x, Z_{-\ell} = z_{-\ell}^k \mid Z_\ell = z_\ell^j) \\
&= \mathbb{P}(Y = 1, X = x \mid Z_\ell = z_\ell^j) \\
&+ \sum_{k=1}^{J^{L-1}} \mathbb{P}(Y_x = 1 \mid X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k) \frac{\mathbb{P}(X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k)}{\mathbb{P}(Z_\ell = z_\ell^j)}. \tag{16}
\end{aligned}$$

The terms

$$\mathbb{P}(Y = 1, X = x \mid Z_\ell = z_\ell^j) \quad \text{and} \quad \frac{\mathbb{P}(X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k)}{\mathbb{P}(Z_\ell = z_\ell^j)}$$

are observed directly in the data. In contrast, the probabilities

$$\mathbb{P}(Y_x = 1 \mid X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k)$$

are not observed. We do know that they must lie in the interval $[0, 1]$, however. For each $x \in \{0, 1\}$, there are J^L of these unknown probabilities. Thus for each $x \in \{0, 1\}$ the LJ probabilities $\mathbb{P}(Y_x = 1 \mid Z_\ell = z_\ell^j)$ are determined by a system of LJ known linear functions of J^L unknowns which are all bounded between 0 and 1. Denote this set by

$$\mathcal{H}_x = \left\{ \mathbf{a} \in [0, 1]^{LJ} : \mathbf{a} = \mathbf{b}_x + \mathbf{A}_x \mathbf{q} \text{ for some } \mathbf{q} \in [0, 1]^{J^L} \right\} \tag{17}$$

where $\mathbf{b}_x \in \mathbb{R}^{LJ}$ is defined by

$$[\mathbf{b}_x]_{(\ell-1)J+j} = \mathbb{P}(Y = 1, X = x \mid Z_\ell = z_\ell^j)$$

for $\ell \in \{1, \dots, L\}$, $j \in \{1, \dots, J\}$ and $\mathbf{A}_x \in \mathbb{R}^{LJ \times LJ}$ is defined by

$$[\mathbf{A}_x]_{(\ell-1)J+j, m} = \frac{\mathbb{P}(X = 1 - x, Z = z^m)}{\mathbb{P}(Z_\ell = z_\ell^j)} \mathbb{1}([z^m]_\ell = z_\ell^j)$$

for $\ell \in \{1, \dots, L\}$, $j \in \{1, \dots, J\}$, $m \in \{1, \dots, J^L\}$. Here we use z^m to denote the m th element of the support of Z .

Although this notation is a bit complicated, it is simply the vector description of equation (16) above: The elements of \mathbf{b}_x are the intercepts, while the non-zero elements of the matrix \mathbf{A}_x are elements of the form

$$\mathbb{P}(X = 1 - x, Z_{-\ell} = z_{-\ell}^k \mid Z_\ell = z_\ell^j)$$

for some $\ell \in \{1, \dots, L\}$, $j \in \{1, \dots, J\}$, and $k \in \{1, \dots, J^{L-1}\}$.

As in the single instrument case, this set \mathcal{H}_x is the identified set for the joint vector of $\mathbb{P}(Y_x = 1 \mid Z_\ell = z_\ell^j)$ probabilities with no assumption on the dependence between Y_x and each instrument Z_ℓ . We next consider constraints on this dependence. As before, we relax independence using c -dependence. Specifically, we make the following assumption.

Assumption B1'' (Instrument partial exogeneity). For each $\ell \in \{1, \dots, L\}$, let c_ℓ be a scalar between 0 and 1. Then

$$|\mathbb{P}(Z_\ell = z_\ell^j \mid Y_x = y_x) - \mathbb{P}(Z_\ell = z_\ell^j)| \leq c_\ell$$

for each $j \in \{1, \dots, J\}$, $\ell \in \{1, \dots, L\}$, $y_x \in \text{supp}(Y_x)$, and $x \in \{0, 1\}$.

This assumption generalizes c -dependence to allow for multiple discrete instruments. As in the single binary instrument case, $c_\ell = 0$ implies $Y_x \perp\!\!\!\perp Z_\ell$, while $c_\ell > 0$ allows for some dependence. Importantly, we allow c_ℓ to vary across instruments. Let $c = (c_1, \dots, c_L)$.

As shown in theorem 6 below, the set

$$\mathcal{D}(c) = \prod_{\ell=1}^L \mathcal{D}_\ell(c_\ell)$$

where

$$\mathcal{D}_\ell(c_\ell) = \left\{ a \in [0, 1]^J : \text{For all } j = 1, \dots, J, \right. \quad (18)$$

$$\begin{aligned} & \mathbb{P}(Z_\ell = z_\ell^j) a_j - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k \leq 0, \\ & \mathbb{P}(Z_\ell = z_\ell^j) a_j - \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c_\ell, 0\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k \geq 0, \\ & \mathbb{P}(Z_\ell = z_\ell^j) (1 - a_j) - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) (1 - a_k) \leq 0, \\ & \left. \mathbb{P}(Z_\ell = z_\ell^j) (1 - a_j) - \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c_\ell, 0\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) (1 - a_k) \geq 0 \right\} \end{aligned}$$

characterizes the constraints that B1'' imposes on the probabilities $\mathbb{P}(Y_x = 1 \mid Z_\ell = z_\ell^j)$. This set $\mathcal{D}(c)$ generalizes the parallelogram defined by equation (14). As in the single binary instrument case, it is defined by a finite number of linear inequalities.

Let

$$\Theta_x(c) = \mathcal{D}(c) \cap \mathcal{H}_x$$

denote the intersection of these generalized diamond and box sets.

Theorem 6. Consider the binary outcome, binary treatment model (11). Suppose the joint distribution of (Y, X, Z) is known. Suppose B1'' and B2' hold. For $x \in \{0, 1\}$, let

$$p_x = \{\mathbb{P}(Y_x = 1 \mid Z_\ell = z_\ell^j) : \ell \in \{1, \dots, L\}, j \in \{j, \dots, J\}\}.$$

Then the identified set for (p_0, p_1) is $\Theta_0(c) \times \Theta_1(c)$. Consequently, the model is falsified if and only if this set is empty.

This result is a direct generalization of theorem 5. All of the discussion and interpretation there applies here as well. One new issue arises: The baseline model, $c = (0, \dots, 0)$, is falsifiable even if none of the instruments are falsifiable themselves. That is, consider the model that imposes $Y_x \perp\!\!\!\perp Z_\ell$ but no constraints on the statistical dependence between the other instruments $Z_{-\ell}$ and Y_x . We can obtain this model by setting $c = \iota - e_\ell$ where ι is a vector of 1's and e_ℓ is the unit vector. Suppose this model is not falsified, regardless of which instrument ℓ we choose to impose exogeneity on. Despite this, it is nonetheless possible that the baseline model $c = (0, \dots, 0)$ is falsified. In this case the instruments are incompatible with each other.

As in proposition 6 in the single binary instrument case, we next state some useful properties of the identified set given in theorem 6.

Proposition 7. Suppose the assumptions of theorem 6 hold. Then:

1. There is a closed set $\mathcal{C} \subseteq [0, 1]^L$ such that $\Theta_0(c) \times \Theta_1(c)$ is non-empty for all $c \in \mathcal{C}$. The set \mathcal{C} is point identified.
2. For each $c \in \mathcal{C}$, the set $\Theta_0(c) \times \Theta_1(c)$ is a closed, convex polytope.
3. The correspondence $\phi : \mathcal{C} \rightrightarrows [0, 1]^{2LJ}$ defined by $\phi(c) = \Theta_0(c) \times \Theta_1(c)$ is continuous at any $c \in \mathcal{C}$ such that $\text{int}(\mathcal{H}_x) \cap \text{int}(\mathcal{D}(c)) \neq \emptyset$ for $x \in \{0, 1\}$.

The set \mathcal{C} is the set of all instrument partial exogeneity assumptions which do *not* refute the model. The complement of this set is the set of all instrument partial exogeneity assumptions which are refuted. The falsification frontier in this case is defined as follows. Let

$$S_{<c} = \{\tilde{c} \in [0, 1]^L : \tilde{c} < c\}$$

be the set of points which are smaller than c . Then

$$\text{FF} = \{c \in [0, 1]^L : S_{<c} \subseteq ([0, 1]^L \setminus \mathcal{C})\}. \quad (19)$$

This is the set of all points c such that any points smaller than c are not in \mathcal{C} . This frontier is identified from the data since \mathcal{C} is defined implicitly by $\{\Theta_0(c) \times \Theta_1(c) : c \in [0, 1]^L\}$, which is point identified.

We next show how to use theorem 6 to get identified sets for counterfactual probabilities $\mathbb{P}(Y_x = 1)$ and for the average treatment effect. Define

$$\bar{P}_x(c) = \sup_{\mathbf{a} \in \Theta_x(c)} \sum_{j=1}^J \sum_{\ell=1}^L a_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j)$$

and

$$\underline{P}_x(c) = \inf_{\mathbf{a} \in \Theta_x(c)} \sum_{j=1}^J \sum_{\ell=1}^L a_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j).$$

As in the single binary instrument case, these are both finite dimensional linear programs and hence can be computed easily.

Corollary 6. Suppose the assumptions of theorem 6 hold. Then:

1. Let $c \in \mathcal{C}$. The identified set for $(\mathbb{P}(Y_0 = 1), \mathbb{P}(Y_1 = 1))$ is $I_0(c) \times I_1(c)$ where $I_x = [\underline{P}_x(c), \bar{P}_x(c)]$.
2. For $x \in \{0, 1\}$, the functions $\underline{P}_x(c)$ and $\bar{P}_x(c)$ are continuous at any $c \in \mathcal{C}$ such that $\text{int}(\mathcal{H}_x) \cap \text{int}(\mathcal{D}(c)) \neq \emptyset$. Moreover, they are monotonic in each component of c .
3. Let $c \in \mathcal{C}$. The identified set for ATE is $[\underline{\text{ATE}}(c), \overline{\text{ATE}}(c)]$ where

$$\underline{\text{ATE}}(c) = \underline{P}_1(c) - \bar{P}_0(c) \quad \text{and} \quad \overline{\text{ATE}}(c) = \bar{P}_1(c) - \underline{P}_0(c).$$

In the single binary instrument case, there is a simple closed form expression for the falsification point c^* . Consequently, computing the falsification adaptive set for ATE can be done easily by using linear programming to compute $\overline{P}_x(c^*)$ and $\underline{P}_x(c^*)$. In the multiple discrete instrument case, there is not a simple closed form expression for the falsification frontier. Nonetheless, it can be computed by a sequence of convex optimization problems. Specifically, the sets $\mathcal{D}(c)$ and \mathcal{H}_x are polytopes in \mathbb{R}^{LJ} . Define the ‘distance’ between these polytopes as

$$\text{dist}(\mathcal{D}(c), \mathcal{H}_x) = \inf_{(v_1, v_2) \in \mathcal{D}(c) \times \mathcal{H}_x} \|v_1 - v_2\|.$$

The model is refuted if and only if $\text{dist}(\mathcal{D}(c), \mathcal{H}_x) > 0$ for some $x \in \{0, 1\}$. For any c , the distance between these polytopes can be computed via convex optimization. Thus one can do a grid search over c and collect all values where the distance is zero for both $x \in \{0, 1\}$, and hence the model is not refuted. This is a numerical approximation to the set \mathcal{C} . We can then use this set to approximate the falsification frontier.

Given this approximate falsification frontier, we can compute the falsification adaptive set for ATE as follows. For any c , we can compute $\overline{P}_x(c)$ and $\underline{P}_x(c)$ via linear programming, since the objective function is linear and the constraint set is a polytope in \mathbb{R}^{LJ} . Thus we can simply compute these bounds for all $c \in \text{FF}$. The union of those bounds is the falsification adaptive set. Identified sets for ATE at values of c beyond the falsification frontier can be computed similarly.

5.2 Continuous Outcomes

In section 5.1 we considered binary outcomes, binary treatment, and multiple discrete instruments. In this case, the joint distribution of the data is characterized by a finite dimensional vector. It is well known that with discretely distributed data identified sets can often be computed using linear programming; see our literature review in appendix A. There are two advantages, however, to deriving analytical results like those of theorems 5 and 6. First, they can explain the role of identifying assumptions, as illustrated in figures 6 and 7. Second, they allow us to generalize the results to cases where brute force computation is costly or impossible. In this section, we show that the analytical results we derived under binary outcomes generalize to continuous outcomes. This leads us to a relatively simple and feasible approach for computing identified sets, falsification frontiers, and the falsification adaptive set under relaxations of instrument exogeneity with continuous outcomes.

We begin by assuming that outcomes are continuously distributed.

Assumption C1. For any $x, x' \in \{0, 1\}$ and $z \in \{0, 1\}^L$, the distribution of $Y_x \mid X = x', Z = z$ is continuous with respect to the Lebesgue measure.

Assumption C1 supposes that, conditional on the treatment and instruments, potential outcomes are continuously distributed. It implies that, conditional on the treatment and instruments, observed outcomes are also continuously distributed. We also suppose there are L binary instru-

ments $Z = (Z_1, \dots, Z_L)$. We can allow for discrete instruments as in section 5.1, but we only consider binary instruments to simplify the notation.

Let $\mathcal{P}(A)$ denote the set of densities (with respect to the Lebesgue measure) with support contained in A . This set can be written as

$$\mathcal{P}(A) = \left\{ p : \int_A p(y) dy = 1, p(y) \geq 0 \text{ for all } y \in \mathbb{R} \right\}.$$

We begin by considering the baseline case where the instruments are exogenous. When there is just a single instrument, this setting was studied in Kitagawa (2009). The following result is essentially his proposition 3.1.

Proposition 8. Let $L = 1$. Suppose C1 (continuous outcomes) and B1 (instrument exogeneity) hold. For each $x \in \{0, 1\}$, let \mathcal{Y}_x be a known set and suppose $\text{supp}(Y_x) \subseteq \mathcal{Y}_x$. Then the identified set for (f_{Y_1}, f_{Y_0}) is

$$\left\{ f_1 \in \mathcal{P}(\mathcal{Y}_1) : f_1(\cdot) \geq \max_{z=0,1} f_{Y,X|Z}(\cdot, 1 | z) \right\} \times \left\{ f_0 \in \mathcal{P}(\mathcal{Y}_0) : f_0(\cdot) \geq \max_{z=0,1} f_{Y,X|Z}(\cdot, 0 | z) \right\}.$$

Consequently, the model is refuted if

$$\int_{\mathcal{Y}_x} \max_{z=0,1} f_{Y,X|Z}(y, x | z) dy > 1$$

for some $x \in \{0, 1\}$.

The assumption that $\text{supp}(Y_x) \subseteq \mathcal{Y}_x$ is not restrictive since we can let $\mathcal{Y}_x = \mathbb{R}$, but this notation allows us to impose bounds on the support if they are known.

We generalize proposition 8 in two ways. First, we consider an arbitrary number L of binary instruments. Second, we relax instrument exogeneity B1 to instrument partial exogeneity B1''. As discussed above, our proof strategy generalizes the result in section 5.1. Independence of each instrument with each potential outcome requires that

$$f_{Y_x|Z_\ell}(y | 1) = f_{Y_x|Z_\ell}(y | 0)$$

for all $y \in \mathbb{R}$ and $\ell \in \{1, \dots, L\}$. This constraint is analogous to the diagonal constraint in the left plots of figures 6 and 7. The c -dependence assumption B1'' does not require the densities $f_{Y_x|Z_\ell}(\cdot | z)$ to be the same for all $z \in \{0, 1\}$, but it bounds the variation across z by a magnitude determined by the vector c of sensitivity parameters. We make this constraint precise below.

Analogously to section 5.1, we will derive the joint identified set for the conditional densities

$$\{f_{Y_x|Z_\ell}(\cdot | z) : x \in \{0, 1\}, z \in \{0, 1\}, \ell \in \{1, \dots, L\}\}$$

under instrument partial exogeneity. We define the following notation for these densities. For all

$\ell \in \{1, \dots, L\}$ and $x \in \{0, 1\}$, let

$$\mathbf{f}_{Y_x|Z_\ell} = \begin{pmatrix} f_{Y_x|Z_\ell}(\cdot | 0) \\ f_{Y_x|Z_\ell}(\cdot | 1) \end{pmatrix} \in \mathcal{P}(\text{supp}(Y_x | Z_\ell = 0)) \times \mathcal{P}(\text{supp}(Y_x | Z_\ell = 1))$$

and

$$(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell} = \begin{pmatrix} \mathbf{f}_{Y_x|Z_1} \\ \vdots \\ \mathbf{f}_{Y_x|Z_L} \end{pmatrix} \in \prod_{\ell=1}^L \left(\mathcal{P}(\text{supp}(Y_x | Z_\ell = 0)) \times \mathcal{P}(\text{supp}(Y_x | Z_\ell = 1)) \right).$$

Thus we will derive the joint identified set for the vectors $(\mathbf{f}_{Y_0|Z_\ell})_{\text{all } \ell}$ and $(\mathbf{f}_{Y_1|Z_\ell})_{\text{all } \ell}$.

This identified set has a structure analogous to that in section 5.1. Thus we first define the continuous outcome generalization of the set $\mathcal{D}(c)$ defined in equations (14) and (18). To do this, define

$$k_z^\ell(c_\ell) = \frac{\mathbb{P}(Z_\ell = z) \max\{\mathbb{P}(Z_\ell = 1 - z) - c_\ell, 0\}}{\mathbb{P}(Z_\ell = 1 - z) \min\{\mathbb{P}(Z_\ell = z) + c_\ell, 1\}}$$

for $\ell \in \{1, \dots, L\}$ and $z \in \{0, 1\}$. This function generalizes equation (13). For each $x \in \{0, 1\}$, let \mathcal{Y}_x be a known set. We assume this set contains the support of $Y_x | X, Z$ for all values of the conditioning variables. Define

$$\mathcal{D}_{x,\ell}(c_\ell) = \left\{ (f(\cdot | 0), f(\cdot | 1)) \in \mathcal{P}(\mathcal{Y}_x)^2 : f(\cdot | 0)k_0^\ell(c_\ell) \leq f(\cdot | 1) \text{ and } f(\cdot | 1)k_1^\ell(c_\ell) \leq f(\cdot | 0) \right\}$$

for $\ell \in \{1, \dots, L\}$ and $x \in \{0, 1\}$. This set generalizes equations (14) and (18). It depends on $x \in \{0, 1\}$ only via the possibility that \mathcal{Y}_x might depend on x . Finally, take the product to get

$$\mathcal{D}_x(c) = \prod_{\ell=1}^L \mathcal{D}_{x,\ell}(c_\ell)$$

We will show that $\mathcal{D}_x(c)$ is the set of densities $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell}$ consistent with c -dependence (B1''). Note that setting $c_\ell = 0$ implies that $f_{Y_x|Z_\ell} = f_{Y_x}$, as required by the baseline model.

Next we generalize the set \mathcal{H}_x defined in equations (15) and (17). With continuous outcomes this will be the set of proper densities $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell}$ consistent with the observed data. Let \mathbf{A}_x be the $2L \times 2^L$ matrix such that for $\ell \in \{1, \dots, L\}$ and $m \in \{1, \dots, 2^L\}$:

$$[\mathbf{A}_x]_{(\ell-1)2+1,m} = \frac{\mathbb{P}(X = 1 - x, Z = z^m)}{\mathbb{P}(Z_\ell = 0)} \mathbb{1}([z^m]_\ell = 0)$$

and

$$[\mathbf{A}_x]_{(\ell-1)2+2,m} = \frac{\mathbb{P}(X = 1 - x, Z = z^m)}{\mathbb{P}(Z_\ell = 1)} \mathbb{1}([z^m]_\ell = 1).$$

Here we use z^m to denote the m th element of the support of Z . Let

$$\mathbf{f}_{Y,X|Z_\ell}(\cdot, x) = \begin{pmatrix} f_{Y,X|Z_\ell}(\cdot, x | 0) \\ f_{Y,X|Z_\ell}(\cdot, x | 1) \end{pmatrix} \quad \text{and} \quad (\mathbf{f}_{Y,X|Z_\ell}(\cdot, x))_{\text{all } \ell} = \begin{pmatrix} \mathbf{f}_{Y,X|Z_1}(\cdot, x) \\ \vdots \\ \mathbf{f}_{Y,X|Z_L}(\cdot, x) \end{pmatrix}.$$

These are vectors of observed densities of outcomes and treatment conditional on different instruments. Define

$$\mathcal{H}_x = \left\{ \mathbf{f}_x \in \mathcal{P}(\mathcal{Y}_x)^{2L} : \mathbf{f}_x = (\mathbf{f}_{Y,X|Z_\ell}(\cdot, x))_{\text{all } \ell} + \mathbf{A}_x \mathbf{q} \text{ for some } \mathbf{q} \in \mathcal{P}(\mathcal{Y}_x)^{2L} \right\}.$$

This set generalizes equation (17). Finally, let

$$\Theta_x(c) = \mathcal{D}_x(c) \cap \mathcal{H}_x.$$

Theorem 7. Consider the continuous outcome, binary treatment model (11). Suppose the joint distribution of (Y, X, Z) is known. Suppose C1, B1'', and B2' hold. For each $x \in \{0, 1\}$, let \mathcal{Y}_x be a known set and suppose $\text{supp}(Y_x | X = x', Z = z^m) \subseteq \mathcal{Y}_x$ for all $x' \in \{0, 1\}$ and $m \in \{1, \dots, 2^L\}$. Then the identified set for $((\mathbf{f}_{Y_0|Z_\ell})_{\text{all } \ell}, (\mathbf{f}_{Y_1|Z_\ell})_{\text{all } \ell})$ is

$$\Theta_0(c) \times \Theta_1(c).$$

Consequently, the model is falsified if and only if this set is empty.

As in proposition 8, knowledge of the set \mathcal{Y}_x is not restrictive since we can let $\mathcal{Y}_x = \mathbb{R}$, but this notation allows us to impose bounds on its support if they are known. We could also generalize this set to allow it to vary with x' and z^m , at the cost of additional notation.

The identified set $\Theta_0(c) \times \Theta_1(c)$ is a convex set of functions defined by linear equality and inequality constraints. It weakly expands as components of c increase. It nests the identified set under the baseline independence assumption ($c = 0$) and the identified set under no assumptions on the dependence between potential outcomes and instruments ($c = \iota_L$, an L -vector of ones).

For a given c , this identified set can be empty, which implies the model at that c is falsified. Define

$$\mathcal{C} = \{c \in [0, 1]^L : \Theta_0(c) \times \Theta_1(c) \neq \emptyset\}.$$

This is the set of all partial exogeneity assumptions which are not refuted. It has the property that $c \in \mathcal{C}$ and $c' \geq c$ implies $c' \in \mathcal{C}$. Using this set, we can compute the falsification frontier as in equation (19).

As in section 5.1, we can obtain the identified set for functionals of the densities $f_{Y_x|Z_\ell}$ via

convex optimization. For example, the identified set for $(f_{Y_0}(\cdot), f_{Y_1}(\cdot))$ is $I_0(c) \times I_1(c)$ where

$$I_x(c) = \{f \in \mathcal{P}(\mathcal{Y}_x) : f(\cdot) = f_{Y_x|Z_1}(\cdot | 0)\mathbb{P}(Z_1 = 0) + f_{Y_x|Z_1}(\cdot | 1)\mathbb{P}(Z_1 = 1)$$

where $(f_{Y_x|Z_1}(\cdot | 0), f_{Y_x|Z_1}(\cdot | 1))$ is the first component of some $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell}$ in $\Theta_x(c)$ }.

Given this identified set, suppose we are interested in a linear functional

$$\Gamma(f_1, f_0) = \int_{\mathcal{Y}_1} \omega_1(y_1) f_1(y_1) dy_1 + \int_{\mathcal{Y}_0} \omega_0(y_0) f_0(y_0) dy_0$$

where ω_1 and ω_0 are known weight functions. For example, with $\omega_1(y) = y$ and $\omega_0(y) = -y$, $\Gamma(f_1, f_0)$ equals the average treatment effect. Let

$$\bar{\Gamma}(c) = \sup_{f_1 \in I_1(c)} \int_{\mathcal{Y}_1} \omega_1(y_1) f_1(y_1) dy_1 + \sup_{f_0 \in I_0(c)} \int_{\mathcal{Y}_0} \omega_0(y_0) f_0(y_0) dy_0$$

and

$$\underline{\Gamma}(c) = \inf_{f_1 \in I_1(c)} \int_{\mathcal{Y}_1} \omega_1(y_1) f_1(y_1) dy_1 + \inf_{f_0 \in I_0(c)} \int_{\mathcal{Y}_0} \omega_0(y_0) f_0(y_0) dy_0.$$

Then $[\underline{\Gamma}(c), \bar{\Gamma}(c)]$ is the closure of the identified set for $\Gamma(f_{Y_1}, f_{Y_0})$. This follows by convexity of the identified set in theorem 7. These bounds are weakly increasing in each component of c . We discuss computation next.

Computation

The identified set $\Theta_0(c) \times \Theta_1(c)$ is infinite dimensional, since it is a set of nonparametric densities. Suppose we are interested in linear functionals $\Gamma(f_1, f_0)$. Even though the corresponding identified set—when it is nonempty—is an interval, we have characterized it via optimization over the infinite dimensional spaces $\Theta_x(c)$. Here we discuss one approach to computing these identified sets in practice. When the baseline model is falsified, this will also allow us to compute approximations to the falsification frontier and the falsification adaptive set. Specifically, we will approximate the infinite dimensional space of densities of interest by a finite dimensional sieve space. Similar approximations of identified sets have been used in Mogstad, Santos, and Torgovitsky (2018), for example. Alternatively, it is likely that the computational approach developed in Christensen and Connault (2019) could be adapted to our setting. Unlike the sieve based approach we consider below, the dimension of their optimization problem does not depend on the precision of the density approximation. We leave the application of their approach to our problem to future work.

For simplicity, let $\mathcal{Y}_x = [0, 1]$ for $x \in \{0, 1\}$. This restriction can be relaxed by transforming the outcome variable to have support in the unit interval. We also assume that for any $x, x' \in \{0, 1\}$ and $z \in \{0, 1\}^L$ the density $f_{Y_x|X,Z}(\cdot | x', z)$ is continuous on $[0, 1]$. Let $\tilde{\mathcal{P}}([0, 1])$ denote this set of

continuous pdfs. Let

$$\begin{aligned}\mathcal{F}^M &= \left\{ f^M(y) = \sum_{m=0}^M a_m b_m(y) : (a_1, \dots, a_M) = \mathbf{a}^M \in \mathcal{A}_M \right\} \\ &= \{ f^M(y) = \mathbf{a}^M \mathbf{b}^M(y) : \mathbf{a}^M \in \mathcal{A}_M \}\end{aligned}$$

where $\mathcal{A}_M \subseteq \mathbb{R}^M$ and $\{b_m(\cdot)\}$ are known basis functions. We assume that every element of $\tilde{\mathcal{P}}([0, 1])$ can be approximated by a sequence of elements in \mathcal{F}_M as $M \rightarrow \infty$. For example, \mathcal{F}^M could be a set of Bernstein polynomials, which can uniformly approximate any continuous function on $[0, 1]$. For Bernstein polynomials,

$$\mathcal{A}_M = \left\{ \mathbf{a}^M : a_m \geq 0, \sum_{m=0}^M a_m = 1 \right\}.$$

Hence \mathcal{A}_M is a closed, bounded, and convex subset of \mathbb{R}^M .

To approximate $\Theta_x(c) = \mathcal{D}_x(c) \cap \mathcal{H}_x$, we approximate each of the sets $\mathcal{D}_x(c)$ and \mathcal{H}_x . We first consider a finite dimensional approximation to the set $\mathcal{D}_x(c)$. Since $\mathcal{Y}_x = [0, 1]$ does not depend on x , we let $\mathcal{D}(c) = \mathcal{D}_x(c)$. Let

$$\begin{aligned}\mathcal{D}_\ell^M(c_\ell) &= \left\{ (f(\cdot | 0), f(\cdot | 1)) = (\mathbf{a}^M(0)' \mathbf{b}^M(\cdot), \mathbf{a}^M(1)' \mathbf{b}^M(\cdot)) : \mathbf{a}^M(0), \mathbf{a}^M(1) \in \mathcal{A}_M, \right. \\ &\quad \left. \begin{aligned} &(\mathbf{a}^M(0) k_0^\ell(c_\ell) - \mathbf{a}^M(1))' \mathbf{b}^M(y) \leq 0 \text{ and} \\ &(\mathbf{a}^M(1) k_1^\ell(c_\ell) - \mathbf{a}^M(0))' \mathbf{b}^M(y) \leq 0 \text{ for all } y \in [0, 1] \end{aligned} \right\}.\end{aligned}$$

Define the set

$$\begin{aligned}\mathcal{A}_\ell^M(c_\ell) &= \left\{ (\mathbf{a}^M(0), \mathbf{a}^M(1)) \in \mathcal{A}_M \times \mathcal{A}_M : \begin{aligned} &(\mathbf{a}^M(0) k_0^\ell(c_\ell) - \mathbf{a}^M(1))' \mathbf{b}^M(y) \leq 0 \text{ and} \\ &(\mathbf{a}^M(1) k_1^\ell(c_\ell) - \mathbf{a}^M(0))' \mathbf{b}^M(y) \leq 0 \text{ for all } y \in [0, 1] \end{aligned} \right\}.\end{aligned}$$

Then

$$\mathcal{D}_\ell^M(c_\ell) = \{(f(\cdot | 0), f(\cdot | 1)) = (\mathbf{a}^M(0)' \mathbf{b}^M(\cdot), \mathbf{a}^M(1)' \mathbf{b}^M(\cdot)) : (\mathbf{a}^M(0), \mathbf{a}^M(1)) \in \mathcal{A}_\ell^M(c_\ell)\}.$$

The set $\mathcal{A}_\ell^M(c_\ell)$ is closed and convex since it is defined by a finite number of weak inequalities. There are a continuum of inequalities, however. In practice, we also approximate $\mathcal{A}_\ell^M(c_\ell)$ as follows. Pick a grid of points $\{y_1, \dots, y_N\} \subseteq [0, 1]$ that becomes dense in $[0, 1]$ as $N \rightarrow \infty$. Then let $\mathcal{A}_\ell^{M,N}(c_\ell)$

be the set of coefficients $\mathbf{a}^M(0)$ and $\mathbf{a}^M(1)$ satisfying

$$\left(\mathbf{a}^M(0)k_0^\ell(c_\ell) - \mathbf{a}^M(1)\right)' \mathbf{b}^M(y_n) \leq 0$$

and

$$\left(\mathbf{a}^M(1)k_1^\ell(c_\ell) - \mathbf{a}^M(0)\right)' \mathbf{b}^M(y_n) \leq 0$$

for all $n \in \{1, \dots, N\}$. This is a finite set of linear inequalities. We keep the N implicit in the notation from here on.

Approximate the overall set $\mathcal{D}(c)$ by the product

$$\mathcal{D}^M(c) = \prod_{\ell=1}^L \mathcal{D}_\ell^M(c_\ell).$$

This set is characterized by

$$\mathcal{A}^M(c) = \prod_{\ell=1}^L \mathcal{A}_\ell^M(c_\ell),$$

a closed, convex, and bounded set in \mathbb{R}^{2ML} .

We similarly approximate \mathcal{H}_x by a finite dimensional set. Let

$$\mathbf{f}_{Y|X,Z_\ell}(\cdot | x) = \begin{pmatrix} f_{Y|X,Z_\ell}(\cdot | x, 0) \\ f_{Y|X,Z_\ell}(\cdot | x, 1) \end{pmatrix} \quad \text{and} \quad (\mathbf{f}_{Y|X,Z}(\cdot | x))_{\text{all } \ell} = \begin{pmatrix} \mathbf{f}_{Y|X,Z_1}(\cdot | x) \\ \vdots \\ \mathbf{f}_{Y|X,Z_L}(\cdot | x) \end{pmatrix}.$$

These conditional pdfs are point identified directly from the data. For each $z \in \{0, 1\}$, we approximate $f_{Y|X,Z_\ell}(\cdot | x, z)$ by

$$\begin{aligned} f_{Y|X,Z_\ell}^M(\cdot | x, z) &= \sum_{m=0}^M \xi_\ell(x, z)_m b_m(\cdot) \\ &= \xi_\ell(x, z)^{M'} \mathbf{b}^M(\cdot) \end{aligned}$$

where $\{b_m(\cdot)\}$ are known basis functions. When these basis functions are orthonormal, so that

$$\int_0^1 b_m(y) b_n(y) dy = \mathbb{1}(m = n),$$

this approximation can be obtained choosing

$$\xi_\ell(x, z)_m = \int_0^1 f_{Y|X,Z_\ell}(y | x, z) b_m(y) dy.$$

More concretely, if we use a Bernstein polynomial approximation, then

$$\xi_\ell(x, z)_m = f_{Y|X, Z_\ell} \left(\frac{m}{M} \mid x, z \right).$$

Here we have approximated the conditional distribution of $Y \mid X, Z$. But \mathcal{H}_x is defined in terms of distributions of $(Y, X) \mid Z$. We can write the vector $(\mathbf{f}_{Y, X|Z}(y, x))_{\text{all } \ell}$ as

$$(\mathbf{f}_{Y, X|Z}(y, x))_{\text{all } \ell} = \mathbf{B}_x (\mathbf{f}_{Y|X, Z}(y \mid x))_{\text{all } \ell}$$

where

$$\mathbf{B}_x = \begin{pmatrix} \mathbb{P}(X = x \mid Z_1 = 0) & 0 & \dots & 0 \\ 0 & \mathbb{P}(X = x \mid Z_1 = 1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathbb{P}(X = x \mid Z_L = 1) \end{pmatrix}.$$

Hence its approximation can be written as

$$(\mathbf{f}_{Y, X|Z}(y, x))_{\text{all } \ell} = \mathbf{B}_x \Xi^M \mathbf{b}^M(y)$$

where

$$\Xi^M = \begin{pmatrix} \xi_1^1(x, 0) & \xi_1^2(x, 0) & \dots & \xi_1^M(x, 0) \\ \xi_1^1(x, 1) & \xi_1^2(x, 1) & \dots & \xi_1^M(x, 1) \\ \vdots & \vdots & \ddots & \vdots \\ \xi_L^1(x, 0) & \xi_L^2(x, 0) & \dots & \xi_L^M(x, 0) \\ \xi_L^1(x, 1) & \xi_L^2(x, 1) & \dots & \xi_L^M(x, 1) \end{pmatrix}.$$

Next, approximate $\mathcal{P}(\mathcal{Y}_x)^{2^L}$ by

$$\{\mathbf{f}(\cdot) = \Gamma \mathbf{b}^M(\cdot) : \gamma_1^M, \dots, \gamma_{2^L}^M \in \mathcal{A}^M\}$$

where

$$\Gamma = \begin{pmatrix} \gamma_1^M \\ \vdots \\ \gamma_{2^L}^M \end{pmatrix}$$

is a $2^L \times M$ matrix of coefficients. Thus we approximate \mathcal{H}_x by

$$\begin{aligned} \mathcal{H}_x^M &= \{\mathbf{f}(y) = (\mathbf{B}_x \Xi^M + \mathbf{A}_x \Gamma) \mathbf{b}^M(y) : \gamma_1, \dots, \gamma_{2^L} \in \mathcal{A}_M\} \\ &= \{\mathbf{f}(y) = \mathbf{E}^M \mathbf{b}^M(y) : \mathbf{E}^M \in \mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2^L}\}. \end{aligned}$$

If \mathcal{A}_M is closed, convex, and bounded, the affine mapping

$$\mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2^L}$$

will also be closed, convex, and bounded.

Thus we approximate the identified set for $((\mathbf{f}_{Y_0|Z_\ell})_{\text{all } \ell}, (\mathbf{f}_{Y_1|Z_\ell})_{\text{all } \ell})$ by

$$\Theta_0^M(c) \times \Theta_1^M(c) = (\mathcal{D}^M(c) \cap \mathcal{H}_0^M) \times (\mathcal{D}^M(c) \cap \mathcal{H}_1^M)$$

where

$$\mathcal{D}^M(c) \cap \mathcal{H}_x^M = \left\{ \mathbf{f}(y) = \mathbf{E}^M \mathbf{b}^M(y) : \mathbf{E}^M \in \mathcal{A}^M(c) \cap (\mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2L}) \right\}.$$

This approximate identified set is empty—and hence the approximate model is falsified—if the set

$$\mathcal{A}^M(c) \cap (\mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2L})$$

is empty for either $x = 0$ or $x = 1$. These two sets are finite dimensional polytopes. Hence it is computationally tractable to compute the distance between them,

$$\text{dist}(\mathcal{A}^M(c), \mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2L}) = \inf_{(\mathbf{v}_1, \mathbf{v}_2) \in \mathcal{A}^M(c) \times (\mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2L})} \|\mathbf{v}_1 - \mathbf{v}_2\|,$$

for any given c . Since they are closed, convex, and bounded, the approximate model will be non-refuted if and only if

$$\text{dist}(\mathcal{A}^M(c), \mathbf{B}_x \Xi^M + \mathbf{A}_x \mathcal{A}_M^{2L}) = 0.$$

This characterization can be used to compute an approximate falsification frontier as well as an approximate falsification adaptive set.

Although we omit a full analysis, we expect the identified sets $\Theta_0^M(c) \times \Theta_1^M(c)$ will converge to $\Theta_0(c) \times \Theta_1(c)$ as $M \rightarrow \infty$ under suitable regularity conditions. Consequently, the approximate falsification frontier and falsification adaptive sets should also converge as $M \rightarrow \infty$. Finally, the value of M can be chosen large enough so that the approximate objects of interest change by less than a preset tolerance for further increases in M .

6 Conclusion

Summary

In this paper we outlined a systematic and constructive answer to the question “What should researchers do when their baseline model is refuted?” Our answer focuses on what can be learned from falsified models, rather than treating falsification as a nuisance to be ignored or as a fatal flaw which dooms a study.

We gave four recommendations. First: Measure the extent of falsification. Do this by defining continuous relaxations of the key identifying assumptions of interest, and then relaxing these assumptions until the model is no longer falsified. We call the set of points at which this happens the falsification frontier. Second: Present the *falsification adaptive set*, the identified set for the parameter of interest under the assumption that the true model lies somewhere on the falsifica-

tion frontier. This second recommendation is a generalization of standard practice for non-refuted baseline models. Moreover, it does not require the researcher to select or calibrate sensitivity parameters. Third: Present the identified set for selected points of interest on the falsification frontier. Fourth: Present identified sets for points beyond the falsification frontier, as a further sensitivity analysis.

We illustrated these four recommendations in two different overidentified instrumental variable models. The first model imposes homogeneous treatment effects but allows for continuous treatments, while the second model allows for heterogeneous treatment effects but focuses on binary treatments. In both models, multiple instruments are observed and the key identifying assumptions are exclusion and exogeneity of each instrument. We considered continuous relaxations of these assumptions. We then characterized the falsification frontier, the falsification adaptive set, and identified sets for points beyond the frontier. In the homogeneous treatment effect model, the falsification adaptive set has a particularly simple closed form expression, depending only on the value of a handful of 2SLS regression coefficients. In the heterogeneous treatment effect model, the falsification adaptive set can be quickly computed using convex optimization.

We showed how to use our results in four substantively different empirical applications. There we emphasized that that falsification adaptive set is an informative complement to traditional overidentification test p -values: It directly summarizes the range of estimates corresponding to non-falsified alternative models. We also emphasized the importance of controlling for the possibly invalid instruments when considering alternative models.

Future Research

In this paper we outlined our general approach and then applied it to the classical linear instrumental variable model. The falsification frontier and falsification adaptive sets, however, must generally be derived separately for each baseline model and class of relaxations from those baseline assumptions. Thus in future work we plan to do this kind of analysis for different baseline models and different classes of relaxations. Here we focused on instrumental variable models. Other possible applications include various kinds of placebo tests, like those discussed in Heckman and Hotz (1989). It may also be possible to derive results for a general class of models, by applying the results of Chesher and Rosen (2017) or Torgovitsky (2019), for example. Such extensions are important since—just like the choice of the baseline assumptions—the choice of relaxation matters. Hence comparing empirical results for different relaxations will be informative.

Finally, in this paper we focused on population level analysis. We briefly discussed estimation and inference in sections 2.5 and 3.4, but a full analysis of this remains for future work.

References

ACEMOGLU, D., S. JOHNSON, AND J. A. ROBINSON (2001): “The colonial origins of comparative development: An empirical investigation,” *American Economic Review*, 91, 1369–1401.

- ALBOUY, D. Y. (2012): “The colonial origins of comparative development: an empirical investigation: comment,” *American Economic Review*, 102, 3059–76.
- ALESINA, A., P. GIULIANO, AND N. NUNN (2013): “On the origins of gender roles: Women and the plough,” *The Quarterly Journal of Economics*, 128, 469–530.
- ALIPRANTIS, C. D. AND K. C. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, Springer, 3rd ed.
- 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.
- 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.
- 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.
- ANDREWS, I., J. ROTH, AND A. PAKES (2019): “Inference for linear conditional moment inequalities,” *arXiv preprint arXiv:1909.10062*.
- ANGRIST, J. AND A. B. KRUEGER (1994): “Why do World War II veterans earn more than nonveterans?” *Journal of Labor Economics*, 12, 74–97.
- ARIEW, R. (2018): “Pierre Duhem,” in *The Stanford Encyclopedia of Philosophy*, ed. by E. N. Zalta, Metaphysics Research Lab, Stanford University, fall 2018 ed.
- ARMSTRONG, T. AND M. KOLESÁR (2019): “Sensitivity analysis using approximate moment condition models,” *Working paper*.
- ASHLEY, R. (2009): “Assessing the credibility of instrumental variables inference with imperfect instruments via sensitivity analysis,” *Journal of Applied Econometrics*, 24, 325–337.
- 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.
- BALKE, A. AND J. PEARL (1997): “Bounds on treatment effects from studies with imperfect compliance,” *Journal of the American Statistical Association*, 92, 1171–1176.
- BAUM-SNOW, N. (2007): “Did highways cause suburbanization?” *The Quarterly Journal of Economics*, 122, 775–805.
- BERESTEANU, A., I. MOLCHANOV, AND F. MOLINARI (2012): “Partial identification using random set theory,” *Journal of Econometrics*, 166, 17–32.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica*, 841–890.
- BERRY, S. T. (1994): “Estimating discrete-choice models of product differentiation,” *The RAND Journal of Economics*, 242–262.
- BLUNDELL, R., A. GOSLING, H. ICHIMURA, AND C. MEGHIR (2007): “Changes in the distribution of male and female wages accounting for employment composition using bounds,” *Econometrica*, 75, 323–363.

- BONHOMME, S. AND M. WEIDNER (2018): “Minimizing sensitivity to model misspecification,” *arXiv preprint arXiv:1807.02161*.
- BONK, T. (2008): *Underdetermination: An Essay on Evidence and the Limits of Natural Knowledge*, vol. 261, Springer.
- BONTEMPS, C., T. MAGNAC, AND E. MAURIN (2012): “Set identified linear models,” *Econometrica*, 80, 1129–1155.
- BORDER, K. C. (1985): *Fixed Point Theorems with Applications to Economics and Game Theory*, Cambridge University Press.
- (2019): “Alternative linear inequalities,” V. 2019.10.23::09.57.
- BOSERUP, E. (1970): *Women’s Role in Economic Development*, George Allen and Unwin Ltd.
- 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.
- BREUSCH, T. S. (1986): “Hypothesis testing in unidentified models,” *Review of Economic Studies*, 53, 635–651.
- BUGNI, F. A., I. A. CANAY, AND X. SHI (2015): “Specification tests for partially identified models defined by moment inequalities,” *Journal of Econometrics*, 185, 259–282.
- CANAY, I. A., A. SANTOS, AND A. M. SHAIKH (2013): “On the testability of identification in some nonparametric models with endogeneity,” *Econometrica*, 81, 2535–2559.
- CANAY, I. A. AND A. M. SHAIKH (2017): “Practical and theoretical advances in inference for partially identified models,” in *Advances in Economics and Econometrics: Eleventh World Congress*, ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson, Cambridge University Press, vol. 2, 271–306.
- CARD, D. (2001): “Estimating the return to schooling: Progress on some persistent econometric problems,” *Econometrica*, 69, 1127–1160.
- CHAN, J. C. AND J. L. TOBIAS (2015): “Priors and posterior computation in linear endogenous variable models with imperfect instruments,” *Journal of Applied Econometrics*, 30, 650–674.
- CHEN, X., C. FLORES, AND A. FLORES-LAGUNES (2016): “Bounds on average treatment effects with an invalid instrument: An application to the Oregon Health Insurance Experiment,” *Working paper*.
- CHERNOZHUKOV, V., H. HONG, AND E. TAMER (2007): “Estimation and confidence regions for parameter sets in econometric models,” *Econometrica*, 75, 1243–1284.
- CHERNOZHUKOV, V., S. LEE, AND A. M. ROSEN (2013): “Intersection bounds: Estimation and inference,” *Econometrica*, 81, 667–737.
- CHESHER, A. AND A. M. ROSEN (2017): “Generalized instrumental variable models,” *Econometrica*, 85, 959–989.

- CHO, J. AND T. M. RUSSELL (2019): “Simple inference on functionals of set-identified parameters defined by linear moments,” *arXiv preprint arXiv:1810.03180*.
- CHRISTENSEN, T. AND B. CONNAULT (2019): “Counterfactual sensitivity and robustness,” *arXiv preprint arXiv:1904.00989*.
- COMPIANI, G. (2018): “Nonparametric demand estimation in differentiated products markets,” *Working paper*.
- CONLEY, T. G., C. B. HANSEN, AND P. E. ROSSI (2012): “Plausibly exogenous,” *The Review of Economics and Statistics*, 94, 260–272.
- COTTERILL, R. W. AND L. E. HALLER (1997): “An econometric analysis of the demand for RTE cereal: Product market definition and unilateral market power effects,” *Food Marketing Policy Center Research Report No. 35*.
- D’HAULTFOEUILLE, X., C. GAILLAC, AND A. MAUREL (2018): “Rationalizing rational expectations? Tests and deviations,” *Working paper*.
- DI TRAGLIA, F. AND C. GARCÍA-JIMENO (2019): “A framework for eliciting, incorporating, and disciplining identification beliefs in linear models,” *Working paper*.
- DURANTON, G., P. M. MORROW, AND M. A. TURNER (2014): “Roads and trade: Evidence from the US,” *Review of Economic Studies*, 81, 681–724.
- FESSLER, P. AND M. KASY (2019): “How to use economic theory to improve estimators: Shrinking toward theoretical restrictions,” *The Review of Economics and Statistics*, 101, 681–698.
- FISHER, F. M. (1961): “On the cost of approximate specification in simultaneous equation estimation,” *Econometrica*, 29, 139–170.
- (2005): *Microeconomics: Essays in Theory and Applications*, Cambridge University Press.
- FLORES, C. AND X. CHEN (2018): *Average Treatment Effect Bounds with an Instrumental Variable: Theory and Practice*, Springer.
- FREYBERGER, J. (2017): “On completeness and consistency in nonparametric instrumental variable models,” *Econometrica*, 85, 1629–1644.
- GAFAROV, B. (2019): “Inference in high-dimensional set-identified affine models,” *arXiv preprint arXiv:1904.00111*.
- GUGGENBERGER, P. (2012): “A note on the (in)consistency of the test of overidentifying restrictions and the concepts of true and pseudo-true parameters,” *Economics Letters*, 117, 901–904.
- GUNDERSEN, C., B. KREIDER, AND J. PEPPER (2012): “The impact of the National School Lunch Program on child health: A nonparametric bounds analysis,” *Journal of Econometrics*, 166, 79–91.
- HAILE, P. A. AND E. TAMER (2003): “Inference with an incomplete model of English auctions,” *Journal of Political Economy*, 111, 1–51.
- HANSEN, L. P. (1982): “Large sample properties of generalized method of moments estimators,” *Econometrica*, 50, 1029–1054.

- HANSEN, L. P., J. HEATON, AND E. G. LUTTMER (1995): “Econometric evaluation of asset pricing models,” *The Review of Financial Studies*, 8, 237–274.
- HANSEN, L. P. AND R. JAGANNATHAN (1991): “Implications of security market data for models of dynamic economies,” *Journal of Political Economy*, 99, 225–262.
- (1997): “Assessing specification errors in stochastic discount factor models,” *The Journal of Finance*, 52, 557–590.
- HAUSMAN, J. A. (1996): “Valuation of new goods under perfect and imperfect competition,” in *The Economics of New Goods*, University of Chicago Press, 209–247.
- HECKMAN, J. J. AND V. J. HOTZ (1989): “Choosing among alternative nonexperimental methods for estimating the impact of social programs: The case of manpower training,” *Journal of the American Statistical Association*, 84, 862–874.
- HECKMAN, J. J. AND S. URZUA (2010): “Comparing IV with structural models: What simple IV can and cannot identify,” *Journal of Econometrics*, 156, 27–37.
- HECKMAN, J. J. AND E. VYTLACIL (1998): “Instrumental variables methods for the correlated random coefficient model: Estimating the average rate of return to schooling when the return is correlated with schooling,” *The Journal of Human Resources*, 33, 974–987.
- HOTZ, V. J., C. H. MULLIN, AND S. G. SANDERS (1997): “Bounding causal effects using data from a contaminated natural experiment: Analysing the effects of teenage childbearing,” *The Review of Economic Studies*, 64, 575–603.
- HSIEH, Y.-W., X. SHI, AND M. SHUM (2017): “Inference on estimators defined by mathematical programming,” *arXiv preprint arXiv:1709.09115*.
- HUBER, M. (2014): “Sensitivity checks for the local average treatment effect,” *Economics Letters*, 123, 220–223.
- HUBER, P. J. AND E. M. RONCHETTI (2009): *Robust Statistics*, John Wiley & Sons.
- IMBENS, G. W. (2003): “Sensitivity to exogeneity assumptions in program evaluation,” *American Economic Review P&P*, 126–132.
- (2010): “Better LATE than nothing: Some comments on Deaton (2009) and Heckman and Urzua (2009),” *Journal of Economic Literature*, 48, 399–423.
- IMBENS, G. W. AND J. D. ANGRIST (1994): “Identification and estimation of local average treatment effects,” *Econometrica*, 62, 467–475.
- KADANE, J. B. AND T. W. ANDERSON (1977): “A comment on the test of overidentifying restrictions,” *Econometrica*, 45, 1027–1031.
- KÉDAGNI, D. AND I. MOURIFIÉ (2019): “Generalized instrumental inequalities: Testing the IV independence assumption,” *Biometrika (forthcoming)*.
- KITAGAWA, T. (2009): “Identification region of the potential outcome distributions under instrument independence,” *Working paper*.

- KRAAY, A. (2012): “Instrumental variables regressions with uncertain exclusion restrictions: A Bayesian approach,” *Journal of Applied Econometrics*, 27, 108–128.
- KREIDER, B. AND J. V. PEPPER (2007): “Disability and employment: reevaluating the evidence in light of reporting errors,” *Journal of the American Statistical Association*, 102, 432–441.
- KREIDER, B., J. V. PEPPER, C. GUNDERSEN, AND D. JOLLIFFE (2012): “Identifying the effects of SNAP (food stamps) on child health outcomes when participation is endogenous and misreported,” *Journal of the American Statistical Association*, 107, 958–975.
- KREIDER, B., J. V. PEPPER, AND M. ROY (2016): “Identifying the effects of WIC on food insecurity among infants and children,” *Southern Economic Journal*, 82, 1106–1122.
- LAFFÉRS, L. (2018): “Bounding average treatment effects using linear programming,” *Empirical Economics*, 1–41.
- (2019): “Identification in models with discrete variables,” *Computational Economics*, 53, 657–696.
- LECHICKI, A. AND A. SPAKOWSKI (1985): “A note on intersection of lower semicontinuous multifunctions,” *Proceedings of the American Mathematical Society*, 95, 119–122.
- LIU, T.-C. (1955): “A simple forecasting model for the US economy,” *International Monetary Fund Staff Papers*, 4, 434–466.
- (1960): “Underidentification, structural estimation, and forecasting,” *Econometrica*, 28, 855–865.
- (1963): “Structural estimation and forecasting: A critique of the Cowles commission method,” *Tsing Hua Journal of Chinese Studies*, 4, 152–171.
- LUDVIGSON, S. C. (2013): “Advances in consumption-based asset pricing: Empirical tests,” in *Handbook of the Economics of Finance*, Elsevier, vol. 2, 799–906.
- MACHADO, C., A. SHAIKH, AND E. VYTLACIL (2019): “Instrumental variables and the sign of the average treatment effect,” *Journal of Econometrics*, 212, 522–555.
- MANSKI, C. F. (1990): “Nonparametric bounds on treatment effects,” *American Economic Review P&P*, 80, 319–323.
- (2003): *Partial Identification of Probability Distributions*, Springer.
- MANSKI, C. F. AND J. V. PEPPER (2000): “Monotone instrumental variables: With an application to the returns to schooling,” *Econometrica*, 68, 997–1010.
- (2009): “More on monotone instrumental variables,” *The Econometrics Journal*, 12, S200–S216.
- (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.
- MANSKI, C. F. AND E. TAMER (2002): “Inference on regressions with interval data on a regressor or outcome,” *Econometrica*, 70, 519–546.

- MASTEN, M. A. AND A. POIRIER (2018a): “Identification of treatment effects under conditional partial independence,” *Econometrica*, 86, 317–351.
- (2018b): “Salvaging falsified instrumental variable models,” *arXiv preprint arXiv:1812.11598*.
- (2019): “Inference on breakdown frontiers,” *Quantitative Economics (forthcoming)*.
- MATZKIN, R. L. (2007): “Nonparametric identification,” *Handbook of Econometrics*, 6, 5307–5368.
- MOGSTAD, M., A. SANTOS, AND A. TORGOVITSKY (2018): “Using instrumental variables for inference about policy relevant treatment parameters,” *Econometrica*, 86, 1589–1619.
- MOLINARI, F. (2019): “Econometrics with partial identification,” *The Handbook of Econometrics (forthcoming)*.
- MOON, H. R. AND F. SCHORFHEIDE (2012): “Bayesian and frequentist inference in partially identified models,” *Econometrica*, 80, 755–782.
- NEVO, A. (2000): “Mergers with differentiated products: The case of the ready-to-eat cereal industry,” *The RAND Journal of Economics*, 31, 395–421.
- (2001): “Measuring market power in the ready-to-eat cereal industry,” *Econometrica*, 69, 307–342.
- (2003): “New products, quality changes, and welfare measures computed from estimated demand systems,” *The Review of Economics and Statistics*, 85, 266–275.
- (2011): “Empirical models of consumer behavior,” *Annual Reviews of Economics*, 3, 51–75.
- NEVO, A. AND A. M. ROSEN (2012): “Identification with imperfect instruments,” *The Review of Economics and Statistics*, 94, 659–671.
- NEWKEY, W. K. (1985): “Generalized method of moments specification testing,” *Journal of Econometrics*, 29, 229–256.
- PARENTE, P. M. AND J. SANTOS SILVA (2012): “A cautionary note on tests of overidentifying restrictions,” *Economics Letters*, 115, 314–317.
- POPPER, K. R. (1934): *The Logic of Scientific Discovery*.
- PRYOR, F. L. (1985): “The invention of the plow,” *Comparative Studies in Society and History*, 27, 727–743.
- RAMSAHAI, R. R. (2012): “Causal bounds and observable constraints for non-deterministic models,” *Journal of Machine Learning Research*, 13, 829–848.
- ROCKAFELLAR, R. T. (1970): *Convex Analysis*, Princeton University Press.
- SARGAN, J. D. (1958): “The estimation of economic relationships using instrumental variables,” *Econometrica*, 26, 393–415.
- SLICHTER, D. (2014): “Testing instrument validity and identification with invalid instruments,” *Working paper*.

- SMALL, D. S. (2007): “Sensitivity analysis for instrumental variables regression with overidentifying restrictions,” *Journal of the American Statistical Association*, 102, 1049–1058.
- SPANOS, A. (2018): “Mis-specification testing in retrospect,” *Journal of Economic Surveys*, 32, 541–577.
- STANFORD, K. (2017): “Underdetermination of scientific theory,” in *The Stanford Encyclopedia of Philosophy*, ed. by E. N. Zalta, Metaphysics Research Lab, Stanford University, winter 2017 ed.
- SUNDARAM, R. K. (1996): *A First Course in Optimization Theory*, Cambridge University Press.
- 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.
- TEBALDI, P., A. TORGOVITSKY, AND H. YANG (2019): “Nonparametric estimates of demand in the California Health Insurance Exchange,” *Working paper*.
- TORGOVITSKY, A. (2019): “Partial identification by extending subdistributions,” *Quantitative Economics*, 10, 105–144.
- VAN KIPPERSLUIJ, H. AND C. A. RIETVELD (2017): “Pleiotropy-robust Mendelian randomization,” *International Journal of Epidemiology*, 47, 1279–1288.
- (2018): “Beyond plausibly exogenous,” *The Econometrics Journal*, 21, 316–331.
- WHITE, H. (1982): “Maximum likelihood estimation of misspecified models,” *Econometrica*, 50, 1–25.
- (1994): *Estimation, Inference and Specification Analysis*, Cambridge University Press.
- WOOLDRIDGE, J. M. (1997): “On two stage least squares estimation of the average treatment effect in a random coefficient model,” *Economics Letters*, 56, 129–133.
- (2003): “Further results on instrumental variables estimation of average treatment effects in the correlated random coefficient model,” *Economics Letters*, 79, 185–191.
- (2008): “Instrumental variables estimation of the average treatment effect in the correlated random coefficient model,” in *Modelling and Evaluating Treatment Effects in Econometrics (Advances in Econometrics, Volume 21)*, ed. by T. Fomby, R. C. Hill, D. L. Millimet, J. A. Smith, and E. J. Vytlačil, Emerald Group Publishing Limited, 93–116.

Appendices to Accompany “Salvaging Falsified Instrumental Variable Models”

Matthew A. Masten
Duke University

Alexandre Poirier
Georgetown University

Table of Contents

This document contains twelve appendices:

- A. [A detailed review](#) of how our results relate to the previous literature.
- B. [A discussion](#) of five alternative responses to falsification from the literature.
- C. [An alternative approach](#) to defining falsification points.
- D. [A discussion](#) of Bayesian approaches to falsification, and how it relates to our results.
- E. [An analysis](#) of the implications of using affine transformations of one’s original instruments, in the the linear model.
- F. [An extension](#) of our linear model analysis to allow for covariates. Here we consider both linear and nonparametric controls. We then use these results to show how to use our approach to falsification to analyze subgroups.
- G. [A formal explanation](#) that exclusion and exogeneity are technically equivalent in linear models, although they are substantively distinct.
- H. [An analysis](#) of the 2SLS estimand in falsified models, and a comparison of it with the falsification adaptive set.
- I. [A discussion](#) of how to interpret non-falsified models, and how our results relate to the previous literature on this topic.
- J. [Additional empirical results](#) for the Duranton et al. (2014) application.
- K. [Proofs](#) for section 3.
- L. [Proofs](#) for section 5.

A Related Literature

In this section, we review the related literature.

Falsification

We begin with the literature on falsification. Typically a model entails many different assumptions. Although we can sometimes falsify all assumptions jointly, it is often impossible to know *which* specific assumption is false. This result is known in philosophy as the Duhem-Quine thesis (Bonk 2008, Stanford 2017). For example, Ariew (2018, section 2.1) states that “the thesis formed from the two sub-theses, that (i) since empirical statements are interconnected, they cannot be singly disconfirmed, and (ii) if we wish to hold a particular statement true we can always adjust another statement, has become known as the Duhem-Quine thesis.” That is, the general idea that one can consider trade-offs between assumptions to avoid falsification of any given assumption is well known in philosophy. Our paper builds on this literature by providing a specific formalization of this trade-off using identification theory.

In econometrics, three papers are most closely related. Of these, the conceptually closest is Manski and Pepper (2018). They study identification of treatment effects under combinations of different “bounded variation” assumptions, which are continuous relaxations of certain baseline assumptions. In an empirical analysis of right-to-carry laws, they compute identified sets for various combinations of these relaxations; see their table 2. Our recommended responses to baseline falsification can be read from this table. First, the boundary between the dark gray identified sets in that table and the rest is essentially a discrete approximation to the falsification frontier in their model. Second, the collection of the singleton identified sets on this boundary is what we call the falsification adaptive set. Third, given all of the separate presented singleton identified sets on the frontier, one can focus on any single set of interest. Finally, they also present identified sets beyond the frontier. In our paper, we formally define the falsification frontier and falsification adaptive set in a general model and relate it to the problem of responding to specification test rejection. We emphasize that these concepts and the four responses apply to any falsified model, not just the specific one considered by Manski and Pepper (2018). Finally, we illustrate these responses using refutable instrumental variable models, whereas Manski and Pepper (2018) studied a model with cross sectional and time series variation in treatment rather than variation due to instrumental variables.

Ramsahai (2012) studies a heterogeneous treatment effect model with a binary outcome, a binary treatment, and a binary instrument. He defines a continuous relaxation of the instrument exogeneity assumption and then shows how to numerically compute identified sets for a single value of this relaxation. He notes that the model can be falsified even if exogeneity does not hold exactly, and that the model is not falsified if instrument exogeneity is completely relaxed. He does not formally define and derive the falsification point, however. He also does not consider multi-dimensional relaxations. On pages 842–843, he notes that “it is not obvious how the methods described in [his] paper can be extended to compute bounds” as a function of his relaxation. In our analysis, we show how to compute bounds for any value of our continuous relaxation, and thus also are able to derive falsification points, falsification frontiers, and the falsification adaptive set. Unlike him, we also consider multiple instruments as well as models for continuous outcomes and continuous treatments.

Machado, Shaikh, and Vytlacil (2019) study the testable implications of instrument exogeneity, sometimes combined with various monotonicity assumptions. They only consider the case where the outcome, treatment, and instrument are all binary. Their proposition 4.1 characterizes the set of

distributions of unobservables (potential outcomes, potential treatments, and the instrument) which falsify instrument exogeneity. In their appendix A, they illustrate this result by parameterizing the distribution of unobservables. From the set of falsified distributions of unobservables, they pick the distribution which has the smallest relaxation of the baseline assumption, in a certain sense. They call this the “minimal detectable violation.” This is equivalent to what we call a falsification point. We explain this equivalence formally in our appendix C. Beyond this similarity, however, there are several main differences between our paper and theirs. First, they only consider the case where outcomes are binary, whereas we also consider the continuous outcome case. Second, even in the binary outcome case their falsification point analysis is based on parametric index model assumptions. Our analysis of that case does not use such assumptions. Third, they do not consider the linear case, which allows for continuous treatments. Fourth, they do not consider the multiple instrument case. Finally, unlike their paper, we characterize identified sets under relaxations of the baseline assumptions. This allows us to characterize the falsification adaptive set, as well as identified sets beyond the falsification frontier. In particular, this allows researchers to do sensitivity analysis even when the baseline model is not falsified.

Before these three papers, a large literature in asset pricing beginning with Hansen and Jagannathan (1991, 1997) developed in response to refutation of models via overidentification tests. For example, Hansen, Heaton, and Luttmer (1995, page 239) state that their goal “is to shift the focus of statistical analysis of asset pricing models away from whether the models are correctly specified and *toward measurement of the extent to which they are misspecified*” (emphasis added). Our paper is motivated by the same concern—the falsification frontier provides one way of measuring the extent to which one’s baseline model is misspecified. Besides this shared motivation, however, the technical approach in our paper is substantially different from that taken in this asset pricing literature. For an explanation and further discussion of the approach in that literature, see Ludvigson (2013, section 3.3).

Recent work by D’Haultfoeuille, Gaillac, and Maurel (2018) also has a similar motivation. They characterize the constraints that the rational expectations assumption places on the marginal distributions of realized outcomes and subjective expectations. When rational expectations is refuted, they define minimal deviations from the rational expectations assumption. They then propose using these minimal deviations as an input to structural estimation. Again, besides a shared motivation, the technical approach in our paper is substantially different. We also study a different setting, instrumental variable models.

Finally, many researchers have studied falsifiable partially identified models. As with point identified models, however, the existing literature focuses on testing the baseline model. For example, see Bontemps, Magnac, and Maurin (2012) and Bugni, Canay, and Shi (2015). The results in Chernozhukov et al. (2013) can also be used for this purpose. The falsification frontier we develop builds on such baseline tests by characterizing how much the baseline assumptions must be relaxed before the model is no longer falsified.

The Testable Implications of Instrumental Variable Models

A large literature has studied the testable implications of instrumental variable models, especially for models with heterogeneous treatment effects. Flores and Chen (2018) and Swanson, Hernán, Miller, Robins, and Richardson (2018) provide excellent surveys. Here we discuss the results most related to our paper.

Our analysis of the heterogeneous treatment effect case starts from Manski (1990). He derived identified sets on average treatment effects assuming potential outcomes are mean independent of the instrument. When the outcome, treatment, and instrument are all binary, Balke and Pearl

(1997) characterize when Manski’s bounds are empty, and hence when the model is falsified.⁸ Kitagawa (2009, proposition 3.1) generalizes this characterization to allow for continuous outcomes, still requiring the treatment and instrument to be binary. As Kitagawa (2009) notes, his extension is an adaptation of corollary 2.2.1 in Manski’s (2003) analysis of missing data. Beresteanu, Molchanov, and Molinari (2012, proposition 2.4) further generalize this characterization to allow for continuous instruments and discrete treatment, for discrete or continuous outcomes. Kédagni and Mourifié (2019, proposition 1) provide an alternative characterization when instruments and outcomes are continuous, treatment is binary, and under the stronger assumption that the instrument is independent of the potential outcomes jointly; also see proposition 2.5 of Beresteanu et al. (2012) for a result under this stronger independence assumption.

Our paper builds on these results in two ways. First, and most importantly: These papers only consider falsification of the baseline model with valid instruments. We consider falsification of a class of models which nests the baseline case along with continuous relaxations of the baseline assumptions. This allows us to define and derive the falsification frontier and falsification adaptive set. Second, for non-falsified models, our results allow researchers to do sensitivity analysis by relaxing the instrument exogeneity assumption.

Our paper complements these results in one subtle way: These previous papers focus on the single instrument case whereas we study multiple instruments. If one assumes that all instruments are jointly independent of potential outcomes, and that instruments are all discrete, then the multiple instrument case can be reduced to a single instrument, and thus the previous results apply.⁹ If even one of the instruments is invalid, however, this new instrument will also typically be invalid. For this reason, we purposely do not combine multiple instruments into one new instrument. This allows us to relax each instrument exogeneity assumption separately. In our baseline model, we only assume that each instrument is independent of each potential outcome, rather than assuming joint independence of all instruments with each potential outcome. Thus our baseline model with multiple instruments is slightly weaker than the model obtained by combining the instruments into a single discrete instrument and then applying one of the papers mentioned above. Hence even the testable implications we derive in our baseline case are apparently new, although this is a minor point since our primary contribution is the falsification frontier analysis rather than the baseline analysis.

Finally, a large literature on the testable implications of instrument exogeneity combined with other assumptions has developed. Most notably, many papers have studied the testable implications of the monotonicity assumption of Imbens and Angrist (1994). Flores and Chen (2018) give a comprehensive review. In this paper we focus on instrument exogeneity only.

Sensitivity Analysis in Homogeneous Effect Instrumental Variable Models

In a series of papers from the mid-1950’s to early 1960’s, Ta-Chung Liu criticized the identifying assumptions used in the Cowles Commission approach to simultaneous equations models (Liu 1955, 1960, 1963). Among other criticisms, he argued that exclusion restrictions are unlikely to

⁸Balke and Pearl (1997) assume the instrument is independent of the potential outcomes jointly, whereas Manski (1990) only assumed the instrument is independent of each potential outcome separately. (Here we suppose outcomes are binary so that mean independence is equivalent to statistical independence.) This difference does not affect whether the identified set is empty or not, given any fixed distribution of observables. Hence it does not change the testable implications of the model. When the identified set is not empty, however, this difference *can* affect the size of the identified set. See the second paragraph of section 3 in Swanson et al. (2018) for further discussion.

⁹Another way to directly apply previous results is by looking at one instrument at a time, obtaining bounds using that instrument only, and then intersecting each of those instrument-specific bounds. This generally does not yield sharp bounds, however.

hold exactly, and thus no valid instruments actually exist. Frank Fisher called this a “disturbing argument,” since “its premises apparently cannot be doubted and because its conclusions, if accepted, imply that the hope of structural estimation by any techniques whatsoever is forlorn indeed” (Fisher 1961, page 139). Fisher countered this argument, however, by providing perhaps the first sensitivity analysis of an instrumental variable model. He considered a linear model where the supposedly excluded variable actually has a nonzero coefficient, with the magnitude of this coefficient measuring the scale of the deviation from the baseline assumption. He showed that the bias of k -class estimators—which includes 2SLS—is continuous in this deviation from the exclusion restriction. He thus argued that it suffices that failures of the exclusion restriction are sufficiently small—it is not necessary for exclusion to hold exactly. Although Fisher called this paper his “most important contribution to econometrics” (Fisher 2005, page 545), it apparently was mostly forgotten for over 40 years.

Around the early 2000’s, however, several papers renewed the study of sensitivity analysis in linear instrumental variable models. This includes Angrist and Krueger (1994), Altonji, Elder, and Taber (2005), Small (2007), Conley et al. (2012), Ashley (2009), Kraay (2012), Ashley and Parmeter (2015), and van Kippersluis and Rietveld (2017, 2018). In section 3, we study the same classical linear model as all of these papers, along with similar relaxations of the baseline assumptions. In that section, our main contribution is the analysis of falsification points, frontiers, and the falsification adaptive set. In contrast, this previous literature has focused on sensitivity analysis for statistical inference on the treatment variable coefficient. Our results build on Small’s (2007) equation (8), which characterizes the set of instrument coefficients consistent with the data. Small focuses on the implications of this characterization for *non-rejection* of the overidentification test (see his proposition 1), and then uses this result to do sensitivity analysis for the treatment effect. In contrast, we focus on the implications of this characterization for falsification.

Sensitivity Analysis in Heterogeneous Effect Instrumental Variable Models

Unlike the literature on models with homogeneous treatment effects, there are far fewer papers on sensitivity analysis in instrumental variable models with heterogeneous treatment effects. Specifically, few papers consider continuous relaxations of the baseline instrumental variable assumptions while still allowing for heterogeneous treatment effects.

An important early paper was Hotz, Mullin, and Sanders (1997), who use a mixture model to allow for relaxations of the baseline assumptions. There are three main differences with our paper: First, they only study sensitivity, not falsification. Second, their mixture relaxation is different from the kind of relaxation we consider. Finally, they focus on the average effect of treatment on the treated, whereas our sensitivity analysis allows for a broader set of parameters of interest.

When all variables are discrete, it is well known that identified sets can often be computed using linear programming. For example, see the literature review in Torgovitsky (2019). This approach has been used in several papers to do sensitivity analysis. One paper is Ramsahai (2012), which we already discussed above. Laffers (2019, section 4) considers continuous relaxations of instrument exogeneity. He then computes identified sets for ATE for several values of this relaxation. In Laffers (2018) he applies this approach to various additional forms of continuous relaxations. These three papers all require all variables to be discrete. A key contribution of our paper is that our results allow for continuous outcome variables. Furthermore, these papers focus on sensitivity analysis, rather than falsification (see our discussion of Ramsahai 2012 above, however).

Finally, the paper by Huber (2014) is somewhat related. He studies deviations from the baseline instrument exogeneity, exclusion, and monotonicity assumptions. These deviations are not relaxations of the baseline assumption; rather, they point identify LATE given the values of the

sensitivity parameters. He does not consider parameters besides LATE. Moreover, his sensitivity analysis does not account for the possibility that the model is falsified. Hence his analysis is substantially different from ours.

Discrete Relaxations of Instrumental Variable Assumptions

Above we focused on the previous literature which studied continuous relaxations of the baseline instrumental variables assumptions, like we do. Several papers instead consider discrete relaxations. Manski and Pepper (2000, 2009) propose and study the monotone instrumental variable assumption. Blundell, Gosling, Ichimura, and Meghir (2007) use this assumption in their study of wage distributions. Kreider and Pepper (2007), Kreider, Pepper, Gundersen, and Jolliffe (2012), Gundersen, Kreider, and Pepper (2012), and Kreider, Pepper, and Roy (2016) use this assumption combined with assumptions on measurement error to do sensitivity analysis in various empirical applications. Chen, Flores, and Flores-Lagunes (2016) also study identification under monotonicity type assumptions on the instrument. Nevo and Rosen (2012) suppose the correlation between the instrument and the unobservable is of the same sign and smaller than the correlation between treatment and the unobservable.

As with the other papers discussed above, these papers’ primary goal is to obtain identified sets on various parameters under assumptions weaker than the standard instrumental variables assumptions. While our results allow for similar analyses, our primary focus is on falsified models.

Local Versus Global Approaches to Sensitivity Analysis

Several papers use local asymptotics to study various kinds of sensitivity to misspecification. For example, see Andrews, Gentzkow, and Shapiro (2017), Bonhomme and Weidner (2018), and Armstrong and Kolesár (2019). The local asymptotic approach assumes the baseline model is approximately correct, in the sense that the magnitude of model misspecification is about the same size as the magnitude of sampling uncertainty. When analyzing falsifiable models, this means that the local approach is designed to handle misspecification which is difficult to detect with specification tests. For example, Bonhomme and Weidner (2018) cite the following quote from Huber and Ronchetti (2009, page 294):

“[such local] neighborhoods make eminent sense, since the standard goodness-of-fit tests are just able to detect deviations of this order. *Larger deviations should be taken care of by diagnostic and modeling*, while smaller ones are difficult to detect and should be covered (in the insurance sense) by [local] robustness.” [emphasis added]

In this paper we focus on models which are clearly falsified—the “larger deviations” which Huber and Ronchetti refer to. In these larger deviations, it is known that the model is not approximately correct. For this reason we take a global approach to sensitivity analysis, which does not rely on linking the size of model misspecification to the size of sampling uncertainty.

Finally, note that this distinction is empirical relevant since many baseline models are strongly rejected in practice. For example, both Andrews et al. (2017) and Armstrong and Kolesár (2019) study the classic automobile data used by Berry, Levinsohn, and Pakes (1995) to estimate a baseline random coefficients logit model. This model is strongly rejected by the data, however. Specifically, Armstrong and Kolesár (2019) report that “The J -statistic for testing the hypothesis that all moments are correctly specified equals 404.7. Consequently, the hypothesis is rejected at the usual significance levels” (page 34). We deal with empirical findings of large misspecification like this by taking a global approach.

B Alternative Responses to Falsification

In this paper we provided four constructive responses to refutation of a baseline model. In this section, we briefly discuss five alternative responses from the literature.

1. *Present the no assumption bounds.* This response is nested as a special case of our analysis. Specifically, in the heterogeneous treatment effects model, the no assumption bounds are obtained by allowing c_ℓ to be sufficiently large for all $\ell \in \{1, \dots, L\}$. In the homogeneous treatment effects model, the no assumption bounds are $(-\infty, \infty)$, which is obtained as $\delta_\ell \rightarrow \infty$ for all $\ell \in \{1, \dots, L\}$.
2. *Present identified sets under alternative, discrete relaxations of the baseline assumptions.* For example, one could use the analysis of Manski and Pepper (2000, 2009) or Nevo and Rosen (2012). Even these relaxations are falsifiable, however. When these weaker assumption are falsified, researchers again face the question of what to do. Our identification assumptions are not nested with these alternative relaxations. Hence they are complementary and can be used together with these previous results.
3. *Focus on a non-falsified subgroup.* We discuss this point formally in section F, so here we give an informal explanation. Suppose a vector of covariates W is observed. Then the model may be refuted conditional on some values $w \in \text{supp}(W)$ but not conditional on other values. So one could present results for those subgroups where the model is not refuted. This is reasonable under either of two assumptions:
 - (a) You assume treatment effects do not depend on covariates. Thus the subgroup effect equals the effect for the entire population.
 - (b) You are specifically interested in a certain subgroup, which happens to be one which does not refute the model, and you are not interested in other subgroups.

If, however, you care about other subgroups, or the entire population, and you want to allow for coefficients that are functions of the covariates, then it is not sufficient to focus on non-falsified subgroups. In section F we formally discuss how you can use our results to handle the subgroups where the model is refuted and then aggregate up to get identified sets for an overall treatment effect.

4. *Use a specific estimator to redefine the parameter of interest.* For example, in a fully parametric model one could declare the maximum likelihood estimand to be the parameter of interest. When the model is falsified, this parameter is called the pseudo-true parameter, while its estimator is called the quasi-MLE. Note that this approach absolves the researcher from any need to do specification testing, since interest in the parameter is motivated by the estimator rather than an underlying model. While this approach is sometimes used, it was not advocated in White’s original analysis of quasi-MLE. To the contrary, he used quasi-MLE as a tool for constructing specification tests, emphasizing that

“[if] one has an indication that the parameter estimator is inconsistent for the parameters of interest” then “the model specification must be carefully re-examined.”
(White 1982, pages 16–17)

He reaches a similar conclusion in his 1994 book:

“In many cases, the goal of empirical economic research is to test hypotheses about parameters to which one wishes to attribute economic meaning. It is our view that this is inappropriate and unjustified without first establishing that the model within which the hypotheses are being tested is correctly specified to at least some extent. Otherwise, one may only have confidence that one is testing hypotheses about [the pseudo-true parameter] θ^* ; the economic interpretation desired is untenable.” (White 1994, pages 346–347)

We concur with this sentiment and focus on methods which hold the parameter of interest fixed but vary the model assumptions to avoid falsification.

5. *Drop the research and move on to a different project.* This response has three problems: (a) It creates a kind of publication bias, since we would never observe analyses with refuted baseline models. (b) It prevents us from answering certain empirical questions when there are no good alternative models or datasets. And (c) It treats all falsified models equally; in contrast, our analysis distinguishes between models which are just barely falsified and those which are strongly falsified.

C An Alternative Definition of the Falsification Point

In this section we discuss an alternative method of defining the falsification point or frontier. This alternative approach is motivated by the analysis in appendix A of Machado et al. (2019). For concreteness, we focus on the binary outcome heterogeneous treatment effects model of section 5.1. Let

$$\mathcal{Q} = \{F_{Y_1, Y_0, X, Z} : \text{This distribution is consistent with the observed data } F_{Y, X, Z}\}.$$

That is, $Q \in \mathcal{Q}$ if and only if $Q_{X, Z} = F_{X, Z}$ and

$$Q_{Y_x | X, Z}(\cdot | x, \cdot) = F_{Y | X, Z}(\cdot | x, \cdot)$$

for each $x \in \{0, 1\}$, so that equation (11) holds. For each $Q \in \mathcal{Q}$, define

$$c_x(Q) = \sup_{y \in \{0, 1\}} |\mathbb{P}_Q(Z = 1 | Y_x = y) - \mathbb{P}_Q(Z = 1)|$$

Then compute

$$c_x^* = \inf_{Q \in \mathcal{Q}} c_x(Q).$$

Proposition 9. Consider the model of section 5.1. The falsification point is $c^* = \max\{c_0^*, c_1^*\}$.

In sections 3 and 5, we obtain falsification frontiers by characterizing identified sets under a given relaxation of the assumptions, and then checking when they are empty. A benefit of this approach is that we can present identified sets for values of the sensitivity parameters on or beyond the falsification frontier. The approach discussed in the present section provides the falsification point but does not characterize identified sets for values $c \geq c^*$.

Proof of proposition 9. Let $c \geq c^*$. Then, $c \geq c_x^*$ for $x \in \{0, 1\}$. Thus, by definition of c_x^* , there exists a distribution $Q^x \in \mathcal{Q}$ under which Z is c -dependent with Y_x , for each $x \in \{0, 1\}$. Define $Q_{Y_1, Y_0, X, Z}$ as follows: Let $Q_{X, Z} = F_{X, Z}$. Next, define

$$Q_{Y_1, Y_0 | X, Z} = Q_{Y_1 | X, Z}^1 \cdot Q_{Y_0 | X, Z}^0.$$

Then (a) $Q \in \mathcal{Q}$, so that it is consistent with the data, and (b) under Q , Z is c -dependent with Y_x for each $x \in \{0, 1\}$. Thus the model under partial exogeneity with c -dependence is not falsified.

Next suppose $c < c^*$. Then $c_x(Q) > c$ for all $Q \in \mathcal{Q}$. That is, all distributions of the unobservables that are consistent with the data have

$$\sup_{y \in \{0,1\}} |\mathbb{P}_Q(Z = 1 \mid Y_x = y) - \mathbb{P}_Q(Z = 1)| > c$$

for each $x \in \{0, 1\}$, and so cannot satisfy c -dependence. Thus the model under partial exogeneity with c -dependence is falsified. \square

D A Bayesian Perspective on Falsification

In this section we briefly discuss how one might respond to falsification of a baseline model from a Bayesian perspective. To do this, it is first helpful to describe our analysis from a different perspective than that used in the main paper. After that, we describe an alternative Bayesian perspective. We then compare the two approaches.

For concreteness, we focus on the classical linear instrumental variable model of section 3. Several papers study sensitivity analysis in this model from a Bayesian perspective, including Ashley (2009), Conley et al. (2012), Kraay (2012), and Chan and Tobias (2015). Also see DiTraglia and García-Jimeno (2019) for a detailed analysis of priors in instrumental variable models. Here we focus on the problem of falsification.

An Alternative Perspective of Our Analysis

Suppose we impose A1, A2, and A3, but we make no assumptions about γ . Then the model is no longer falsifiable. Instead, there is always a nonempty identified set for (β, γ) . Proposition 15 in appendix I below shows that the identified set for γ is a line in \mathbb{R}^L . Figure 11 shows two examples in the two instrument case. Falsification of the baseline model occurs when this line does not pass through the origin.

For simplicity, focus on the two instrument case, as depicted in figure 11. Let γ_1^* denote the horizontal intercept of this line. Let γ_2^* denote the vertical intercept. Then $\delta_1^* = \gamma_1^*$ and $\delta_2^* = \gamma_2^*$, where δ_1^* and δ_2^* are the falsification points defined in section 3.2. The falsification frontier corresponds to the line segment with endpoints $(0, \gamma_2^*)$ and $(\gamma_1^*, 0)$.

Each point in the identified set for γ corresponds to a unique value of β , namely, the value b that satisfies

$$\gamma = \psi - b\Pi.$$

See proposition 15. Thus the falsification adaptive set is simply the values of b corresponding to γ 's in the line segment between $(0, \gamma_2^*)$ and $(\gamma_1^*, 0)$. As we emphasize in our fourth recommendation, researchers may want to do sensitivity analysis beyond the falsification frontier. That is, they may want to start by presenting the falsification adaptive set, but also present larger identified sets corresponding to values of δ that lie beyond the falsification frontier. In this case, the identified set for β simply expands to include values of b that correspond to γ 's beyond the line segment between $(0, \gamma_2^*)$ and $(\gamma_1^*, 0)$.

A Bayesian Perspective

With that as background, we next discuss a Bayesian approach. The baseline model corresponds to a dogmatic prior that places all of its mass on $\gamma = 0$. Instead, suppose we specify a prior π on γ

which is centered at zero, but has full support. Suppose now that we observe the distribution of the data $F_{Y,X,Z}$. And we see that this distribution is incompatible with the baseline model. That is, the identified set for γ described in proposition 15 does not pass through the origin. A prior which was dogmatic on zero would be refuted. But a full support prior would not. Instead, we would update our prior so that our posterior will place support only on the identified set for γ . Given that our prior had full support, our posterior will have full support on this line in \mathbb{R}^L . Since each point on this line corresponds to a unique value of β , this posterior on γ will imply a posterior on β . This posterior will have full support on \mathbb{R} , but of course will give larger weight to some values than to others.

Note that there is no need to specify a prior on β , since we are working with population level data. The only uncertainty that is being captured by the prior is the model uncertainty over the precise value of γ , rather than sampling uncertainty. That said, the prior on γ implies a prior on β , which depends on the values ψ and Π . These values themselves come from the data. Thus, before seeing the population data, it may still be useful to consider one's prior on β and on the values of ψ and Π when thinking about the choice of the prior on γ . We leave a full analysis of prior selection to future work.

A Comparison and Integration of Both Approaches

Now that we have described both approaches, let's compare them. As noted above, the Bayesian approach yields a posterior for β that has full support on \mathbb{R} . Assuming the prior on γ places most of its mass near zero, the posterior on β will put more mass on the values corresponding to γ 's closer to the origin. Thus means that it will generally put large amount of mass on the falsification adaptive set. In general, researchers taking a Bayesian perspective may find it useful to compute the posterior probability that β lies within the falsification adaptive set. If we view the falsification adaptive set as a partial summary of the identified set for γ , reporting this posterior probability can help researchers understand which aspects of the conclusions come from the data and which come from the prior. This is similar to Moon and Schorfheide's (2012) suggestion that researchers reported estimated identified sets along with Bayesian posteriors.

As researchers relax their model beyond the falsification frontier, the identified set for β will expand beyond the falsification adaptive set. Bayesian researchers may also find it helpful to compare the size of credible sets to the falsification adaptive set. More generally, they could compute a kind of Bayesian falsification frontier: For $\alpha \in (0, 1)$, what is the set of δ 's that yield a posterior for β which places probability $1 - \alpha$ on the identified set $\mathcal{B}(\delta)$?

In this section we have briefly compared and contrasted our approach with a Bayesian approach, and shown that they can complement each other. There are many open questions, however, like how to choose priors and how to analyze finite sample versions of this analysis. In particular, in finite samples we never know for sure that our model is falsified. Thus it will be interesting to see how Bayesian shrinkage methods, like those in Fessler and Kasy (2019), can be applied to analyze falsifiable models. We leave this and related questions to future work.

E Affine Transformations of the Instruments in the Linear Model

In this section we discuss the implications of using affine transformations of the original instruments Z . Specifically, let A denote an invertible $L \times L$ matrix. Let $b \in \mathbb{R}^L$. Then $\tilde{Z} = AZ + b$ is an affine transformation of Z . Let

$$\tilde{\psi} = \text{var}(\tilde{Z})^{-1} \text{cov}(\tilde{Z}, Y) \quad \text{and} \quad \tilde{\pi} = \text{var}(\tilde{Z})^{-1} \text{cov}(\tilde{Z}, X).$$

Here we assume $K = 1$ for simplicity.

Proposition 10. Suppose A is diagonal. Then

$$\frac{\tilde{\psi}_\ell}{\tilde{\pi}_\ell} = \frac{\psi_\ell}{\pi_\ell}.$$

Hence the FAS using Z as instruments equals the FAS using \tilde{Z} as instruments.

Proof of proposition 10. We have

$$\begin{aligned} \tilde{\psi}_\ell &= e'_\ell \tilde{\psi} \\ &= e'_\ell \text{var}(\tilde{Z})^{-1} \text{cov}(\tilde{Z}, Y) \\ &= e'_\ell \text{var}(AZ + b)^{-1} \text{cov}(AZ + b, Y) \\ &= e'_\ell \text{var}(AZ)^{-1} \text{cov}(AZ, Y) \\ &= e'_\ell [A \text{var}(Z) A']^{-1} A \text{cov}(Z, Y) \\ &= e'_\ell (A')^{-1} \text{var}(Z)^{-1} \text{cov}(Z, Y) \\ &= e'_\ell (A')^{-1} \psi. \end{aligned}$$

Similarly,

$$\tilde{\pi}_\ell = e'_\ell (A')^{-1} \pi.$$

Thus

$$\begin{aligned} \frac{\tilde{\psi}_\ell}{\tilde{\pi}_\ell} &= \frac{e'_\ell (A')^{-1} \psi}{e'_\ell (A')^{-1} \pi} \\ &= \frac{a_\ell e'_\ell \psi}{a_\ell e'_\ell \pi} \\ &= \frac{e'_\ell \psi}{e'_\ell \pi} \\ &= \frac{\psi_\ell}{\pi_\ell}. \end{aligned}$$

The second line follows since $A = \text{diag}(a_1, \dots, a_L)$, so that $e'_\ell (A')^{-1} = a_\ell e'_\ell$. \square

Proposition 10 shows that the just-identified 2SLS estimand using \tilde{Z}_ℓ as an instrument and $\tilde{Z}_{-\ell}$ as controls is equivalent to the just-identified 2SLS estimand using Z_ℓ as an instrument and $Z_{-\ell}$ as controls. As stated, this implies that the falsification adaptive set is the same too. When A is invertible and diagonal, $\tilde{Z}_\ell = a_\ell Z + b_\ell$ for some $a_\ell \neq 0$. So the proposition shows that the falsification adaptive set is invariant to rescaling each instrument and changing their location.

When A is not diagonal, however, these just-identified 2SLS estimands ψ_ℓ/π_ℓ are not invariant to affine transformations. As we show next, this is not surprising since using a non-diagonal affine transformation is equivalent to using the original instruments, but with a different way of relaxing the exclusion restriction.

To see this, consider the case where A is a rotation matrix and $b = 0$. From the outcome

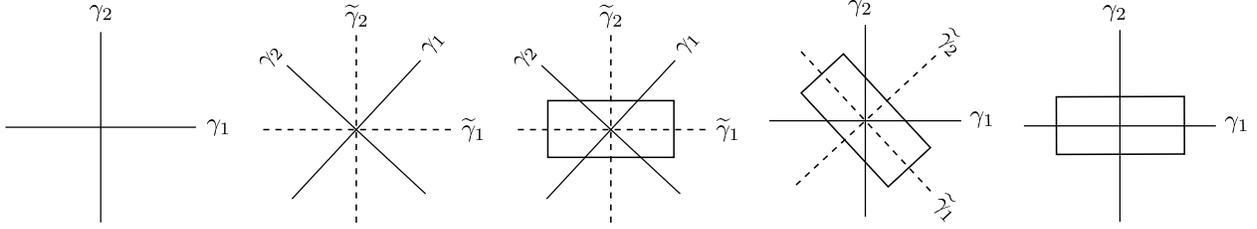


Figure 8: Rotating the instruments leads to a different class of exclusion restriction relaxations. See text for explanation.

equation we have

$$\begin{aligned}
 Y &= \beta' X + \gamma' Z + U \\
 &= \beta' X + \gamma' A^{-1} A Z + U \\
 &= \beta' X + (A^{-1} \gamma)' \tilde{Z} + U \\
 &= \beta' X + \tilde{\gamma}' \tilde{Z} + U.
 \end{aligned}$$

In the last line we defined $\tilde{\gamma} = A^{-1} \gamma$. A^{-1} is also a rotation matrix, and hence $\tilde{\gamma}$ is a rotation of γ . To illustrate this, consider the two instrument case and suppose A represents a 45 degree counter-clockwise rotation:

$$A = \begin{pmatrix} \cos(\pi/4) & -\sin(\pi/4) \\ \sin(\pi/4) & \cos(\pi/4) \end{pmatrix}.$$

Then

$$A^{-1} = A' = \begin{pmatrix} \cos(-\pi/4) & -\sin(-\pi/4) \\ \sin(-\pi/4) & \cos(-\pi/4) \end{pmatrix}$$

represents a 45 degree clockwise rotation. The first plot in figure 8 shows the original axes for measuring (γ_1, γ_2) . The second plot shows the rotated axes. Suppose we now apply our results directly using \tilde{Z} as the instruments. That is, we define relaxations of exclusion $\tilde{\gamma} = 0$ by

$$|\tilde{\gamma}_\ell| \leq \delta_\ell.$$

Then we measure relaxations of $\tilde{\gamma} = 0$ by placing rectangles around the origin, in the rotated coordinate system. The third plot in figure 8 shows an example of this constraint set, for a given value of δ . The fourth plot in figure 8 shows this same constraint set, but now in our original coordinate system. Thus we see that relaxations of the exclusion restriction $\gamma = 0$ are now measured in terms of rotated rectangles around the origin. In contrast, suppose we instead used our original instruments Z and measured relaxations of exclusion by $|\gamma_\ell| \leq \delta_\ell$. The fifth and final plot in figure 8 shows this constraint set, for the same δ .

Comparing the last two plots in figure 8 thus shows why rotating the instruments will generally yield different a falsification adaptive sets: It is equivalent to using a different class of relaxations of the exclusion restriction with the original instruments. Different classes of relaxations of one's assumptions will typically lead to different falsification adaptive sets. This is exactly the same situation as one's baseline analysis, where different baseline assumptions typically lead to different baseline estimands.

F Including Covariates in the Linear Model and Subgroup Analysis

In this section we discuss how to include additional covariates in the analysis of section 3. We consider two approaches: In the first approach we include covariates linearly. In the second approach we condition nonparametrically on the covariates. Finally, we discuss how to use our results in combination with the subgroup analysis approach discussed in section B.

Linear Controls

Generalize equation (1) as follows:

$$Y = X'\beta + Z'\gamma + W'\alpha + U \quad (1')$$

where $W \in \mathbb{R}^{\dim(W)}$ is a $\dim(W)$ -vector of additional covariates and α is a $\dim(W)$ -vector of unknown constant coefficients. In addition to assumptions A1–A3, we make two further assumptions:

1. The controls are exogenous: $\text{cov}(W, U) = 0$.
2. The vector $\tilde{Z} = (Z, W)$ has a nonsingular variance matrix.

In the framework considered in section 3, we can think of the vector W as a vector of instruments whose exclusion violation is not restricted at all. That is, we set $\delta_\ell = \infty$ for the elements of \tilde{Z} corresponding to the W covariates. Hence we do not place any constraints on the magnitude of α . It is nonetheless helpful to keep Z and W separate in the notation.

Under equation (1') and the exogenous controls assumption, we can obtain the identified set for β by simply applying the analysis of section 3, but with an initial step where the exogenous covariates are partialled out from instruments Z and endogenous regressors X . To see this, write

$$\begin{aligned} \text{cov}(Z, Y) &= \text{cov}(Z, X)\beta + \text{var}(Z)\gamma + \text{cov}(Z, W)\alpha \\ \text{cov}(W, Y) &= \text{cov}(W, X)\beta + \text{cov}(W, Z)\gamma + \text{var}(W)\alpha. \end{aligned}$$

Solve the second equation for α to get

$$\alpha = \text{var}(W)^{-1} \text{cov}(W, Y) - \text{var}(W)^{-1} \text{cov}(W, X)\beta - \text{var}(W)^{-1} \text{cov}(W, Z)\gamma.$$

Substitute this into the first equation to get

$$\begin{aligned} \text{cov}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W, Y) &= \text{cov}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W, X)\beta \\ &\quad + \text{var}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W)\gamma. \end{aligned}$$

By the assumption that \tilde{Z} has a nonsingular variance matrix, solve for γ to get

$$\begin{aligned} \gamma &= \text{var}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W)^{-1} \left(\text{cov}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W, Y) \right. \\ &\quad \left. - \text{cov}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W, X)\beta \right). \end{aligned}$$

Thus we can write the identified set for β as

$$\mathcal{B}(\delta) = \{b \in \mathbb{R}^K : -\delta \leq \tilde{\Psi} - \tilde{\Pi}\beta \leq \delta\}$$

where

$$\tilde{\Psi} = \text{var}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W)^{-1} \text{cov}(Z - \text{cov}(Z, W) \text{var}(W)^{-1} W, Y)$$

is the vector of population OLS coefficients on Z in a regression of Y on $(1, Z, W)$ and

$$\tilde{\Pi} = \text{var}(Z - \text{cov}(Z, W) \text{var}(W)^{-1}W)^{-1} \text{cov}(Z - \text{cov}(Z, W) \text{var}(W)^{-1}W, X)$$

is the vector of population OLS coefficients on Z in a regression of X on $(1, Z, W)$.

From this characterization of the identified set for β , we can adapt the proof of theorem 2 to see that the falsification adaptive set is

$$\bigcup_{\delta \in \text{FF}} \mathcal{B}(\delta) = \left[\min_{\ell=1, \dots, L} \frac{\tilde{\psi}_\ell}{\tilde{\pi}_\ell}, \max_{\ell=1, \dots, L} \frac{\tilde{\psi}_\ell}{\tilde{\pi}_\ell} \right]$$

where $\tilde{\psi}_\ell$ is the ℓ th component of $\tilde{\Psi}$ and $\tilde{\pi}_\ell$ is the ℓ th component of $\tilde{\Pi}$. Here we assume $\tilde{\pi}_\ell \neq 0$ for all $\ell \in \{1, \dots, L\}$. Moreover, as in lemma 1, $\tilde{\psi}_\ell/\tilde{\pi}_\ell$ equals the population 2SLS coefficient on X using Z_ℓ as the excluded instrument and $(Z_{-\ell}, W)$ as controls. Thus, under the above assumptions, the presence of these additional controls poses no additional technical challenges for estimation and inference since the falsification adaptive set can still be computed via standard linear 2SLS regressions.

Nonparametric Controls

Above we included controls linearly and assumed they were exogenous. An alternative approach is to condition on them nonparametrically. This allows the value of the covariates to affect the coefficients on X and Z , and also allows the controls to be endogenous. Thus we now generalize equation (1) as follows:

$$Y = X'\beta(W) + Z'\gamma(W) + U \quad (1'')$$

where W is a $\text{dim}(W)$ -vector of controls, and $\beta(\cdot)$ and $\gamma(\cdot)$ are unknown functions. We replace A1–A4 with their conditional-on- W versions:

A1*. $\text{cov}(Z, X \mid W = w)$ has rank K for almost all $w \in \text{supp}(W)$.

A2*. $\text{var}(Z \mid W = w)$ has full rank for almost all $w \in \text{supp}(W)$.

A3*. For all $\ell \in \{1, \dots, L\}$, $\text{cov}(Z_\ell, U \mid W = w) = 0$ for almost all $w \in \text{supp}(W)$.

A4*. For all $\ell \in \{1, \dots, L\}$, $\gamma_\ell(w) = 0$ for almost all $w \in \text{supp}(W)$.

We relax A4* by

A4**. There are known functions $\delta_\ell(w) \geq 0$ such that $|\gamma_\ell(w)| \leq \delta_\ell(w)$ for almost all $w \in \text{supp}(W)$, for all $\ell \in \{1, \dots, L\}$. Let $\delta(w) = (\delta_1(w), \dots, \delta_L(w))'$.

Write

$$\gamma(w) = \text{var}(Z \mid W = w)^{-1} (\text{cov}(Z, Y \mid W = w) - \text{cov}(Z, X \mid W = w)\beta(w))$$

for any $w \in \text{supp}(W)$.

By derivations similar to those in section 3, the identified set for $\beta(w)$ is

$$\begin{aligned} & \mathcal{B}(\delta(w) \mid w) \\ &= \{b \in \mathbb{R}^K : -\delta(w) \leq \text{var}(Z \mid W = w)^{-1} (\text{cov}(Z, Y \mid W = w) - \text{cov}(Z, X \mid W = w)b) \leq \delta(w)\} \end{aligned}$$

for any $w \in \text{supp}(W)$. For a given $\delta(\cdot)$, the model is refuted if the set $\{w \in \text{supp}(W) : \mathcal{B}(\delta(w) | w) = \emptyset\}$ has positive probability.

Subgroup Analysis

Suppose that the baseline model discussed in section 3 is refuted. Or suppose the model with covariates based on equation (1') is refuted. As discussed in section B, one response is to switch to the nonparametric controls model based on equation (1''), and then find the set of conditioning covariates for which the model is not refuted. Specifically, define

$$\mathcal{W}(0) = \{w \in \text{supp}(W) : \mathcal{B}(0 | w) \neq \emptyset\}.$$

This set is point identified. If we impose A1*-A4* only for $w \in \mathcal{W}(0)$, then the model is not refuted. Thus one approach to salvaging the model is to focus on the subgroups $\mathcal{W}(0)$, and to report $\beta(w)$ for $w \in \mathcal{W}(0)$.

Alternatively, one could make specific assumptions to extrapolate from $\beta(w)$ for $w \in \mathcal{W}(0)$ to $\beta(w)$ for w in the complement $\mathcal{W}(0)^c$. For example, suppose $\beta(w)$ is observed to be constant across $w \in \mathcal{W}(0)$. Then one could assume that it is also constant across $w \in \mathcal{W}(0)^c$, and that this constant value is the same as the observed constant value across $w \in \mathcal{W}(0)$.

Suppose, however, that we want to allow for $\beta(w)$ to vary over w , and that we are interested in the entire population, not just the subgroups in $\mathcal{W}(0)$. In this case, a natural parameter is the mean coefficient, $\mathbb{E}[\beta(W)]$. We can construct the falsification adaptive set for this parameter as follows. For simplicity, suppose $K = 1$. We have

$$\begin{aligned} \mathbb{E}[\beta(W)] &= \int_{w \in \text{supp}(W)} \beta(w) dF_W(w) \\ &= \int_{w \in \mathcal{W}(0)} \beta(w) dF_W(w) + \int_{w \notin \mathcal{W}(0)} \beta(w) dF_W(w). \end{aligned}$$

The first term is point identified. For the second term, we can construct the falsification adaptive set for $\beta(w)$ for each $w \in \mathcal{W}(0)$ by applying the results of section 3. Denote these bounds by $[L(w), U(w)]$. Then the falsification adaptive set for $\mathbb{E}[\beta(W)]$ is $[L, U]$ where

$$L = \int_{w \in \mathcal{W}(0)} \beta(w) dF_W(w) + \int_{w \notin \mathcal{W}(0)} L(w) dF_W(w)$$

and

$$U = \int_{w \in \mathcal{W}(0)} \beta(w) dF_W(w) + \int_{w \notin \mathcal{W}(0)} U(w) dF_W(w).$$

Keep in mind that this falsification adaptive set relaxes exclusion by using values $\delta_\ell(w)$ that vary across w . In particular, $\delta_\ell(w) = 0$ for all $w \in \mathcal{W}(0)$. If we require $\delta_\ell(w) = \delta_\ell$ to be constant across w , then the falsification adaptive set will be weakly larger.

G Technical Equivalence of Exclusion and Exogeneity in Linear Models

In section 3 we explicitly modeled relaxations of the instrument exclusion restriction, holding instrument exogeneity fixed. In this section we show that the same formal results can be used to relax exogeneity alone or to relax exogeneity and exclusion simultaneously. This analysis is essentially a variation of Small (2007).

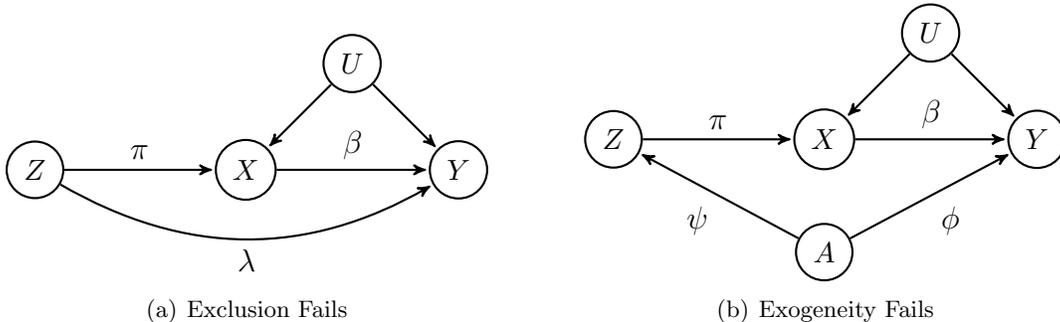


Figure 9: Two ways for the standard instrumental variable assumptions to fail.

Let A and U denote unobserved random variables. Suppose the true model for potential outcomes is

$$Y(x, z, a) = x'\beta + z'\lambda + \phi a + U.$$

Thus observed outcomes are

$$\begin{aligned} Y &= Y(X, Z, A) \\ &= X'\beta + Z'\lambda + (\phi A + U). \end{aligned}$$

This equation allows for violations of exclusion if λ is nonzero. It also allows for violations of exogeneity if A and Z are correlated. Figure G shows these two kinds of failures. Let ψ be the coefficients on Z in a linear projection of A onto Z and let \tilde{U} be the projection residual. Thus

$$A = Z'\psi + \tilde{U}$$

where $\text{cov}(Z_\ell, \tilde{U}) = 0$ by construction. Then

$$\begin{aligned} Y &= X'\beta + Z'\lambda + \phi A + U \\ &= X'\beta + Z'\lambda + \phi(Z'\psi + \tilde{U}) + U \\ &= X'\beta + Z'(\lambda + \phi\psi) + (\tilde{U} + U) \\ &\equiv X'\beta + Z'\gamma + V. \end{aligned}$$

In the last line we defined

$$\gamma = \lambda + \phi\psi$$

and $V = \tilde{U} + U$. For all $\ell \in \{1, \dots, L\}$, we have $\text{cov}(Z_\ell, V) = 0$ since $\text{cov}(Z_\ell, U) = 0$ by A3 and $\text{cov}(Z_\ell, \tilde{U}) = 0$ by construction.

Thus the formal results in section 3 apply to relaxations of exogeneity as well. The key difference is that the *interpretation* of the coefficient on Z changes depending on whether we are interested in exclusion or exogeneity failure. For example, suppose $\lambda = 0$ so that we are only worried about exogeneity failures. Then the interpretation of γ depends on the size of the causal effect of ϕ and the relationship between Z and A . This is the case that Small (2007) focuses on. To aid interpretation of these parameters, he adapts the ideas of Imbens (2003) to phrase constraints on ϕ and $\|\psi\|$ in terms of R^2 measures.

H Comparing the 2SLS Estimand with the Falsification Adaptive Set

The following proposition gives the relationship between the 2SLS estimand

$$\beta_{2\text{SLS}} = \frac{\text{cov}(X, Z) \text{var}(Z)^{-1} \text{cov}(Z, Y)}{\text{cov}(X, Z) \text{var}(Z)^{-1} \text{cov}(Z, X)}$$

and the falsification adaptive set when there is a single endogenous variable. Here we assume all instruments are relevant.

Proposition 11. Suppose $K = 1$. Suppose $\pi_\ell \neq 0$ for all $\ell = 1, \dots, L$. Then

$$\beta_{2\text{SLS}} = \sum_{\ell=1}^L W_\ell \frac{\psi_\ell}{\pi_\ell}$$

where

$$W_\ell = \frac{\text{cov}(\sum_{j=1}^L \pi_j Z_j, \pi_\ell Z_\ell)}{\text{var}(\sum_{j=1}^L \pi_j Z_j)}.$$

Recall that ψ_ℓ/π_ℓ is the 2SLS estimand using Z_ℓ as an instrument and the remaining instruments $Z_{-\ell}$ as controls. The falsification adaptive set is the interval between the smallest and largest such estimands. Proposition 11 shows that the 2SLS estimand is a weighted average of the estimands ψ_ℓ/π_ℓ . The weights sum to one, $\sum_{\ell=1}^L W_\ell = 1$, but the weights are not necessarily positive. For example, this may happen if covariances between instruments are negative, or if some $\pi_\ell < 0$.

Thus an important implication of this result is that the 2SLS estimand does not have to be inside the falsification adaptive set. Indeed, it can be arbitrarily far from it. Put differently, in a falsified linear IV model, the 2SLS estimand can be arbitrarily far from the set of parameters which correspond to minimally non-falsified models.

Proof of proposition 11. Recall that

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

Thus

$$\beta_{2\text{SLS}} = \frac{\text{cov}(X, Z) \psi}{\text{cov}(X, Z) \pi}.$$

Since π is the vector of coefficients from a linear projection of X onto the instruments, we can write

$$X = \sum_{\ell=1}^L \pi_\ell Z_\ell + V$$

where $\text{cov}(Z_\ell, V) = 0$ for all ℓ . Consider the numerator of $\beta_{2\text{SLS}}$:

$$\begin{aligned}
\text{cov}(X, Z)\psi &= \sum_{\ell=1}^L \text{cov}(X, Z_\ell)\psi_\ell \\
&= \sum_{\ell=1}^L \text{cov}\left(\sum_{j=1}^L \pi_j Z_j + V, Z_\ell\right) \psi_\ell \\
&= \sum_{\ell=1}^L \text{cov}\left(\sum_{j=1}^L \pi_j Z_j, Z_\ell\right) \psi_\ell \\
&= \sum_{\ell=1}^L \frac{\psi_\ell}{\pi_\ell} \text{cov}\left(\sum_{j=1}^L \pi_j Z_j, \pi_\ell Z_\ell\right).
\end{aligned}$$

Next consider the denominator of $\beta_{2\text{SLS}}$:

$$\begin{aligned}
\text{cov}(X, Z)\pi &= \sum_{\ell=1}^L \text{cov}(X, Z_\ell)\pi_\ell \\
&= \sum_{\ell=1}^L \text{cov}\left(\sum_{s=1}^L \pi_s Z_s + V, Z_\ell\right) \pi_\ell \\
&= \sum_{\ell=1}^L \sum_{s=1}^L \text{cov}(\pi_s Z_s, \pi_\ell Z_\ell) \\
&= \text{var}\left(\sum_{j=1}^L \pi_j Z_j\right).
\end{aligned}$$

Putting these two results together yields the result. □

I Further Remarks on Interpreting Non-Falsified Models

In this section we elaborate on the discussion in section 2.2. We only consider the linear instrumental variable model of section 3 with one endogenous variable and two instruments. Newey (1985) analyzes the problem in a much broader class of moment equality models.

With just two instruments, the outcome equation is

$$Y = \alpha + \beta X + \gamma_1 Z_1 + \gamma_2 Z_2 + U. \tag{20}$$

Here we include a constant intercept α , unlike in section 3. So we also impose the normalization $\mathbb{E}(U) = 0$. This normalization combined with the exogeneity assumption A3 gives us the following three moment conditions:

$$\begin{aligned}
\mathbb{E}[U] &= 0 \\
\mathbb{E}[Z_1 U] &= 0 \\
\mathbb{E}[Z_2 U] &= 0.
\end{aligned}$$

We can use these moments to define a GMM objective function

$$Q(b) = m(b)'Wm(b)$$

where W is a symmetric positive definite weighting matrix, $m(b) = (m_1(b), m_2(b))'$, and

$$m_\ell(b) = \text{cov}(Z_\ell, Y) - b \text{cov}(Z_\ell, X).$$

for $\ell \in \{1, 2\}$. This leads to the following standard result.

Proposition 12. Consider the outcome equation (20). Suppose $\mathbb{E}(U) = 0$. Suppose A1, A3, and A4 hold; that is, suppose the baseline model is true. Then Q is uniquely minimized at β . Moreover, $Q(\beta) = 0$.

This result immediately provides a way of checking whether the model is falsified. Let

$$Q^* = \min_{b \in \mathbb{R}} Q(b).$$

Corollary 7. Suppose $Q^* > 0$. Then the model is false.

As we discussed in section 2.2, however, the inverse statement is *not* true. That is, observing $Q^* = 0$ does not imply the model is true. It merely implies that there exist values of the unobservables that are consistent with the assumptions and the data. But, as is well known, there *also* exist values of the unobservables that are consistent with the data but which are *not* consistent with the assumptions. To see this, we will study the value of Q^* as the true coefficients (γ_1, γ_2) on the instruments vary. For simplicity, suppose we use the identity weight matrix. Then

$$Q(b) = (\text{cov}(Z_1, Y) - b \text{cov}(Z_1, X))^2 + (\text{cov}(Z_2, Y) - b \text{cov}(Z_2, X))^2. \quad (21)$$

Proposition 13. Suppose the true model is given by (20) with coefficients $(\beta, \gamma_1, \gamma_2)$. That is, the exclusion restriction does not necessarily hold. Suppose A1 and A3 hold. Then the value of Q evaluated at its optimizer is

$$Q^*(\gamma_1, \gamma_2) = f(\gamma_1, \gamma_2)^2 \left[(1 - \omega)^2 + \omega^2 \left(\frac{\text{cov}(Z_2, X)}{\text{cov}(Z_1, X)} \right)^2 \right]$$

where

$$f(\gamma_1, \gamma_2) = \gamma_1 \left[\text{var}(Z_1) - \text{cov}(Z_1, Z_2) \frac{\text{cov}(Z_1, X)}{\text{cov}(Z_2, X)} \right] + \gamma_2 \left[\text{cov}(Z_1, Z_2) - \text{var}(Z_2) \frac{\text{cov}(Z_1, X)}{\text{cov}(Z_2, X)} \right]$$

and

$$\omega = \frac{\text{cov}(X, Z_1)^2}{\text{cov}(X, Z_1)^2 + \text{cov}(X, Z_2)^2}.$$

Notice that $Q^*(\gamma_1, \gamma_2)$ does not depend on the value of β . The left plot in figure 10 shows an example of $Q^*(\gamma_1, \gamma_2)$.¹⁰ From this proposition, we see that $Q^*(\gamma_1, \gamma_2) = 0$ if and only if $f(\gamma_1, \gamma_2) = 0$. This equation defines a line in \mathbb{R}^2 . Thus the model is not refuted if and only if (γ_1, γ_2) falls on this line. Only one of these points—the origin—corresponds to the baseline

¹⁰This figure plots the function $Q_{\text{MT}}^*(\delta, \gamma_1, \gamma_2)$ which we discuss below. When evaluated at $\delta = 0$, and when $\text{var}(Z_1) = \text{var}(Z_2) = 1$ as used in the figure, we have $Q_{\text{MT}}^*(0, \gamma_1, \gamma_2) = Q^*(\gamma_1, \gamma_2)$. In general, $Q_{\text{MT}}^*(0, \gamma_1, \gamma_2)$ equals the GMM objective function with weight matrix $W = \text{diag}(1/\text{var}(Z_1)^2, 1/\text{var}(Z_2)^2)$.

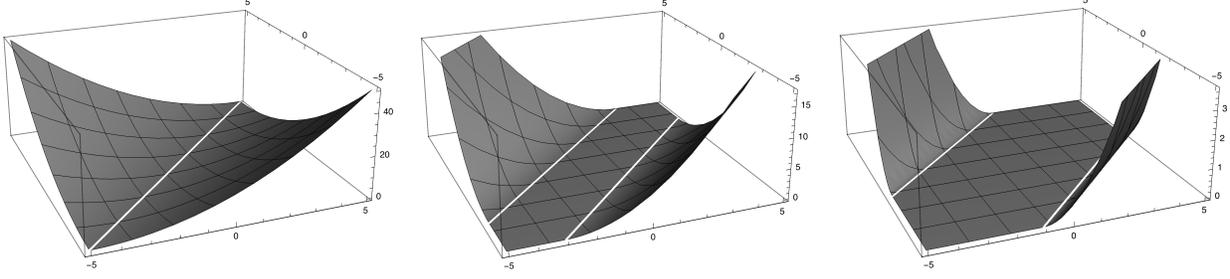


Figure 10: Plots of $Q_{\text{MT}}^*(\delta, \gamma_1, \gamma_2)$ as functions of (γ_1, γ_2) for three different values of δ . Left: The baseline case, $\delta_1 = \delta_2 = 0$. Middle: A relaxed model with $\delta_1 = \delta_2 = 1$. Right: A relaxed model with $\delta_1 = \delta_2 = 2$. All three plots set $\text{var}(Z_1) = \text{var}(Z_2) = 1$, $\text{cov}(Z_1, Z_2) = 0$, and $\text{cov}(Z_1, X) = \text{cov}(Z_2, X) = 0.5$.

model. All other points on this line are false models which are observationally equivalent to a set of structural parameters which satisfy the model assumptions.

Next we consider the relaxed model which replaces exclusion A4 with partial exclusion A4'. Following Manski and Tamer (2002) and Chernozhukov, Hong, and Tamer (2007), we can again characterize falsification of the model in terms of a certain objective function. Under partial exclusion, theorem 1 shows that the identified set $\mathcal{B}(\delta)$ is characterized by four moment inequalities:

$$\begin{aligned} (\psi_\ell - \pi_\ell \beta) + \delta_\ell &\geq 0 \\ -(\psi_\ell - \pi_\ell \beta) + \delta_\ell &\geq 0 \end{aligned}$$

for $\ell \in \{1, 2\}$. Recall that ψ and π were defined on page 13. Define

$$Q_{\text{MT}}(b, \delta) = \sum_{\ell=1}^2 \left(\max \{ -[(\psi_\ell - \pi_\ell b) + \delta_\ell], 0 \}^2 + \max \{ -[-(\psi_\ell - \pi_\ell b) + \delta_\ell], 0 \}^2 \right).$$

Then $Q_{\text{MT}}(b, \delta) = 0$ if and only if $b \in \mathcal{B}(\delta)$. This objective function generalizes equation (21). Recall the notation $B_\ell(\delta) = [\underline{b}_\ell(\delta), \bar{b}_\ell(\delta)]$ from equation (4). The following proposition characterizes this objective function at its minimal value.

Proposition 14. Let $Q_{\text{MT}}^*(\delta) = \min_{b \in \mathbb{R}} Q_{\text{MT}}(b, \delta)$. Then

$$Q_{\text{MT}}^*(\delta) = \begin{cases} 0 & \text{if } \max\{\underline{b}_1(\delta), \underline{b}_2(\delta)\} \leq \min\{\bar{b}_1(\delta), \bar{b}_2(\delta)\} \\ \frac{1}{2}(\underline{b}_2(\delta) - \bar{b}_1(\delta))^2 & \text{if } \bar{b}_1(\delta) < \underline{b}_2(\delta) \\ \frac{1}{2}(\underline{b}_1(\delta) - \bar{b}_2(\delta))^2 & \text{if } \bar{b}_2(\delta) < \underline{b}_1(\delta). \end{cases}$$

The identified set $\mathcal{B}(\delta)$ is the intersection of two intervals, $B_1(\delta)$ and $B_2(\delta)$. When these two intervals overlap, the model is not refuted and $Q_{\text{MT}}^*(\delta) = 0$. This is the first case. When the intervals do not overlap, $Q_{\text{MT}}^*(\delta)$ is strictly positive. The second case above occurs when $B_1(\delta)$ is completely below $B_2(\delta)$. The third case above is the reverse. That is, if $Q_{\text{MT}}^*(\delta) > 0$, then the model is false.

As in proposition 13 above, we can view $Q_{\text{MT}}^*(\delta)$ as a function of the true underlying parameters, which we write as

$$Q_{\text{MT}}^*(\delta, \gamma_1, \gamma_2).$$

We give an explicit expression for this function on page 91. The middle and right plots in figure 10 show two examples of $Q_{\text{MT}}^*(\delta, \gamma_1, \gamma_2)$, for two different values of δ . Like the baseline case, $(\delta_1, \delta_2) = (0, 0)$, there is a set of values (γ_1, γ_2) where the objective function is zero. This set contains the values in the square $[-\delta_1, \delta_1] \times [-\delta_2, \delta_2]$ which are consistent with partial exclusion A4'. But this set also contains values outside of this square. Each of these values is a false model which is observationally equivalent to a set of structural parameters which satisfy the model assumptions. As the components of δ increase, this set of (γ_1, γ_2) values such that $Q_{\text{MT}}^*(\delta, \gamma_1, \gamma_2) = 0$ grows.

Finally, we relate these results to the falsification frontier. In figure 10, as we vary (γ_1, γ_2) in the domain, the distribution of the observed data changes. This follows since $\text{cov}(Z_1, Y)$ and $\text{cov}(Z_2, Y)$ change when (γ_1, γ_2) change while holding all else fixed. In our analysis in sections 2 and 3 we fix a single distribution of the observed data (Y, X, Z_1, Z_2) . We then check whether the model is falsified by that distribution. For example, we check whether the objective function in equation (21) is zero at its optimizer. Suppose it is not. Then the model is falsified. This means that, in the left plot of figure 10, we observed a distribution of (Y, X, Z_1, Z_2) off the diagonal defined by $f(\gamma_1, \gamma_2) = 0$. As we increase the components of δ , the value function $Q_{\text{MT}}^*(\delta)$ will flatten, until eventually we have $Q_{\text{MT}}^*(\delta) = 0$. The smallest set of δ 's at which this happens is what we call the falsification frontier.

The Identified Set for (β, γ)

An alternative way to view these arguments is to study the identified set for (β, γ) when no constraints on the values of γ are imposed. Specifically, we have the following result.

Proposition 15. Suppose A1, A2, and A3 hold. Suppose $K = 1$. Then

$$\Gamma = \{\gamma \in \mathbb{R}^L : \gamma = \psi - b\Pi \text{ for some } b \in \mathbb{R}\}$$

is the identified set for γ . Moreover, the identified set for (β, γ) is

$$\{(\beta, \gamma) \in \mathbb{R}^{1+L} : \psi - \beta\Pi = \gamma \text{ for some } \gamma \in \Gamma\}.$$

Recall from page 13 that ψ is the vector of reduced form regression coefficients of Y on Z . Similarly, Π is the vector of first stage regression coefficients of X on Z . Proposition 15 is essentially a restatement of equation (8) in Small (2007).

The identified set Γ is never empty, which means that without any constraints on the direct causal effects of the instruments, the model is not refutable. With a single instrument ($L = 1$), the identified set Γ always contains 0. That is, the exclusion restriction $\gamma = 0$ is always consistent with the data and assumptions A1, A2, and A3. This is simply the well known result that the classical linear IV model with one instrument is not refutable.

With $L \geq 2$ instruments, however, the identified set Γ does not necessarily contain the origin. Thus the baseline model (the outcome equation (1) plus assumptions A1–A4) with two or more instruments is refutable: The baseline model is refuted if and only if the origin is not an element of the identified set Γ . Figure 11 shows two examples of the identified sets Γ , for the two instrument case. On the left, the baseline model is not refuted since the set Γ passes through the origin. On the right, the baseline model is refuted since the set Γ does not pass through the origin.

Importantly, even when the baseline model is not refuted, the identified set for γ always contains elements which are inconsistent with the exclusion restriction. This is the same point we discussed above, by analyzing the GMM objective function.

To relate this discussion to the falsification frontier, fix a distribution of the observed data

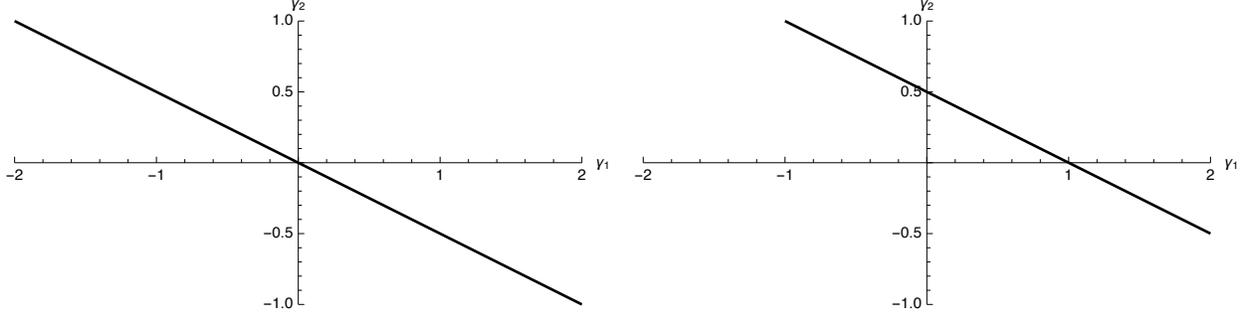


Figure 11: Example identified sets for (γ_1, γ_2) . Left: A distribution of (Y, X, Z_1, Z_2) which does not refute the baseline model. Right: A distribution of (Y, X, Z_1, Z_2) which does refute the baseline model.

(Y, X, Z) . We then check whether the baseline model is falsified by that distribution. For example, we check whether Γ contains the origin. Suppose it does not. Then the baseline model is falsified. We then compute the falsification frontier. While this set was given by proposition 2, it can alternatively be obtained as the set

$$\text{FF} = \left\{ \delta \in [0, \infty)^L : \left(\prod_{\ell=1}^L [-\delta_\ell, \delta_\ell] \right) \cap \Gamma \text{ is a singleton} \right\}.$$

Essentially, we consider boxes around the origin. We then find all boxes that have a singleton intersection with the line Γ .

At any point δ on the falsification frontier, the relaxed model (the outcome equation (1) plus assumptions A1–A3 and A4') is not refuted and there is a unique γ value consistent with the relaxed model assumptions. If we drop the partial exclusion assumption A4', however, there are infinitely many values of γ consistent with the remaining assumptions and the data—all the elements of the set Γ .

Proofs for section I

Proof of proposition 12. This result is standard, but we include the proof for completeness. Substituting the outcome equation (20) into our system of three moments and using exclusion yields the moments

$$\begin{aligned} \mathbb{E}[Y - \alpha - \beta X] &= 0 \\ \mathbb{E}[Z_1(Y - \alpha - \beta X)] &= 0 \\ \mathbb{E}[Z_2(Y - \alpha - \beta X)] &= 0. \end{aligned}$$

Next we profile out the intercept. Solving the first equation for α yields $\alpha = \mathbb{E}(Y) - \beta\mathbb{E}(X)$. Substituting this into the second two equations gives

$$\begin{aligned} \mathbb{E}[Z_1(Y - \mathbb{E}(Y) + \beta\mathbb{E}(X) - \beta X)] &= 0 \\ \mathbb{E}[Z_2(Y - \mathbb{E}(Y) + \beta\mathbb{E}(X) - \beta X)] &= 0. \end{aligned}$$

We can rewrite this system as

$$\begin{aligned}\text{cov}(Z_1, Y) - \beta \text{cov}(Z_1, X) &= 0 \\ \text{cov}(Z_2, Y) - \beta \text{cov}(Z_2, X) &= 0.\end{aligned}$$

This is a system of two moment equalities with one unknown, β . The left hand side is the vector $m(\beta)$. This explains the definition of Q and shows that $Q(\beta) = 0$. The fact that Q is uniquely minimized at β follows by the relevance assumption A1 and since the moments are linear in β . \square

Proof of corollary 7. This is simply the contrapositive of proposition 12. For concreteness, however, notice that the model is falsified if and only if both instruments are relevant, so that $\text{cov}(Z_\ell, X) \neq 0$ for $\ell \in \{1, 2\}$, and

$$\frac{\text{cov}(Z_1, Y)}{\text{cov}(Z_1, X)} \neq \frac{\text{cov}(Z_2, Y)}{\text{cov}(Z_2, X)}.$$

This happens when there does not exist a solution to the population moment equations. \square

Proof of proposition 13. Solving the first order conditions for minimizing Q yields

$$\begin{aligned}b^* &= \frac{\text{cov}(Z_1, Y) \text{cov}(Z_1, X) + \text{cov}(Z_2, Y) \text{cov}(Z_2, X)}{\text{cov}(Z_1, X)^2 + \text{cov}(Z_2, X)^2} \\ &= \frac{\text{cov}(Y, Z_1)}{\text{cov}(X, Z_1)} \frac{\text{cov}(X, Z_1)^2}{\text{cov}(X, Z_1)^2 + \text{cov}(X, Z_2)^2} + \frac{\text{cov}(Y, Z_2)}{\text{cov}(X, Z_2)} \frac{\text{cov}(X, Z_2)^2}{\text{cov}(X, Z_1)^2 + \text{cov}(X, Z_2)^2} \\ &\equiv \beta_{Z_1}^{2\text{SLS}} \omega + \beta_{Z_2}^{2\text{SLS}} (1 - \omega)\end{aligned}$$

where $\omega \in [0, 1]$. That is, the optimizer is a convex combination of the two individual 2SLS estimands. Thus the value function is

$$\begin{aligned}Q^* &= Q(b^*) \\ &= (\text{cov}(Z_1, Y) - b^* \text{cov}(Z_1, X))^2 + (\text{cov}(Z_2, Y) - b^* \text{cov}(Z_2, X))^2 \\ &= (\text{cov}(Z_1, Y) - [\beta_{Z_1}^{2\text{SLS}} \omega + \beta_{Z_2}^{2\text{SLS}} (1 - \omega)] \text{cov}(Z_1, X))^2 \\ &\quad + (\text{cov}(Z_2, Y) - [\beta_{Z_1}^{2\text{SLS}} \omega + \beta_{Z_2}^{2\text{SLS}} (1 - \omega)] \text{cov}(Z_2, X))^2 \\ &= (\text{cov}(Z_1, Y) - \text{cov}(Z_1, Y) \omega - (1 - \omega) \beta_{Z_2}^{2\text{SLS}} \text{cov}(Z_1, X))^2 \\ &\quad + (\text{cov}(Z_2, Y) - (1 - \omega) \text{cov}(Z_2, Y) - \omega \beta_{Z_1}^{2\text{SLS}} \text{cov}(Z_2, X))^2 \\ &= (1 - \omega)^2 (\text{cov}(Z_1, Y) - \text{cov}(Z_1, X) \beta_{Z_2}^{2\text{SLS}})^2 + \omega^2 (\text{cov}(Z_2, Y) - \text{cov}(Z_2, X) \beta_{Z_1}^{2\text{SLS}})^2 \\ &= (1 - \omega)^2 \left(\text{cov}(Z_1, Y) - \text{cov}(Z_2, Y) \frac{\text{cov}(Z_1, X)}{\text{cov}(Z_2, X)} \right)^2 \\ &\quad + \omega^2 \left(\text{cov}(Z_2, Y) - \text{cov}(Z_1, Y) \frac{\text{cov}(Z_2, X)}{\text{cov}(Z_1, X)} \right)^2.\end{aligned}$$

To understand how this value function changes as the underlying dgp changes, we substitute the outcome equation (20) for Y . Since Y only enters via its covariance with Z_ℓ , this yields

$$\begin{aligned}\text{cov}(Z_\ell, Y) &= \text{cov}(Z_\ell, \alpha + \beta X + \gamma_1 Z_1 + \gamma_2 Z_2 + U) \\ &= \beta \text{cov}(Z_\ell, X) + \sum_{j \in \{1, 2\}} \gamma_j \text{cov}(Z_\ell, Z_j).\end{aligned}$$

Thus

$$\begin{aligned}
& Q^*(\gamma_1, \gamma_2) \\
&= (1 - \omega)^2 \left(\beta \operatorname{cov}(Z_1, X) + \gamma_1 \operatorname{var}(Z_1) + \gamma_2 \operatorname{cov}(Z_1, Z_2) \right. \\
&\quad \left. - [\beta \operatorname{cov}(Z_2, X) + \gamma_1 \operatorname{cov}(Z_1, Z_2) + \gamma_2 \operatorname{var}(Z_2)] \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right)^2 \\
&\quad + \omega^2 \left(\beta \operatorname{cov}(Z_2, X) + \gamma_1 \operatorname{cov}(Z_1, Z_2) + \gamma_2 \operatorname{var}(Z_2) \right. \\
&\quad \left. - [\beta \operatorname{cov}(Z_1, X) + \gamma_1 \operatorname{var}(Z_1) + \gamma_2 \operatorname{cov}(Z_1, Z_2)] \frac{\operatorname{cov}(Z_2, X)}{\operatorname{cov}(Z_1, X)} \right)^2 \\
&= (1 - \omega)^2 \left(\gamma_1 \left[\operatorname{var}(Z_1) - \operatorname{cov}(Z_1, Z_2) \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right] + \gamma_2 \left[\operatorname{cov}(Z_1, Z_2) - \operatorname{var}(Z_2) \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right] \right)^2 \\
&\quad + \omega^2 \left(\gamma_1 \left[\operatorname{cov}(Z_1, Z_2) - \operatorname{var}(Z_1) \frac{\operatorname{cov}(Z_2, X)}{\operatorname{cov}(Z_1, X)} \right] + \gamma_2 \left[\operatorname{var}(Z_2) - \operatorname{cov}(Z_1, Z_2) \frac{\operatorname{cov}(Z_2, X)}{\operatorname{cov}(Z_1, X)} \right] \right)^2 \\
&= (1 - \omega)^2 \left(\gamma_1 \left[\operatorname{var}(Z_1) - \operatorname{cov}(Z_1, Z_2) \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right] + \gamma_2 \left[\operatorname{cov}(Z_1, Z_2) - \operatorname{var}(Z_2) \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right] \right)^2 \\
&\quad + \omega^2 \left(\frac{\operatorname{cov}(Z_2, X)}{\operatorname{cov}(Z_1, X)} \right)^2 \left(\gamma_1 \left[\operatorname{var}(Z_1) - \operatorname{cov}(Z_1, Z_2) \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right] + \gamma_2 \left[\operatorname{cov}(Z_1, Z_2) - \operatorname{var}(Z_2) \frac{\operatorname{cov}(Z_1, X)}{\operatorname{cov}(Z_2, X)} \right] \right)^2.
\end{aligned}$$

□

Proof of proposition 14. When the model is not refuted, $Q_{\text{MT}}(b, \delta) = 0$ by definition. This gives us the first case in the definition of $Q_{\text{MT}}^*(\delta)$. So suppose the model is refuted. Then the intervals $[\underline{b}_1(\delta), \bar{b}_1(\delta)]$ and $[\underline{b}_2(\delta), \bar{b}_2(\delta)]$ do not intersect. Hence one of them lies strictly to the left of the other. Suppose that

$$[L_1, U_1] \equiv [\underline{b}_1(\delta), \bar{b}_1(\delta)]$$

lies to the left of

$$[L_2, U_2] \equiv [\underline{b}_2(\delta), \bar{b}_2(\delta)],$$

meaning that $U_1 < L_2$. First we show that $Q_{\text{MT}}(b, \delta)$ is uniquely minimized at the midpoint between U_1 and L_2 ,

$$b^* = \frac{U_1 + L_2}{2}.$$

First note that any value b outside of $[L_1, U_2]$ can never be a minimizer. Furthermore, b^* must lie inside $[U_1, L_2]$. For b 's in this range, the objective function is

$$\begin{aligned}
Q_{\text{MT}}(b, \delta) &= \max\{b - U_1, 0\}^2 + \max\{-(b - L_2), 0\}^2 \\
&= (b - U_1)^2 + (b - L_2)^2.
\end{aligned}$$

The FOC is

$$2(b - U_1)(-1) + 2(b - L_2)(-1) = 0.$$

Solve for b to get the midpoint. This essentially follows since our objective function gives equal

weight to violations of each of the two inequalities. Next, to obtain the value of the objective function at the optimizer, simply plug b^* back in to get

$$\begin{aligned} Q_{\text{MT}}(b^*) &= \left(\frac{U_1 + L_2}{2} - U_1 \right)^2 + \left(\frac{U_1 + L_2}{2} - L_2 \right)^2 \\ &= \frac{(L_2 - U_1)^2}{4} + \frac{(U_1 - L_2)^2}{4} \\ &= \frac{1}{2}(U_1 - L_2)^2. \end{aligned}$$

This gives us the second case in the definition of $Q_{\text{MT}}^*(\delta)$. The third case occurs when the interval $[\underline{b}_2(\delta), \bar{b}_2(\delta)]$ lies to the left of $[\underline{b}_1(\delta), \bar{b}_1(\delta)]$. The result follows by following the derivations of the second case, but reversing the role of the two intervals. \square

Derivation of $Q_{\text{MT}}^(\delta, \gamma_1, \gamma_2)$.* First substitute the definitions of the bounds $\underline{b}_\ell(\delta)$ and $\bar{b}_\ell(\delta)$ into the expression for $Q_{\text{MT}}^*(\delta)$. This yields $Q_{\text{MT}}^*(\delta) = 0$ if

$$\max \left\{ \frac{\psi_1}{\pi_1} - \frac{\delta_1}{|\pi_1|}, \frac{\psi_2}{\pi_2} - \frac{\delta_2}{|\pi_2|} \right\} \leq \min \left\{ \frac{\psi_1}{\pi_1} + \frac{\delta_1}{|\pi_1|}, \frac{\psi_2}{\pi_2} + \frac{\delta_2}{|\pi_2|} \right\}$$

and

$$Q_{\text{MT}}^*(\delta) = \begin{cases} \frac{1}{2} \left[\left(\frac{\psi_2}{\pi_2} - \frac{\delta_2}{|\pi_2|} \right) - \left(\frac{\psi_1}{\pi_1} + \frac{\delta_1}{|\pi_1|} \right) \right]^2 & \text{if } \left(\frac{\psi_2}{\pi_2} - \frac{\delta_2}{|\pi_2|} \right) > \left(\frac{\psi_1}{\pi_1} + \frac{\delta_1}{|\pi_1|} \right) \\ \frac{1}{2} \left[\left(\frac{\psi_1}{\pi_1} - \frac{\delta_1}{|\pi_1|} \right) - \left(\frac{\psi_2}{\pi_2} + \frac{\delta_2}{|\pi_2|} \right) \right]^2 & \text{if } \left(\frac{\psi_1}{\pi_1} - \frac{\delta_1}{|\pi_1|} \right) > \left(\frac{\psi_2}{\pi_2} + \frac{\delta_2}{|\pi_2|} \right). \end{cases}$$

Of all of these terms, only ψ_ℓ depends on $(\beta, \gamma_1, \gamma_2)$. Specifically,

$$\begin{aligned} \psi &= \text{var}(Z)^{-1} \text{cov}(Z, Y) \\ &= \begin{pmatrix} \text{var}(Z_1) & \text{cov}(Z_1, Z_2) \\ \text{cov}(Z_1, Z_2) & \text{var}(Z_2) \end{pmatrix}^{-1} \begin{pmatrix} \text{cov}(Z_1, Y) \\ \text{cov}(Z_2, Y) \end{pmatrix} \\ &= \frac{1}{\text{var}(Z_1) \text{var}(Z_2) - \text{cov}(Z_1, Z_2)^2} \begin{pmatrix} \text{var}(Z_2) & -\text{cov}(Z_1, Z_2) \\ -\text{cov}(Z_1, Z_2) & \text{var}(Z_1) \end{pmatrix} \begin{pmatrix} \text{cov}(Z_1, Y) \\ \text{cov}(Z_2, Y) \end{pmatrix}. \end{aligned}$$

So

$$\psi_1 = \frac{\text{cov}(Z_1, Y) \text{var}(Z_2) - \text{cov}(Z_2, Y) \text{cov}(Z_1, Z_2)}{\text{var}(Z_1) \text{var}(Z_2) - \text{cov}(Z_1, Z_2)^2}$$

and

$$\psi_2 = \frac{\text{cov}(Z_2, Y) \text{var}(Z_1) - \text{cov}(Z_1, Y) \text{cov}(Z_1, Z_2)}{\text{var}(Z_1) \text{var}(Z_2) - \text{cov}(Z_1, Z_2)^2}.$$

Now substitute

$$\text{cov}(Z_\ell, Y) = \beta^* \text{cov}(Z_\ell, X) + \sum_{j \in \{1, 2\}} \gamma_j^* \text{cov}(Z_\ell, Z_j)$$

to obtain the desired function $Q_{\text{MT}}^*(\delta, \gamma_1, \gamma_2)$. As in the baseline case, notice that this function does not actually depend on the value of β . \square

Proof of proposition 15. This follows immediately from the proof of theorem 1. \square

J Duranton et al. (2014): Analysis of export value

Duranton et al. (2014) consider two different outcome variables: Propensity to export weight and propensity to export value. In section 4.1 we focused on export weight. Here we briefly discuss the results for export value. Panel A of table 11 reproduces columns 5–8 of table 5 in Duranton et al. (2014). These are their main results for export value. As with our analysis of export weight, we add the falsification adaptive set to the last row of panel A. In panel B we present results using plan and exploration as instruments, controlling for railroads. Compared to their baseline results, our modified baseline results in panel B suggest that the effect of within city highways on propensity to export value is potentially large and negative. Their theoretical model, however, predicts that this effect should be small and positive (their comparative static 3). That said, there is substantial model uncertainty: The falsification adaptive set in column 4 is $[-0.19, 0.047]$, which includes small positive values along with the large negative values. There is substantial statistical uncertainty as well. Thus the data does not allow us to draw a definitive conclusion about the validity of their model’s predicted comparative static.

K Proofs for section 3

Proof of proposition 1. This result is well known, but we include a proof for completeness. Suppose equation (2) holds for all $m, \ell \in \{1, \dots, L\}$. We will construct a joint distribution (Y, X, Z, \tilde{U}) and a parameter $\tilde{\beta}$ consistent with the data and assumptions A1–A4.

By the relevance assumption A1, there exists an ℓ such that $\text{cov}(X, Z_\ell) \neq 0$. Let

$$\tilde{\beta} = \frac{\text{cov}(Y, Z_\ell)}{\text{cov}(X, Z_\ell)}.$$

Let $\tilde{U} = Y - X\tilde{\beta}$. For every $m \in \{1, \dots, L\}$,

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

Thus A3 holds. A4 holds by definition of \tilde{U} . A2 holds automatically. Thus the model is not refuted.

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

$$\begin{aligned} 0 &= \text{cov}(U, Z_\ell) \\ &= \text{cov}(Y - X\beta, Z_\ell) \\ &= \text{cov}(Y, Z_\ell) - \beta \text{cov}(X, Z_\ell) \end{aligned}$$

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

Suppose $\beta \neq 0$. Suppose $\text{cov}(X, Z_\ell) = 0$. Then the above equation implies that we must have $\text{cov}(Y, Z_\ell) = 0$ as well. Hence equation (2) holds for this ℓ and any other $m \in \{1, \dots, L\}$.

Table 11: Baseline 2SLS results for Duranton et al. (2014): The effect of highways on export value. Non-highlighted parts reproduce results from their paper. Highlighted parts are new. Panel A reproduces columns 1–4 of their table 5. It also shows the estimated falsification adaptive set. Panel B uses only two of their instruments, controlling for the other. See text for discussion.

	Dependent variable: Export value			
	(1)	(2)	(3)	(4)
Panel A. Plan, exploration, and railroads used as instruments				
log highway km	1.10*** (0.17)	0.17 (0.16)	0.070 (0.14)	-0.026 (0.12)
log employment		0.91*** (0.091)	1.19** (0.58)	0.90** (0.42)
Market access (export)		-0.19 (0.12)	-0.38*** (0.14)	-0.36*** (0.11)
log 1920 population			-0.34 (0.30)	-0.23 (0.30)
log 1950 population			0.95* (0.48)	0.49 (0.51)
log 2000 population			-0.84 (0.72)	-0.14 (0.58)
log % manuf. emp.				0.83*** (0.16)
First-stage F stat.	97.5	90.3	80	84.8
Overid. p -value	0.081	0.071	0.28	0.55
FAS	[0.24, 0.78]	[-0.85, -0.085]	[-0.85, -0.01]	[-0.19, 0.047]
Panel B. Plan and exploration used as instruments, controlling for railroads				
log highway km	0.69** (0.30)	-0.22 (0.23)	-0.16 (0.18)	-0.14 (0.15)
log 1898 railroad km	0.41** (0.20)	0.32** (0.14)	0.22* (0.13)	0.11 (0.11)
First-stage F stat.	61.1	65.4	77.8	82.2
Overid. p -value	0.60	0.47	0.41	0.74

Notes: 66 observations per column. All specifications include a constant. Heteroskedasticity robust standard errors in parentheses. ***, **, *: statistically significant at 1%, 5%, 10%.

Finally, consider any pair m and ℓ such that $\text{cov}(X, Z_m) \neq 0$ and $\text{cov}(X, Z_\ell) \neq 0$. Then the above equation implies

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

Thus we have shown that equation (2) holds for all $m, \ell \in \{1, \dots, L\}$. \square

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 the instrument exogeneity A3,

$$\begin{aligned} \text{cov}(Z, Y) &= \text{cov}(Z, X'\beta + Z'\gamma + U) \\ &= \text{cov}(Z, X)\beta + \text{var}(Z)\gamma. \end{aligned}$$

By A2,

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

Since $-\delta \leq \gamma \leq \delta$ (component-wise) 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

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

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

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

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

Proof of corollary 1. Write the identified set from theorem 1 as

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

Equation (4) follows immediately by considering the three cases $\pi_\ell = 0$, $\pi_\ell < 0$, and $\pi_\ell > 0$ separately. \square

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 $\tilde{X} = (X, Z_2, \dots, Z_L)$. Hence

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

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

$$e_1' \text{cov}(Z, \tilde{X}_1)^{-1} = \begin{pmatrix} (\text{cov}(Z_1, X) - \text{cov}(Z_1, Z_{-1}) \text{var}(Z_{-1})^{-1} \text{cov}(Z_{-1}, X))^{-1} \\ -(\text{cov}(Z_1, X) - \text{cov}(Z_1, Z_{-1}) \text{var}(Z_{-1})^{-1} \text{cov}(Z_{-1}, X))^{-1} \text{cov}(Z_1, Z_{-1}) \text{var}(Z_{-1})^{-1} \end{pmatrix}'.$$

Hence

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

Here we defined $\tilde{Z}_1 = Z_1 - \text{cov}(Z_1, Z_{-1}) \text{var}(Z_{-1})^{-1} Z_{-1}$. This is the population residual of Z_1 after removing the projection of Z_1 onto Z_{-1} . The last line thus follows by the partitioned regression formula. \square

Proof of equation (5) on page 14. As in the proof of lemma 1, without loss of generality suppose $\ell = 1$. Suppose the baseline model ($\delta = 0$) holds. Then

$$\begin{aligned} \frac{\psi_1}{\pi_1} &= \frac{\text{cov}(\tilde{Z}_1, Y)}{\text{cov}(\tilde{Z}_1, X)} \\ &= \frac{\text{cov}(\tilde{Z}_1, \beta X + U)}{\text{cov}(\tilde{Z}_1, X)} \\ &= \beta \frac{\text{cov}(\tilde{Z}_1, X)}{\text{cov}(\tilde{Z}_1, X)} \\ &= \beta. \end{aligned}$$

Similarly, $\text{cov}(Z_1, Y)/\text{cov}(Z_1, X) = \beta$. Equation (5) follows. \square

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 (6).

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

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

The first line follows by equation (4), and by the definition of $\delta(b)$. The second line also uses the definition of $\delta(b)$. The remaining lines follow by considering two cases so that we can eliminate the absolute values. \square

Proof of proposition 2. Let FF denote the true falsification frontier. Let

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

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} , $\delta_\ell = |\psi_\ell|$.

- Suppose $\psi_\ell > 0$. Then $\delta_\ell = \psi_\ell$ and hence $[\psi_\ell - \delta_\ell, \psi_\ell + \delta_\ell] = [0, 2\psi_\ell]$. Since $\delta'_\ell < \delta_\ell$, $\psi_\ell - \delta'_\ell > 0$ and so 0 is not an element of $[\psi_\ell - \delta'_\ell, \psi_\ell + \delta'_\ell]$. Hence $B_\ell(\delta'_\ell) = \emptyset$ by equation (4).
- Suppose $\psi_\ell < 0$. Then $\delta_\ell = -\psi_\ell$ and hence $[\psi_\ell - \delta_\ell, \psi_\ell + \delta_\ell] = [2\psi_\ell, 0]$. Since $\delta'_\ell < \delta_\ell$, $\psi_\ell + \delta'_\ell < 0$ and so 0 is not an element of $[\psi_\ell - \delta'_\ell, \psi_\ell + \delta'_\ell]$. Hence $B_\ell(\delta'_\ell) = \emptyset$ by equation (4).
- Suppose $\psi_\ell = 0$. Then $\delta_\ell = 0$ and so we cannot have $\delta'_\ell < \delta_\ell$, since $\delta'_\ell \geq 0$.

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

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

Consider two cases:

- i. Suppose $b^* \leq \psi_\ell/\pi_\ell$. Then

$$\begin{aligned} \frac{\psi_\ell}{\pi_\ell} - \frac{\delta_\ell}{|\pi_\ell|} &= \frac{\psi_\ell}{\pi_\ell} - \frac{|\psi_\ell - b^*\pi_\ell|}{|\pi_\ell|} \\ &= \frac{\psi_\ell}{\pi_\ell} - \left| \frac{\psi_\ell}{\pi_\ell} - b^* \right| \\ &= b^*. \end{aligned}$$

Hence

$$\begin{aligned} \frac{\psi_\ell}{\pi_\ell} - \frac{\delta'_\ell}{|\pi_\ell|} &> \frac{\psi_\ell}{\pi_\ell} - \frac{\delta_\ell}{|\pi_\ell|} \\ &= b^*. \end{aligned}$$

Thus $b^* \notin B_\ell(\delta')$.

- ii. Suppose $b^* > \psi_\ell/\pi_\ell$. Then

$$\begin{aligned} \frac{\psi_\ell}{\pi_\ell} + \frac{\delta_\ell}{|\pi_\ell|} &= \frac{\psi_\ell}{\pi_\ell} + \frac{|\psi_\ell - b^*\pi_\ell|}{|\pi_\ell|} \\ &= \frac{\psi_\ell}{\pi_\ell} + \left| \frac{\psi_\ell}{\pi_\ell} - b^* \right| \\ &= b^*. \end{aligned}$$

Hence

$$\begin{aligned} \frac{\psi_\ell}{\pi_\ell} + \frac{\delta'_\ell}{|\pi_\ell|} &< \frac{\psi_\ell}{\pi_\ell} + \frac{\delta_\ell}{|\pi_\ell|} \\ &= b^*. \end{aligned}$$

Thus $b^* \notin B_\ell(\delta')$.

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 $\delta \notin \text{FF}^{\text{guess}}$. There are three cases to consider.

- (a) Suppose $\mathcal{B}(\delta) = \emptyset$. Then $\delta \notin \text{FF}$ by definition.
(b) Denote

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

Suppose $\mathcal{B}(\delta) \neq \emptyset$ and $\mathcal{B}(\delta) \subseteq [b_{\min}, b_{\max}]$. We'll show that we can find a $\delta' < \delta$ such that $\mathcal{B}(\delta') \neq \emptyset$, and hence $\delta \notin \text{FF}$. We use two observations:

- i. Since $\delta \notin \text{FF}^{\text{guess}}$ we must have $\delta \neq \delta(b)$ for all $b \in [b_{\min}, b_{\max}]$. In particular, $\delta \neq \delta(b)$ for all $b \in \mathcal{B}(\delta)$, since we're considering the case where $\mathcal{B}(\delta) \subseteq [b_{\min}, b_{\max}]$.
- ii. Let $b \in \mathcal{B}(\delta)$. By the characterization of $\mathcal{B}(\delta)$ in theorem 1 (specifically, see equation (22) in our discussion below), this implies that $|\psi_\ell - b\pi_\ell| \leq \delta_\ell$ for all $\ell = 1, \dots, L$. That is, $\delta(b) \leq \delta$.

Let b' be any element of $\mathcal{B}(\delta)$. This exists by assumption. Let $\delta' = \delta(b')$. Observation (i) implies $\delta' \neq \delta$. Observation (ii) implies $\delta' \leq \delta$. Hence $\delta' < \delta$. Moreover, we have $\mathcal{B}(\delta') = \mathcal{B}(\delta(b')) = \{b'\} \neq \emptyset$ by lemma 2. Thus $\delta \notin \text{FF}$, by definition of the falsification frontier.

- (c) Suppose $\mathcal{B}(\delta)$ contains an element $b \notin [b_{\min}, b_{\max}]$. Suppose $b > b_{\max}$. Let $\delta' = \delta(b_{\max})$.
- Observation (i) above implies $\delta' \neq \delta$.
 - Next we show $\delta' \leq \delta$. Observation (ii) above implies $\delta_\ell(b) = |\psi_\ell - b\pi_\ell| \leq \delta_\ell$ for all $\ell = 1, \dots, L$. So it suffices to show

$$\delta_\ell(b_{\max}) = |\psi_\ell - b_{\max}\pi_\ell| \leq |\psi_\ell - b\pi_\ell| = \delta_\ell(b)$$

for all ℓ . This holds immediately for all ℓ with $\pi_\ell = 0$. For the other ℓ 's, it follows by the structure of the $\delta_\ell(\cdot)$ function, the definition

$$b_{\max} = \max_{\ell=1, \dots, L: \pi_\ell \neq 0} \frac{\psi_\ell}{\pi_\ell},$$

and the fact that $b > b_{\max}$. Specifically, notice that for ℓ such that $\pi_\ell \neq 0$, $\delta_\ell(\cdot)$ is zero at ψ_ℓ/π_ℓ and otherwise is strictly increasing as the distance between ψ_ℓ/π_ℓ and its argument increases. (Specifically, it is a V-shaped function, whose slope is determined by π_ℓ , but this specific shape is not important here.) Since $b > b_{\max}$, $b > \psi_s/\pi_s$ for all s with $\pi_s \neq 0$. Thus we must have $\delta_\ell(b) > \delta_\ell(\psi_s/\pi_s)$ for all ℓ and s with $\pi_\ell \neq 0$ and $\pi_s \neq 0$. In particular, it holds for the s such that $b_{\max} = \psi_s/\pi_s$.

Hence $\delta' < \delta$. Moreover, $\mathcal{B}(\delta') = \mathcal{B}(\delta(b_{\max})) = \{b_{\max}\} \neq \emptyset$ by lemma 2. So $\delta \notin \text{FF}$, by definition of the falsification frontier. A similar argument applies if instead we have $b < b_{\min}$.

□

Proof of corollary 2. The falsification frontier is given by

$$\text{FF} = \left\{ \delta \in \mathbb{R}_{\geq 0}^2 : \delta_\ell = |\psi_\ell - b\pi_\ell|, \ell = 1, 2, b \in \left[\min_{\ell=1,2} \frac{\psi_\ell}{\pi_\ell}, \max_{\ell=1,2} \frac{\psi_\ell}{\pi_\ell} \right] \right\},$$

a line segment in \mathbb{R}^2 . We can directly see that FF is the line segment between

$$\left(0, \left| \psi_2 - \frac{\psi_1}{\pi_1} \pi_2 \right| \right) \quad \text{and} \quad \left(\left| \psi_1 - \frac{\psi_2}{\pi_2} \pi_1 \right|, 0 \right)$$

which corresponds the equation

$$|\pi_2| \left| \frac{\psi_1}{\pi_1} - \frac{\psi_2}{\pi_2} \right| - \left| \frac{\pi_2}{\pi_1} \right| \delta_1 = \delta_2$$

for nonnegative values of δ_1 and δ_2 . □

Proof of theorem 2. We have

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

The first equality follows by the characterization of the falsification frontier in proposition 2. The second equality follows by lemma 2. □

Proof of proposition 3. Let

$$m^* = \max_{\ell, \ell' \in \{1, \dots, L\}} \frac{\frac{\psi_\ell}{\pi_\ell} - \frac{\psi_{\ell'}}{\pi_{\ell'}}}{\frac{d_\ell}{|\pi_\ell|} + \frac{d_{\ell'}}{|\pi_{\ell'}|}}.$$

Suppose $m < m^*$. Then there is an s and an s' such that

$$m < \frac{\frac{\psi_s}{\pi_s} - \frac{\psi_{s'}}{\pi_{s'}}}{\frac{d_s}{|\pi_s|} - \frac{d_{s'}}{|\pi_{s'}|}}.$$

Rearranging yields

$$\frac{\psi_s}{\pi_s} - \frac{md_s}{|\pi_s|} > \frac{\psi_{s'}}{\pi_{s'}} + \frac{md_{s'}}{|\pi_{s'}|}.$$

Hence

$$\max_{\ell=1, \dots, L} \left(\frac{\psi_\ell}{\pi_\ell} - \frac{md_\ell}{|\pi_\ell|} \right) > \min_{\ell'=1, \dots, L} \left(\frac{\psi_{\ell'}}{\pi_{\ell'}} + \frac{md_{\ell'}}{|\pi_{\ell'}|} \right).$$

But recall from corollary 4 that for any δ ,

$$\begin{aligned} \mathcal{B}(\delta) &= \bigcap_{\ell=1, \dots, L} \left[\frac{\psi_\ell}{\pi_\ell} - \frac{\delta_\ell}{|\pi_\ell|}, \frac{\psi_\ell}{\pi_\ell} + \frac{\delta_\ell}{|\pi_\ell|} \right] \\ &= \left[\max_{\ell=1, \dots, L} \left(\frac{\psi_\ell}{\pi_\ell} - \frac{md_\ell}{|\pi_\ell|} \right), \min_{\ell'=1, \dots, L} \left(\frac{\psi_{\ell'}}{\pi_{\ell'}} + \frac{md_{\ell'}}{|\pi_{\ell'}|} \right) \right]. \end{aligned}$$

We just showed, however, that for $\delta = m \cdot d$, the left endpoint is larger than the right endpoint. Hence $\mathcal{B}(m \cdot d)$ is empty.

Now, let $\delta = m^* \cdot d$. Let s, s' be such that

$$m^* = \frac{\frac{\psi_s}{d_s} - \frac{\psi_{s'}}{d_{s'}}}{\frac{|\pi_s|}{|\pi_{s'}|}}.$$

Rearranging yields

$$\frac{\psi_s}{\pi_s} - \frac{m^* d_s}{|\pi_s|} = \frac{\psi_{s'}}{\pi_{s'}} + \frac{m^* d_{s'}}{|\pi_{s'}|}.$$

Let b^* denote this common value.

Next let $\ell \in \{1, \dots, L\}$ be arbitrary. We will show that $b^* \in B_\ell(\delta)$. We prove this by contradiction. Consider two cases.

1. Suppose

$$\frac{\psi_\ell}{\pi_\ell} + \frac{m^* d_\ell}{|\pi_\ell|} < b^*.$$

Rearranging this equation yields

$$\frac{\frac{\psi_s}{d_s} - \frac{\psi_\ell}{d_\ell}}{\frac{|\pi_s|}{|\pi_\ell|} + \frac{d_\ell}{d_s}} > m^*.$$

This contradicts the definition of s .

2. Suppose

$$b^* < \frac{\psi_\ell}{\pi_\ell} + \frac{m^* d_\ell}{|\pi_\ell|}.$$

Rearranging this equation yields

$$\frac{\frac{\psi_\ell}{d_\ell} - \frac{\psi_{s'}}{d_{s'}}}{\frac{|\pi_\ell|}{|\pi_{s'}|} + \frac{d_{s'}}{d_\ell}} > m^*.$$

This contradicts the definition of s' .

Thus $b^* \in B_\ell(\delta)$ for any $\ell \in \{1, \dots, L\}$. Hence $b^* \in \mathcal{B}(\delta)$. Thus implies that $\mathcal{B}(\delta)$ is not empty, as desired. \square

Proof of corollary 3. This result follows immediately from corollary 1, proposition 2, and proposition 3. \square

Proofs for section 3.6: K endogenous variables

To prove the results in this section, we first review the structure of the identified set for β . From theorem 1, this set is

$$\begin{aligned}\mathcal{B}(\delta) &= \{b \in \mathbb{R}^K : -\delta \leq \psi - \Pi b \leq \delta\} \\ &= \bigcap_{\ell=1}^L \{b \in \mathbb{R}^K : -\delta_\ell \leq \psi_\ell - \pi'_\ell b \leq \delta_\ell\}\end{aligned}$$

where π'_ℓ is the ℓ th row of the matrix Π . So if we let

$$B_\ell(\delta_\ell) = \{b \in \mathbb{R}^K : -\delta_\ell \leq \psi_\ell - \pi'_\ell b \leq \delta_\ell\}$$

then

$$\mathcal{B}(\delta) = \bigcap_{\ell=1}^L B_\ell(\delta_\ell). \quad (22)$$

In our baseline model with $\delta_\ell = 0$, $B_\ell(0)$ is a hyperplane in \mathbb{R}^K :

$$B_\ell(0) = \{b \in \mathbb{R}^K : \psi_\ell = \pi'_\ell b\}.$$

As we relax exclusion, this set goes from being a hyperplane to the space between two hyperplanes. Figures 12, 13, and 14 plot examples of these spaces. The overall identified set $\mathcal{B}(\delta)$ is just the intersection of these spaces.

With that as background, we next prove our main results: proposition 5, our characterization of the falsification frontier, and theorem 4, our characterization of the falsification adaptive set. We do this using a sequence of lemmas.

We begin by showing a basic geometric fact about the set FAS* when $L = K + 1$. Here and elsewhere we use the notation $\beta_{-\ell}^{2\text{SLS}} = \beta_{\{1, \dots, L\} \setminus \{\ell\}}^{2\text{SLS}}$.

Lemma 3. Suppose A1', A2, and A3 hold. Suppose $L = K + 1$. Then exactly one of the following holds:

1. $\beta_{-\ell}^{2\text{SLS}} = \beta_{-\ell'}^{2\text{SLS}}$ for all $\ell, \ell' \in \{1, \dots, L\}$.
2. $\pi'_\ell \beta_{-\ell}^{2\text{SLS}} \neq \psi_\ell$ for all $\ell \in \{1, \dots, L\}$.

Proof of lemma 3. We will show that (1) holds if and only if (2) does not hold.

1. (\Rightarrow) Suppose statement (1) holds. So $\beta_{-\ell}^{2\text{SLS}} = \beta_{-\ell'}^{2\text{SLS}}$ for all $\ell, \ell' \in \{1, \dots, L\}$. Then $\pi'_\ell \beta_{-\ell}^{2\text{SLS}} = \pi'_{\ell'} \beta_{-\ell'}^{2\text{SLS}} = \psi_{\ell'}$ for any ℓ and any $\ell' \neq \ell$. Therefore statement (2) does not hold.
2. (\Leftarrow). Suppose (2) does not hold. Then there exists an ℓ such that $\pi'_\ell \beta_{-\ell}^{2\text{SLS}} = \psi_\ell$. By definition of $\beta_{-\ell}^{2\text{SLS}}$, we have $\pi'_{\ell'} \beta_{-\ell}^{2\text{SLS}} = \psi_{\ell'}$ for all $\ell' \neq \ell$. But now we know it also holds for $\ell' = \ell$. Thus $\beta_{-\ell}^{2\text{SLS}}$ is an element of *every* hyperplane $B_{\ell'}(0)$. This implies that

$$(\Pi_{\{1, \dots, L\} \setminus \{\ell'\}}) b = \psi_{\{1, \dots, L\} \setminus \{\ell'\}}$$

is satisfied by $b = \beta_{-\ell}^{2\text{SLS}}$. But this equation is also satisfied by $b = \beta_{-\ell'}^{2\text{SLS}}$, by definition of this estimand. By A1', this equation has a unique solution. Thus $\beta_{-\ell}^{2\text{SLS}} = \beta_{-\ell'}^{2\text{SLS}}$. This argument holds for all $\ell' \in \{1, \dots, L\} \setminus \{\ell\}$. Thus (1) holds.

□

The next lemma shows that, when $L = K + 1$ and FAS^* is not a singleton, we can write any element of \mathbb{R}^K as a weighted sum of our just-identified 2SLS estimands.

Lemma 4. Suppose A1' holds. Suppose $L = K + 1$. Assume that FAS^* is not a singleton. Then for any $b \in \mathbb{R}^K$ there exist weights $w_1(b), \dots, w_L(b)$ such that

$$b = \sum_{\ell=1}^L w_{\ell}(b) \beta_{-\ell}^{2\text{SLS}}$$

and the weights that sum to one, $\sum_{\ell=1}^L w_{\ell}(b) = 1$.

Proof of lemma 4. Since we want the weights to sum to one, the last weight is always defined by

$$1 - \sum_{\ell=1}^{L-1} w_{\ell}(b)$$

once we've chosen the first $L - 1$ weights. Thus we want to find $(w_1(b), \dots, w_{L-1}(b))$ such that

$$b = \begin{pmatrix} \beta_{-1}^{2\text{SLS}} & \dots & \beta_{-L}^{2\text{SLS}} \end{pmatrix} \begin{pmatrix} w_1(b) \\ \vdots \\ 1 - \sum_{\ell=1}^{L-1} w_{\ell}(b) \end{pmatrix}$$

holds. Multiplying this out we get

$$b = \sum_{s=1}^{L-1} \beta_{-s}^{2\text{SLS}} w_s(b) + \left(\beta_{-L}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \sum_{s=1}^{L-1} w_s(b) \right)$$

or

$$b - \beta_{-L}^{2\text{SLS}} = \sum_{s=1}^{L-1} (\beta_{-s}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}}) w_s(b).$$

Thus in matrix notation we have

$$b - \beta_{-L}^{2\text{SLS}} = \underbrace{\begin{pmatrix} \beta_{-1}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} & \dots & \beta_{-(L-1)}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \end{pmatrix}}_{(L-1) \times (L-1)} \begin{pmatrix} w_1(b) \\ \vdots \\ w_{L-1}(b) \end{pmatrix}.$$

This system has a unique solution if the matrix

$$\begin{pmatrix} \beta_{-1}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} & \dots & \beta_{-(L-1)}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \end{pmatrix}$$

has full rank. So that's what we'll show.

This matrix doesn't have full rank if we can find a $c \in \mathbb{R}^{L-1}$ with $c \neq \mathbf{0}_{L-1}$ such that

$$\begin{pmatrix} \beta_{-1}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} & \dots & \beta_{-(L-1)}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \end{pmatrix} c = \mathbf{0}_{L-1}. \quad (23)$$

Multiplying out we see that this equality is equivalent to

$$\sum_{\ell=1}^{L-1} c_\ell \beta_{-\ell}^{2\text{SLS}} = \beta_{-L}^{2\text{SLS}} \sum_{\ell=1}^{L-1} c_\ell.$$

Premultiplying by π'_L gives

$$\begin{aligned} \pi'_L \beta_{-L}^{2\text{SLS}} \sum_{\ell=1}^{L-1} c_\ell &= \sum_{\ell=1}^{L-1} c_\ell \pi'_L \beta_{-\ell}^{2\text{SLS}} \\ &= \psi_L \sum_{\ell=1}^{L-1} c_\ell. \end{aligned}$$

The second line follows since $\beta_{-\ell}^{2\text{SLS}}$ lies on the line $\psi_L = \pi'_L \tilde{b}$ for all $\ell \neq L$, by definition of these 2SLS estimands. Thus

$$(\psi_L - \pi'_L \beta_{-L}^{2\text{SLS}}) \sum_{\ell=1}^{L-1} c_\ell = 0.$$

By lemma 3, $\psi_L \neq \pi'_L \beta_{-L}^{2\text{SLS}}$. So we must have

$$\sum_{\ell=1}^{L-1} c_\ell = 0.$$

We can now go back and substitute this into equation (23) to get

$$\left(\beta_{-1}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \quad \dots \quad \beta_{-(L-1)}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \right) \begin{pmatrix} c_1 \\ \vdots \\ -\sum_{\ell=1}^{L-2} c_\ell \end{pmatrix} = \mathbf{0}_{L-1}.$$

Multiplying this out yields

$$\sum_{\ell=1}^{L-2} c_\ell \beta_{-\ell}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \sum_{\ell=1}^{L-2} c_\ell + (\beta_{-(L-1)} - \beta_{-L}) \left(-\sum_{\ell=1}^{L-2} c_\ell \right) = \mathbf{0}_{L-1}.$$

Simplifying gives

$$\beta_{-(L-1)}^{2\text{SLS}} \sum_{\ell=1}^{L-2} c_\ell = \sum_{\ell=1}^{L-2} c_\ell \beta_{-\ell}^{2\text{SLS}}.$$

Premultiply both sides by π_{L-1} and repeat the arguments above to get

$$\sum_{\ell=1}^{L-2} c_\ell = 0.$$

Repeat this argument iteratively to see that $c = \mathbf{0}_{L-1}$. Thus the matrix

$$\left(\beta_{-1}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \quad \dots \quad \beta_{-(L-1)}^{2\text{SLS}} - \beta_{-L}^{2\text{SLS}} \right)$$

is nonsingular. □

Now recall our general definition of the falsification frontier: It is the set of values $\delta \in \mathbb{R}_{\geq 0}^L$ such that $\mathcal{B}(\delta) \neq \emptyset$ and $\mathcal{B}(\delta') = \emptyset$ for any $\delta' < \delta$. We first show that, in the $L = K + 1$ case, the identified set for β is a singleton at any point δ in the conjectured falsification frontier defined in equation (10). For this we define

$$\delta_\ell(b) = |\psi_\ell - \pi'_\ell b|$$

for all $\ell = 1, \dots, L$. Let $\delta(b) = (\delta_1(b), \dots, \delta_L(b))$.

Lemma 5. Suppose A1', A2, and A3 hold. Suppose $L = K + 1$. Let $b \in \text{FAS}^*$. Then $\mathcal{B}(\delta(b)) = \{b\}$.

Figure 12 illustrates the geometric intuition for this lemma.

Proof of lemma 5. By lemma 3, there are two cases to consider: FAS^* is either a singleton or a nondegenerate simplex in \mathbb{R}^K .

Case 1. Suppose $\text{FAS}^* = \{b\}$ is a singleton. By the definition of FAS^* , this implies that $b = \beta_{\mathcal{L}}^{2\text{SLS}}$ for any $\mathcal{L} \subseteq \{1, \dots, L\}$ with $|\mathcal{L}| = K$. Moreover, for any such \mathcal{L} ,

$$\begin{aligned} \bigcap_{\ell \in \mathcal{L}} B_\ell(0) &= \{\tilde{b} \in \mathbb{R}^K : \Pi_{\mathcal{L}} \tilde{b} = \psi_{\mathcal{L}}\} \\ &= \{b\}. \end{aligned}$$

Let \mathcal{L}_1 and \mathcal{L}_2 be subsets of $\{1, \dots, L\}$ such that $\mathcal{L}_1 \cup \mathcal{L}_2 = \mathcal{L}$. Then

$$\begin{aligned} \mathcal{B}(0) &= \bigcap_{\ell=1}^L B_\ell(0) \\ &= \left(\bigcap_{\ell \in \mathcal{L}_1} B_\ell(0) \right) \cap \left(\bigcap_{\ell \in \mathcal{L}_2} B_\ell(0) \right) \\ &= \{b\}. \end{aligned}$$

The first line holds by equation (22). Thus $\mathcal{B}(\delta(b)) = \mathcal{B}(0) = \{b\}$.

Case 2. Suppose FAS^* is not a singleton. Then $\pi'_\ell \beta_{-\ell}^{2\text{SLS}} \neq \psi_\ell$ for all $\ell \in \{1, \dots, L\}$. We prove equality of sets by showing that both directions of set inclusion hold.

Step 1 (\supseteq). First we show that $\mathcal{B}(\delta(b)) \supseteq \{b\}$. By definition of $\delta_\ell(\cdot)$,

$$\psi_\ell - \pi'_\ell b \in [-\delta_\ell(b), \delta_\ell(b)]$$

for all ℓ . Thus, by the characterization of $\mathcal{B}(\cdot)$ in theorem 1, $b \in \mathcal{B}(\delta(b))$.

Step 2 (\subseteq). Next we show that $\mathcal{B}(\delta(b)) \subseteq \{b\}$. First suppose $\delta_\ell(b) = 0$ for all ℓ . In this case the baseline model is not refuted. b is then the unique value that satisfies all of the instrument constraints. So FAS^* is a singleton. This is a contradiction. So we must have $\delta_\ell(b) > 0$ for some ℓ .

We will show that any element $b^* \neq b$ is not in $\mathcal{B}(\delta(b))$. The set FAS^* is a polytope. Consider its alternative half-space representation. The half-spaces correspond to one side of the hyperplanes

$B_\ell(0)$. For example, the shaded areas in figure 5 can be written as the intersection of half-spaces formed by the plotted lines. Formally, write

$$\text{FAS}^* = \bigcap_{\ell=1}^L \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \leq 0\}. \quad (24)$$

Here we assume without loss of generality that all of the inequalities go in the same direction (≤ 0).

- If one of the half-spaces originally has the form $\psi_\ell - \pi'_\ell \tilde{b} \geq 0$, it can be rewritten as $-\psi_\ell - (-\pi'_\ell \tilde{b}) \leq 0$. This is equivalent to replacing Z_ℓ with $-Z_\ell$. As shown in appendix E, changing the sign of an instrument does not affect the estimands $\beta_{\mathcal{L}}^{2\text{SLS}}$ and therefore does not affect the set FAS^* .

Now recall that

$$\begin{aligned} \mathcal{B}(\delta) &= \bigcap_{\ell=1}^L \left\{ \tilde{b} \in \mathbb{R}^K : -\delta_\ell \leq \psi_\ell - \pi'_\ell \tilde{b} \leq \delta_\ell \right\} \\ &= \left(\bigcap_{\ell=1}^L \{ \tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq -\delta_\ell \} \right) \cap \left(\bigcap_{\ell=1}^L \{ \tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \leq \delta_\ell \} \right). \end{aligned}$$

Evaluating this expression at $\delta(b)$ yields

$$\begin{aligned} \mathcal{B}(\delta(b)) &= \left(\bigcap_{\ell=1}^L \{ \tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq -|\psi_\ell - \pi'_\ell b| \} \right) \cap \left(\bigcap_{\ell=1}^L \{ \tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \leq |\psi_\ell - \pi'_\ell b| \} \right) \\ &\equiv \mathcal{P}_1(b) \cap \mathcal{P}_2(b). \end{aligned}$$

By lemma 4, any element in \mathbb{R}^K can be written as a linear combination of our L different just-identified 2SLS estimands. In particular, we can write b^* in this way:

$$b^* = \sum_{\ell=1}^L w_\ell(b^*) \beta_{-\ell}^{2\text{SLS}}$$

where $\beta_{-\ell}^{2\text{SLS}} = \beta_{\{1, \dots, L\} \setminus \{\ell\}}^{2\text{SLS}}$ and where $w_\ell(b^*)$ are weights that sum to one, $\sum_{\ell=1}^L w_\ell(b^*) = 1$. (These weights are not necessarily non-negative.) Remember that our goal is to show that $b^* \neq b$ implies $b^* \notin \mathcal{B}(\delta(b))$. By our decomposition above, it suffices to show that $b^* \notin \mathcal{P}_1(b)$.

Since $b \in \text{FAS}^*$,

$$\psi_\ell - \pi'_\ell b \leq 0$$

for all ℓ . This follows directly from our half-space representation of FAS^* . Thus

$$-|\psi_\ell - \pi'_\ell b| = \psi_\ell - \pi'_\ell b$$

for all ℓ . Hence

$$\begin{aligned}
\mathcal{P}_1(b) &= \bigcap_{\ell=1}^L \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq -|\psi_\ell - \pi'_\ell b|\} \\
&= \bigcap_{\ell=1}^L \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq \psi_\ell - \pi'_\ell b\} \\
&= \bigcap_{\ell=1}^L \{\tilde{b} \in \mathbb{R}^K : \pi'_\ell(\tilde{b} - b) \leq 0\}.
\end{aligned}$$

So $b^* \in \mathcal{P}_1(b)$ if and only if $\pi'_\ell(b^* - b) \leq 0$ for all ℓ . Focus on just one ℓ for a moment. Then

$$\begin{aligned}
\pi'_\ell(b^* - b) &= \sum_{s=1}^L (w_s(b^*) - w_s(b)) \pi'_\ell \beta_{-s}^{2\text{SLS}} \\
&= \sum_{s=1}^L (w_s(b^*) - w_s(b)) \pi'_\ell \beta_{-s}^{2\text{SLS}} - \pi'_\ell \beta_{-\ell}^{2\text{SLS}} \sum_{s=1}^L (w_s(b^*) - w_s(b)) \\
&= \sum_{s \neq \ell} (w_s(b^*) - w_s(b)) \pi'_\ell \beta_{-s}^{2\text{SLS}} - \pi'_\ell \beta_{-\ell}^{2\text{SLS}} \sum_{s \neq \ell} (w_s(b^*) - w_s(b)) \\
&= \sum_{s \neq \ell} (w_s(b^*) - w_s(b)) \psi_\ell - \pi'_\ell \beta_{-\ell}^{2\text{SLS}} \sum_{s \neq \ell} (w_s(b^*) - w_s(b)) \\
&= (\psi_\ell - \pi'_\ell \beta_{-\ell}^{2\text{SLS}}) \sum_{s \neq \ell} (w_s(b^*) - w_s(b)).
\end{aligned}$$

The first line follows by substituting in our linear combination representation of any element in \mathbb{R}^K for both b^* and b . The second line follows since the weights sum to one, so the sum of the difference in weights is zero. The third line comes from cancelling the ℓ th term in the two summations. The fourth line follows since, for all $s \neq \ell$, $\beta_{-s}^{2\text{SLS}}$ lies on the line $\psi_\ell = \pi'_\ell \tilde{b}$. This holds by the definition of these 2SLS estimands.

Next notice that $\psi_\ell - \pi'_\ell \beta_{-\ell}^{2\text{SLS}} < 0$. This follows from $\beta_{-\ell}^{2\text{SLS}} \in \text{FAS}^*$, the half-space representation of FAS^* , and the fact that FAS^* is a nondegenerate simplex. So suppose by way of contradiction that $b^* \in \mathcal{P}_1(b)$. Then $\pi'_\ell(b^* - b) \leq 0$ for all ℓ . We've just seen that this implies

$$\sum_{s \neq \ell} (w_s(b^*) - w_s(b)) \geq 0$$

for all ℓ . But now we have

$$\begin{aligned}
0 &= \sum_{s=1}^L (w_s(b^*) - w_s(b)) \\
&= \sum_{s \neq \ell} (w_s(b^*) - w_s(b)) + (w_\ell(b^*) - w_\ell(b)).
\end{aligned}$$

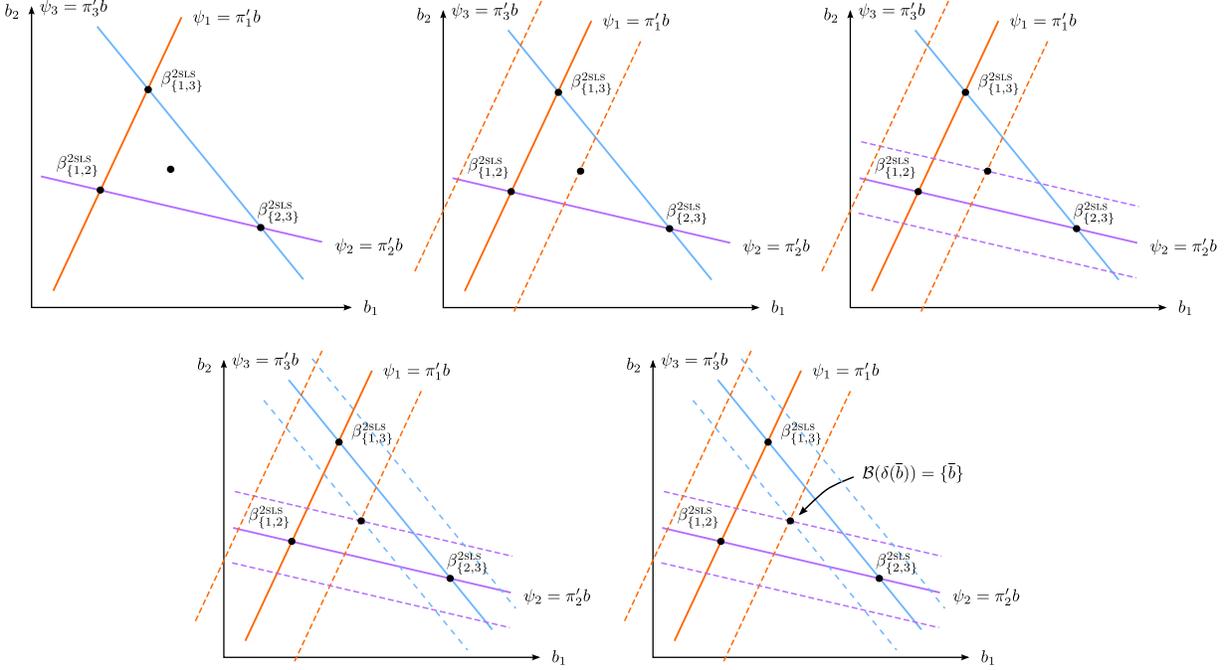


Figure 12: Geometric intuition for a key step in the proof of theorem 3. Start at the top left. Here we have $K = 2$ and $L = 3$, as in the left plot of figure 5. We pick a point inside the convex hull of the three just identified 2SLS estimands. Call it \bar{b} . We'll now illustrate the intuition for lemma 5, which shows that $\mathcal{B}(\delta(\bar{b})) = \{\bar{b}\}$. We next plot the sets $B_\ell(\delta_\ell(\bar{b})) = \{b \in \mathbb{R}^K : -\delta_\ell(\bar{b}) \leq \psi_\ell - \pi'_\ell b \leq \delta_\ell(\bar{b})\}$, adding one at a time. We look at relaxations $\delta_\ell(\bar{b}) = |\psi_\ell - \pi'_\ell \bar{b}|$ which are chosen to be as small as possible, while still making sure that \bar{b} is included in each of the tubes centered at the different hyperplanes $B_\ell(0)$. The result is that the identified set for β given the relaxation $\delta(\bar{b})$ —the intersection of these three tubes—is the singleton $\{\bar{b}\}$.

Thus

$$\begin{aligned} w_\ell(b^*) - w_\ell(b) &= - \sum_{s \neq \ell} (w_s(b^*) - w_s(b)) \\ &\leq 0. \end{aligned}$$

This holds for all ℓ . But the only way for the sum of non-positive numbers to equal zero is if they're all zero. Thus $w_\ell(b^*) = w_\ell(b)$ for all ℓ . This implies that $b^* = b$, a contradiction. Therefore $b^* \notin \mathcal{P}_1(b)$. As discussed above, this suffices to see that $b^* \notin \mathcal{B}(\delta(b))$. \square

Lemma 6. Suppose A1', A2, and A3 hold. Suppose $L \geq K + 1$. Let $b \in \mathcal{P}$. Then $\mathcal{B}(\delta(b)) = \{b\}$.

Proof of lemma 6. As in the proof of lemma 5, we prove set equality by showing that both directions of set inclusion hold.

Step 1 (\supseteq). The proof of this step from lemma 5 applies without modification.

Step 2 (\subseteq). Since $b \in \mathcal{P}$ there is some $\mathcal{L} \subseteq \{1, \dots, L\}$ with $|\mathcal{L}| = K$ such that $b \in \mathcal{P}_\mathcal{L}$. Let

$$\mathcal{B}_\mathcal{L}(\delta) = \bigcap_{\ell \in \mathcal{L}} B_\ell(\delta).$$

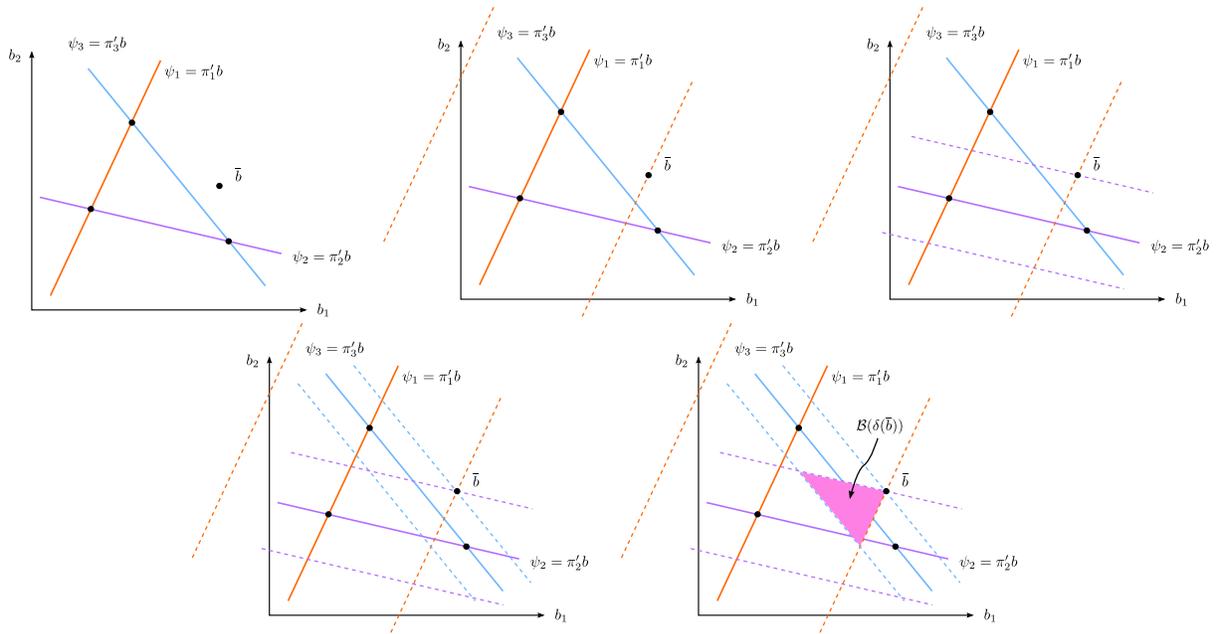


Figure 13: Geometric intuition for a key step in the proof of theorem 3, continued from figure 12. Here we instead pick a point *outside* the convex hull of the three just identified 2SLS estimands. We again relax each exclusion restriction just enough so that \bar{b} is inside each of the tubes centered at the different hyperplanes $B_\ell(0)$. But now when we look at the intersection of these tubes we get a non-singleton set. Because this set has an interior, we've relaxed exclusion too much—we can strengthen some of the exclusion restrictions without obtaining a falsified model. Thus the point \bar{b} is not in the falsification adaptive set.

This is the identified set for β when we impose partial exclusion for only the instruments with indices in \mathcal{L} . By lemma 5, $\mathcal{B}_\mathcal{L}(\delta(b)) = \{b\}$. Moreover, $\mathcal{B}_\mathcal{L}(\delta) \supseteq \mathcal{B}(\delta)$. Thus $\mathcal{B}(\delta(b)) \subseteq \{b\}$. \square

The following variation on Farkas' lemma (for example, corollary 22.3.1 on page 200 of Rockafellar 1970; Border 2019 provides an extensive discussion) is helpful.

Lemma 7 (Variation on Farkas' Lemma). Let $x_1, \dots, x_n \in \mathbb{R}^K$. Then $\mathbf{0} \notin \text{conv}(\{x_1, \dots, x_n\})$ if and only if there exists a point $p \in \mathbb{R}^K$ such that $p'x_i > 0$ for all $i = 1, \dots, n$.

In particular, this lemma shows that if you pick a point outside a convex hull of $\{x_1, \dots, x_n\}$, then you can find a vector p pointing towards all of the vectors x_1, \dots, x_n .

Proof of lemma 7. $\mathbf{0} \notin \text{conv}(\{x_1, \dots, x_n\})$ if and only if the system

$$\begin{pmatrix} 1 & \cdots & 1 \\ x_1 & \cdots & x_n \end{pmatrix} \lambda = \begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix} \\ \lambda \geq 0$$

has no solutions. By Farkas' lemma (for example, theorem 6 in Border 2019), this system does not

have a solution if and only if there exists a point $\tilde{p} = (p_0, p) \in \mathbb{R}^{K+1}$ such that

$$\begin{aligned}\tilde{p}' \begin{pmatrix} 1 & \cdots & 1 \\ x_1 & \cdots & x_n \end{pmatrix} &\geq 0 \\ \tilde{p}' \begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix} &< 0.\end{aligned}$$

That is,

$$\begin{aligned}(p_0 \ p') \begin{pmatrix} 1 & \cdots & 1 \\ x_1 & \cdots & x_n \end{pmatrix} &\geq 0 \\ (p_0 \ p') \begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix} &< 0\end{aligned}$$

which is equivalent to

$$\begin{aligned}p_0 + p'x_i &\geq 0 \quad \text{for all } i = 1, \dots, n \\ p_0 &< 0.\end{aligned}\tag{25}$$

For the \Rightarrow direction, this immediately implies that $p'x_i > 0$ for all $i = 1, \dots, n$. For the \Leftarrow direction, we take $p'x_i > 0$ for all $i = 1, \dots, n$ as our assumption. We can then always pick a $p_0 < 0$ sufficiently close to zero that the system (25) holds. \square

Lemma 8. Suppose A1' holds. Suppose $L = K + 1$. Suppose FAS* is not a singleton. Then $\psi_\ell - \pi'_\ell b = w_\ell(b)(\psi_\ell - \pi'_\ell \beta_{-\ell}^{2\text{SLS}})$ for all $\ell = 1, \dots, L$.

Proof of lemma 8. For $\ell' \in \{1, \dots, L\}$ we have

$$\begin{aligned}\psi_{\ell'} - \pi'_{\ell'} b &= \psi_{\ell'} - \pi'_{\ell'} \left(\sum_{\ell=1}^L w_\ell(b) \beta_{-\ell}^{2\text{SLS}} \right) \\ &= \psi_{\ell'} - \pi'_{\ell'} \left(\sum_{\ell \neq \ell'} w_\ell(b) \beta_{-\ell}^{2\text{SLS}} + w_{\ell'}(b) \beta_{-\ell'}^{2\text{SLS}} \right) \\ &= \psi_{\ell'} - \pi'_{\ell'} \left(\sum_{\ell \neq \ell'} w_\ell(b) \beta_{-\ell}^{2\text{SLS}} + \left(1 - \sum_{\ell \neq \ell'} w_\ell(b) \right) \beta_{-\ell'}^{2\text{SLS}} \right) \\ &= \psi_{\ell'} - \sum_{\ell \neq \ell'} w_\ell(b) \pi'_{\ell'} \beta_{-\ell}^{2\text{SLS}} - \left(1 - \sum_{\ell \neq \ell'} w_\ell(b) \right) \pi'_{\ell'} \beta_{-\ell'}^{2\text{SLS}} \\ &= \psi_{\ell'} - \sum_{\ell \neq \ell'} w_\ell(b) \psi_{\ell'} - \left(1 - \sum_{\ell \neq \ell'} w_\ell(b) \right) \pi'_{\ell'} \beta_{-\ell'}^{2\text{SLS}} \\ &= \left(1 - \sum_{\ell \neq \ell'} w_\ell(b) \right) \psi_{\ell'} - \left(1 - \sum_{\ell \neq \ell'} w_\ell(b) \right) \pi'_{\ell'} \beta_{-\ell'}^{2\text{SLS}} \\ &= \left(1 - \sum_{\ell \neq \ell'} w_\ell(b) \right) (\psi_{\ell'} - \pi'_{\ell'} \beta_{-\ell'}^{2\text{SLS}})\end{aligned}$$

$$= w_{\ell'}(b)(\psi_{\ell'} - \pi'_{\ell'}\beta_{-\ell'}^{2\text{SLS}}).$$

In the first line we used the representation $b = \sum_{\ell=1}^L w_{\ell}(b)\beta_{-\ell}^{2\text{SLS}}$, by lemma 4. The last equality follows since the weights sum to one. \square

Lemma 9. Suppose A1' holds. Suppose $L = K + 1$. Suppose FAS^* is not a singleton. Without loss of generality, write

$$\text{FAS}^* = \bigcap_{\ell=1}^L \{b \in \mathbb{R}^K : \psi_{\ell} - \pi'_{\ell}b \leq 0\}.$$

(See equation (24) in the proof of lemma 5, and the surrounding discussion.) Then there are no $b \in \mathbb{R}^K$ such that $\psi_{\ell} - \pi'_{\ell}b \geq 0$ for all $\ell = 1, \dots, L$.

Proof of lemma 9. Let

$$\mathcal{E} = \bigcap_{\ell=1}^L \{\tilde{b} \in \mathbb{R}^K : \psi_{\ell} - \pi'_{\ell}\tilde{b} \geq 0\}.$$

Suppose by way of contradiction that there is some element $b \in \mathcal{E}$. Write

$$b = \sum_{\ell=1}^L w_{\ell}(b)\beta_{-\ell}^{2\text{SLS}}$$

where $\sum_{\ell=1}^L w_{\ell}(b) = 1$. By lemma 8 (which we can use since $L = K + 1$ and FAS^* is not a singleton),

$$\psi_{\ell} - \pi'_{\ell}b = w_{\ell}(b)(\psi_{\ell} - \pi'_{\ell}\beta_{-\ell}^{2\text{SLS}}).$$

This term is ≥ 0 since $b \in \mathcal{E}$. Since $\beta_{-\ell}^{2\text{SLS}} \in \text{FAS}^*$ and by lemma 3 we have $\psi_{\ell} - \pi'_{\ell}\beta_{-\ell}^{2\text{SLS}} < 0$. This implies $w_{\ell}(b) \leq 0$. This must hold for all ℓ . That is a contradiction to $\sum_{\ell=1}^L w_{\ell}(b) = 1$. Hence $\mathcal{E} = \emptyset$. \square

Lemma 10. Suppose A1', A2, and A3 hold. Suppose $L = K + 1$. Then there exists a normalization of the hyperplanes $\{b \in \mathbb{R}^K : \pi'_{\ell}b = \psi_{\ell}\}$ with $\psi_{\ell} \geq 0$ and such that

$$\mathbf{0} \in \text{conv}(\{\pi_{\ell} : \ell = 1, \dots, L\}) \quad \Rightarrow \quad \mathbf{0} \in \text{conv}(\{\beta_{-\ell}^{2\text{SLS}} : \ell = 1, \dots, L\}).$$

Proof of lemma 10. We prove this by contrapositive. So suppose $\mathbf{0} \notin \text{conv}(\{\beta_{-\ell}^{2\text{SLS}} : \ell = 1, \dots, L\})$. By $L = K + 1$ and lemma 3, there are two cases to consider.

1. Suppose $\text{FAS}^* = \text{conv}(\{\beta_{-\ell}^{2\text{SLS}} : \ell = 1, \dots, L\})$ is a singleton. Let b^* denote this common value. By definition of $\beta_{-\ell}^{2\text{SLS}}$, $\pi'_{\ell}\beta_{-\ell}^{2\text{SLS}} = \psi_{\ell}$ for all ℓ . Thus $\pi'_{\ell}b^* = \psi_{\ell}$ for all ℓ . Since $\mathbf{0}$ is not in the convex hull of the $\beta_{-\ell}^{2\text{SLS}}$ points by assumption, we have $b^* \neq \mathbf{0}$. Recall also that we normalized $\psi_{\ell} \geq 0$ for all ℓ . So we have $\pi'_{\ell}b^* \geq 0$ for all ℓ .

Let \mathcal{L}_0 denote the indices ℓ such that $\psi_{\ell} = 0$. Note that $\pi'_{\ell}b^* = 0$ for $\ell \in \mathcal{L}_0$. The set

$$\bigcup_{\ell \in \mathcal{L}_0} \{b \in \mathbb{R}^K : \pi'_{\ell}b = 0\}$$

does not equal \mathbb{R}^K . This follows since it is a finite union of $K - 1$ dimensional hyperplanes. So there is a point $v \in \mathbb{R}^K$ which is not in this set. That is, there exists a vector $v \in \mathbb{R}^K$ such that $\pi'_{\ell}v \neq 0$ for all $\ell \in \mathcal{L}_0$.

Note that hyperplanes corresponding to $\ell \in \mathcal{L}_0$ are invariant to sign normalizations of π_ℓ since

$$\{b \in \mathbb{R}^K : \pi'_\ell b = 0\} = \{b \in \mathbb{R}^K : (-\pi_\ell)'b = 0\}.$$

The element v is independent of the normalization since if $\pi'_\ell v \neq 0$, then $(-\pi_\ell)'v \neq 0$ as well. Let

$$b^*(t) = b^* + tv$$

for $t \geq 0$. At $t = 0$ we have

$$\pi'_\ell b^*(0) = \psi_\ell > 0$$

for all $\ell \in \mathcal{L} \setminus \mathcal{L}_0$. By continuity, for small enough $t > 0$ we also have

$$\pi'_\ell b^*(t) = \pi'_\ell b^* + t\pi'_\ell v > 0$$

for all $\ell \in \mathcal{L} \setminus \mathcal{L}_0$. Denote this point by $\bar{t} > 0$. Next, for all $\ell \in \mathcal{L}_0$ we have

$$\begin{aligned} \pi'_\ell b^*(\bar{t}) &= \pi'_\ell b^* + \bar{t}\pi'_\ell v \\ &= \bar{t}\pi'_\ell v \\ &\neq 0. \end{aligned}$$

Therefore, given that the sign of π_ℓ for $\ell \in \mathcal{L}_0$ can be renormalized, we can normalize π_ℓ for $\ell \in \mathcal{L}_0$ such that $\pi'_\ell b^*(\bar{t}) > 0$ for all $\ell \in \mathcal{L}_0$. So now we have found a normalization and a point $b^*(\bar{t})$ such that $\pi'_\ell b^*(\bar{t}) > 0$ for all ℓ . Lemma 7 thus implies that $\mathbf{0} \notin \text{conv}(\{\pi_\ell : \ell = 1, \dots, L\})$.

2. Suppose $\text{conv}(\{\beta_{-\ell}^{\text{2SLS}} : \ell = 1, \dots, L\})$ is not a singleton. By A1.2', the equation

$$\begin{pmatrix} 1 & \dots & 1 \\ \pi_1 & \dots & \pi_L \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_L \end{pmatrix} = \begin{pmatrix} 1 \\ \mathbf{0}_K \end{pmatrix} \quad (26)$$

has a unique solution $\lambda \in \mathbb{R}^L$. We want to show that this solution does *not* satisfy $\lambda \geq 0$.

Recall that we can write $\text{FAS}^* = \text{conv}(\{\beta_{-\ell}^{\text{2SLS}} : \ell = 1, \dots, L\})$ as an intersection of halfspaces $\{\tilde{b} \in \mathbb{R}^K : \pi'_\ell \tilde{b} \leq \psi_\ell\}$. Without loss of generality, write

$$\text{FAS}^* = \left(\bigcap_{\ell \in \mathcal{L}_1} \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \leq 0\} \right) \cap \left(\bigcap_{\ell \in \mathcal{L}_2} \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq 0\} \right)$$

for some partition $\mathcal{L}_1 \cup \mathcal{L}_2 = \mathcal{L}$. Since we have normalized $\psi_\ell \geq 0$, we *cannot* have $\mathcal{L}_1 = \emptyset$ because this would imply that $\mathbf{0} \in \text{FAS}^*$. So this implies that there is some index k_1 such that

$$\text{conv}(\{\beta_{-\ell}^{\text{2SLS}} : \ell = 1, \dots, L\}) \subseteq \{b \in \mathbb{R}^K : \pi'_{k_1} b \geq \psi_{k_1}\}.$$

However, we also cannot have $\mathcal{L}_1 = \{1, \dots, L\}$.

- If $\mathcal{L}_1 = \{1, \dots, L\}$ then

$$\text{FAS}^* = \bigcap_{\ell \in \{1, \dots, L\}} \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \leq 0\}.$$

By lemma 9, this implies that

$$\bigcap_{\ell \in \{1, \dots, L\}} \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq 0\}$$

is empty. This is a contradiction since

$$\mathbf{0} \in \bigcap_{\ell \in \{1, \dots, L\}} \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell \tilde{b} \geq 0\}$$

by the normalization $\psi_\ell \geq 0$.

Hence there is some index $k_2 \in \mathcal{L}_2$. That is, there is a k_2 such that

$$\{b \in \mathbb{R}^K : \pi'_{k_2} b \leq \psi_{k_2}\} \supseteq \text{conv}(\{\beta_{-\ell}^{2\text{SLS}} : \ell = 1, \dots, L\}).$$

Now, from the last row of equation (26) we have

$$\sum_{\ell=1}^L \lambda_\ell \pi_\ell = \mathbf{0}_K.$$

Multiply this equation by $(\beta_{-k}^{2\text{SLS}})'$ to get

$$(\beta_{-k}^{2\text{SLS}})' \sum_{\ell=1}^L \pi_\ell \lambda_\ell = 0.$$

Separate the $\ell = k$ from the sum to get

$$\sum_{\ell \neq k} \lambda_\ell \pi'_\ell \beta_{-k}^{2\text{SLS}} + \lambda_k \pi'_k \beta_{-k}^{2\text{SLS}} = 0.$$

Recall that, by the definition of these 2SLS estimands, $\pi'_k \beta_{-\ell}^{2\text{SLS}} = \psi_k$ for $k \neq \ell$. Hence

$$\sum_{\ell \neq k} \lambda_\ell \psi_k + \lambda_k \pi'_k \beta_{-k}^{2\text{SLS}} = 0. \quad (27)$$

Consider this equation for our two different indices $k_1, k_2 \in \{1, \dots, L\}$,

$$\sum_{\ell \neq k_1} \lambda_\ell \psi_{k_1} + \lambda_{k_1} \pi'_{k_1} \beta_{-k_1}^{2\text{SLS}} = 0 \quad \text{and} \quad \sum_{\ell \neq k_2} \lambda_\ell \psi_{k_2} + \lambda_{k_2} \pi'_{k_2} \beta_{-k_2}^{2\text{SLS}} = 0.$$

Subtract the left equation from the right to get

$$0 = \lambda_{k_2} (\pi'_{k_2} \beta_{-k_2}^{2\text{SLS}} - \psi_{k_2}) + \lambda_{k_1} (\psi_{k_1} - \pi'_{k_1} \beta_{-k_1}^{2\text{SLS}}).$$

As mentioned earlier, lemma 3 implies that $\pi'_{k_2} \beta_{-k_2}^{2\text{SLS}} - \psi_{k_2} < 0$ and $\psi_{k_1} - \pi'_{k_1} \beta_{-k_1}^{2\text{SLS}} < 0$. This equation implies that λ_{k_2} and λ_{k_1} cannot both have the same signs. If one of them is > 0 then the other must be < 0 . In this case, we're done. Here we use the fact that λ is the *unique* solution to equation (26).

So suppose both of them are zero. Let $k_3 \in \{1, \dots, L\} \setminus \{k_1, k_2\}$. Then by similar derivations

we have

$$0 = \lambda_{k_2}(\pi'_{k_2}\beta_{-k_2}^{2\text{SLS}} - \psi_{k_2}) + \lambda_{k_3}(\psi_{k_1} - \pi'_{k_3}\beta_{-k_3}^{2\text{SLS}}).$$

We know that $\lambda_{k_2} = 0$ already and thus

$$0 = \lambda_{k_3}(\psi_{k_1} - \pi'_{k_3}\beta_{-k_3}^{2\text{SLS}}).$$

We know that $\psi_{k_1} - \pi'_{k_3}\beta_{-k_3}^{2\text{SLS}} \neq 0$ and thus we must have $\lambda_{k_3} = 0$. Continuing in this way we obtain $\lambda = \mathbf{0}$. This contradicts $\sum_{\ell=1}^L \lambda_\ell = 1$. Thus λ cannot be nonnegative. So equation (26) then implies that $\mathbf{0} \notin \text{conv}(\{\pi_\ell : \ell = 1, \dots, L\})$. □

Lemma 11. Suppose A1', A2, and A3 hold. Suppose $L \geq K + 1$. Let $b \notin \mathcal{P}$. Then there exists a $\delta' < \delta(b)$ such that $\mathcal{B}(\delta') \neq \emptyset$.

Figures 13 and 14 illustrate the geometric intuition for this lemma.

Proof of lemma 11. Without loss of generality, suppose $b = \mathbf{0}$. This follows since we can simply translate our coordinate system so that the origin is at b . Put differently, we map all $x \in \mathbb{R}^K$ to $x - b$. In particular, the hyperplanes

$$B_\ell(0) = \{\tilde{b} \in \mathbb{R}^K : \psi_\ell = \pi'_\ell \tilde{b}\}$$

get mapped to

$$B_\ell(0) - b = \{\tilde{b} \in \mathbb{R}^K : \psi_\ell - \pi'_\ell b = \pi'_\ell \tilde{b}\}.$$

Thus we let $b = \mathbf{0}$. Throughout this proof, pick one of the normalizations given by lemma 10. In these normalizations we have $\psi_\ell \geq 0$ for all ℓ . Next, there are two cases to consider.

1. Suppose $\delta_\ell(b) > 0$ for all ℓ . Since $b \notin \mathcal{P}$, $b \notin \mathcal{P}_\mathcal{L} = \text{conv}(\{\beta_{\mathcal{L} \setminus \{l\}}^{2\text{SLS}} : l \in \mathcal{L}\})$ for any \mathcal{L} with $|\mathcal{L}| = K + 1$. Since $b = \mathbf{0}$, lemma 10 implies that $\mathbf{0} \notin \text{conv}(\{\pi_\ell : \ell \in \mathcal{L}\})$. This holds for any set \mathcal{L} such that $|\mathcal{L}| = K + 1$. This implies that $\mathbf{0} \notin \text{conv}(\{\pi_\ell : \ell = 1, \dots, L\})$.
 - If $\mathbf{0} \in \text{conv}(\{\pi_\ell : \ell = 1, \dots, L\})$, then Caratheodory's theorem (e.g., chapter 17 of Rockafellar 1970) implies that $\mathbf{0}$ is in the convex hull of a $(K + 1)$ -element subset of $\{\pi_\ell : \ell = 1, \dots, L\}$. That is, $\mathbf{0} \in \text{conv}(\{\pi_\ell : \ell \in \mathcal{L}\})$ for some \mathcal{L} with $|\mathcal{L}| = K + 1$. This is a contradiction.

Since $\mathbf{0} \notin \text{conv}(\{\pi_\ell : \ell = 1, \dots, L\})$, lemma 7 implies that there exists a vector \bar{b} such that $\bar{b}'\pi_\ell > 0$ for all $\ell = 1, \dots, L$. Define

$$\begin{aligned} b(\varepsilon) &= b + \varepsilon \bar{b} \\ &= \mathbf{0} + \varepsilon \bar{b}. \end{aligned}$$

This is a point that starts at b and moves in the direction of \bar{b} . Recall that we normalized $\psi_\ell \geq 0$ for all ℓ . Moreover, note that we must actually have $\psi_\ell > 0$ since $0 < \delta_\ell(\mathbf{0}) = |\psi_\ell|$. Thus, since $\psi_\ell > 0$ and $\pi'_\ell \bar{b} > 0$ for all ℓ , there exists an $\bar{\varepsilon} > 0$ such that

$$\psi_\ell - \bar{\varepsilon} \pi'_\ell \bar{b} > 0$$

for all ℓ . This implies that

$$\begin{aligned}
|\psi_\ell - \pi'_\ell b(\bar{\varepsilon})| &= |\psi_\ell - \pi'_\ell(b + \bar{\varepsilon}\bar{b})| \\
&= |(\psi_\ell - \pi'_\ell b) - \bar{\varepsilon}\pi'_\ell\bar{b}| \\
&= (\psi_\ell - \pi'_\ell b) - \bar{\varepsilon}\pi'_\ell\bar{b} \\
&= \delta_\ell(b) - \bar{\varepsilon}\pi'_\ell\bar{b} \\
&< \delta_\ell(b).
\end{aligned}$$

The third line uses $b = \mathbf{0}$ and our specific choice of $\bar{\varepsilon}$ to ensure that this term is positive, and hence we can drop the absolute value. The fourth line uses $\psi_\ell \geq 0$ and $b = \mathbf{0}$. The fifth line uses $\pi'_\ell\bar{b} > 0$ and $\bar{\varepsilon} > 0$.

Let

$$\delta'_\ell = |\psi_\ell - \pi'_\ell b(\bar{\varepsilon})|.$$

We have shown that $\delta'_\ell < \delta_\ell(b)$. Finally, by our characterization of $\mathcal{B}(\cdot)$ and the definition of δ' , we have $b(\bar{\varepsilon}) \in \mathcal{B}(\delta')$. Hence $\mathcal{B}(\delta') \neq \emptyset$.

2. Suppose $\delta_\ell(b) = 0$ for some ℓ 's. There can be at most $K - 1$ such indices, since otherwise we would have $b \in \mathcal{P}$. Let \mathcal{L}_0 denote the set of these indices. As in the proof of lemma 10, we must have $\psi_\ell = 0$ for $\ell \in \mathcal{L}_0$. In this case, consider the subspace

$$\{\tilde{b} \in \mathbb{R}^K : 0 = \pi'_\ell\tilde{b}, \ell \in \mathcal{L}_0\}.$$

This is a linear subspace of dimension at least 1 (by assumption that $\delta_\ell(b) = 0$ for some ℓ 's) and at most $K - 1$ (as noted earlier). Within this subspace, we can look at the remaining indices $\{1, \dots, L\} \setminus \mathcal{L}_0$. We have $\delta_\ell(b) > 0$ for all of these indices. By restricting attention to this subspace we can thus immediately apply the analysis of case 1. □

For the next two lemmas, let

$$\text{FF}^{\text{guess}} = \{\delta \in \mathbb{R}_{\geq 0}^L : \delta_\ell = |\psi_\ell - \pi'_\ell b|, \ell = 1, \dots, L, b \in \mathcal{P}\}$$

and let FF denote the true falsification frontier.

Lemma 12. Suppose A1', A2, and A3 hold. Suppose $L \geq K + 1$. Then $\text{FF}^{\text{guess}} \subseteq \text{FF}$.

Proof of lemma 12. Recall the definition $\delta_\ell(b) = |\psi_\ell - \pi'_\ell b|$. Let $\delta \in \text{FF}^{\text{guess}}$. Then, by definition, there is a $b \in \mathcal{P}$ such that $\delta_\ell(b) = \delta_\ell$ for all ℓ . Thus $\mathcal{B}(\delta) = \{b\}$ by lemma 6.

Let $\delta' < \delta(b)$. Then there is some index ℓ such that $0 < \delta'_\ell < \delta_\ell(b)$. So $\delta_\ell(b) \notin [-\delta'_\ell, \delta'_\ell]$. That is,

$$\psi_\ell - \pi'_\ell b \notin [-\delta'_\ell, \delta'_\ell].$$

So $b \notin B_\ell(\delta')$. This implies that $b \notin \mathcal{B}(\delta')$. But since $\mathcal{B}(\delta') \subseteq \mathcal{B}(\delta) = \{b\}$, we must have $\mathcal{B}(\delta') = \emptyset$. Hence, by the definition of the falsification frontier, $\text{FF}^{\text{guess}} \subseteq \text{FF}$. □

Lemma 13. Suppose A1', A2, and A3 hold. Suppose $L \geq K + 1$. Then $\text{FF}^{\text{guess}} \supseteq \text{FF}$.

Proof of lemma 13. We will show the contrapositive: $\delta \notin \text{FF}^{\text{guess}}$ implies $\delta \notin \text{FF}$. So let $\delta \notin \text{FF}^{\text{guess}}$. There are three cases to consider.

1. Suppose δ is such that $\mathcal{B}(\delta) = \emptyset$.
 - In this case the result follows immediately since $\mathcal{B}(\delta) = \emptyset$ implies $\delta \notin \text{FF}$ by definition.
2. Suppose δ is such that $\mathcal{B}(\delta)$ contains an element $b \notin \mathcal{P}$.
 - By lemma 11 there exists a $\delta' < \delta(b)$ with $\mathcal{B}(\delta') \neq \emptyset$. Moreover, $\delta(b) \leq \delta$ by the characterization of $\mathcal{B}(\delta)$ in theorem 1. Thus $\delta \notin \text{FF}$ by the definition of the falsification frontier.
3. Suppose δ is such that $\mathcal{B}(\delta) \neq \emptyset$ and $\mathcal{B}(\delta) \subseteq \mathcal{P}$.
 - We use two observations:
 - (a) Since $\delta \notin \text{FF}^{\text{guess}}$ we must have $\delta \neq \delta(b)$ for all $b \in \mathcal{P}$. In particular, $\delta \neq \delta(b)$ for all $b \in \mathcal{B}(\delta)$, since we are considering the case where $\mathcal{B}(\delta) \subseteq \mathcal{P}$.
 - (b) Let $b \in \mathcal{B}(\delta)$. By the characterization of $\mathcal{B}(\delta)$ in theorem 1 (see equation (22)), this implies that $|\psi_\ell - \pi'_\ell b| \leq \delta_\ell$. That is, $\delta(b) \leq \delta$.

Let b' be any element of $\mathcal{B}(\delta)$. This exists by assumption. Let $\delta' = \delta(b')$. The previous two observations imply that $\delta' \neq \delta$ and $\delta' \leq \delta$. This implies $\delta' < \delta$. Moreover, we have $\mathcal{B}(\delta') = \mathcal{B}(\delta(b')) = \{b'\} \neq \emptyset$ by lemma 6. Thus $\delta \notin \text{FF}$, by definition of the falsification frontier.

All values of δ must fall in one of these three cases. Therefore $\text{FF}^{\text{guess}} \supseteq \text{FF}$. □

Proof of proposition 5. This follows directly from lemmas 12 and 13. □

Proof of proposition 4. This is a special case of proposition 5. □

Proof of theorem 4. We have

$$\begin{aligned}
\bigcup_{\delta \in \text{FF}} \mathcal{B}(\delta) &= \bigcup_{b \in \mathcal{P}} \mathcal{B}(\delta(b)) \\
&= \bigcup_{b \in \mathcal{P}} \{b\} \\
&= \mathcal{P}.
\end{aligned}$$

The first line follows by proposition 5. The second line follows by lemma 6. □

Proof of theorem 3. This is a special case of theorem 4. In particular, note that $\mathcal{P} = \text{FAS}^*$ in this case. □

To prove corollary 4, we use the following definition: Let P be a $K \times K$ idempotent matrix. Define the linear operator $p : \mathbb{R}^K \rightarrow \mathbb{R}^K$ by $p(a) = Pa$. p is called a *projection*. For $A \subseteq \mathbb{R}^K$, define the projection of the set A as

$$\text{proj}(A) = \{p(a) \in \mathbb{R}^K : a \in A\}.$$

We use the following lemma.

Lemma 14. $\text{proj}(\text{conv}(A)) = \text{conv}(\text{proj}(A))$.

Proof of lemma 14. Step 1 (\supseteq). Since $A \subseteq \text{conv}(A)$, $\text{proj}(A) \subseteq \text{proj}(\text{conv}(A))$. Thus

$$\text{conv}(\text{proj}(A)) \subseteq \text{conv}(\text{proj}(\text{conv}(A))).$$

Since $p(\cdot)$ is a linear operator, $\text{proj}(\text{conv}(A))$ is convex. Hence

$$\text{conv}(\text{proj}(\text{conv}(A))) = \text{proj}(\text{conv}(A)).$$

This with our previous result shows that $\text{conv}(\text{proj}(A)) \subseteq \text{proj}(\text{conv}(A))$.

Step 2 (\subseteq). Let $x \in \text{proj}(\text{conv}(A))$. Then $x = p(y)$ for some $y \in \text{conv}(A)$. Since $y \in \text{conv}(A)$, we can write

$$y = \sum_{i=1}^n \theta_i y_i$$

where $\sum_{i=1}^n \theta_i = 1, \theta_i \geq 0$, and $y_i \in A$ for $i = 1, \dots, n$, where n is a finite integer. Since $p(\cdot)$ is a linear operator,

$$p(y) = \sum_{i=1}^n \theta_i p(y_i).$$

Thus we've written x as a convex combination of elements $p(y_i) \in \text{proj}(A)$. Hence $x \in \text{conv}(\text{proj}(A))$. Thus we've shown that $\text{proj}(\text{conv}(A)) \subseteq \text{conv}(\text{proj}(A))$. \square

Proof of corollary 4. The first part of this result states that $\mathcal{P} \subseteq \text{FAS}^*$. To see this, let $b \in \mathcal{P}$. Then $b \in \mathcal{P}_{\mathcal{L}}$ for some set of indices \mathcal{L} with $|\mathcal{L}| = K + 1$. But if b is a convex combination of $\{\beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} : \ell \in \mathcal{L}\}$, then it is also a convex combination of the larger set of elements $\{\beta_{\mathcal{S}}^{2\text{SLS}} : \mathcal{S} \subseteq \{1, \dots, L\}, |\mathcal{S}| = K\}$. Hence $b \in \text{FAS}^*$.

To prove the second part, consider projections defined by

$$P = \begin{pmatrix} \alpha' \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

where $\alpha \in \mathbb{R}^K$. For any set $A \subseteq \mathbb{R}^K$, let $[A]_1 = \{a_1 \in \mathbb{R} : a = (a_1, \dots, a_K) \in \mathbb{R}^K\}$. Since P maps the $2, \dots, K$ components of any vector $a \in \mathbb{R}^K$ to zero, it suffices to show that

$$[\text{proj}(\mathcal{P})]_1 = [\text{proj}(\text{FAS}^*)]_1$$

and

$$[\text{proj}(\text{FAS}^*)]_1 = \left[\min_{\mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}|=K} \alpha' \beta_{\mathcal{L}}^{2\text{SLS}}, \max_{\mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}|=K} \alpha' \beta_{\mathcal{L}}^{2\text{SLS}} \right]. \quad (28)$$

We have

$$\begin{aligned} \text{proj}(\text{FAS}^*) &= \text{proj}(\text{conv}(\mathcal{P})) \\ &= \text{proj}(\text{conv}(\{\beta_{\mathcal{L}}^{2\text{SLS}} : \mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}| = K\})) \\ &= \text{conv}(\text{proj}(\{\beta_{\mathcal{L}}^{2\text{SLS}} : \mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}| = K\})) \\ &= \text{conv}(\{P\beta_{\mathcal{L}}^{2\text{SLS}} : \mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}| = K\}). \end{aligned}$$

The third line follows from lemma 14. The fourth by the definition of the projection of a set. Using the specific form of the projection P now gives equation (28).

Similarly,

$$\begin{aligned} \text{proj}(\mathcal{P}) &= \text{proj} \left(\bigcup_{\mathcal{L} \subseteq \{1, \dots, L\}: |\mathcal{L}|=K+1} \text{conv}(\{\beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} : \ell \in \mathcal{L}\}) \right) \\ &= \bigcup_{\mathcal{L} \subseteq \{1, \dots, L\}: |\mathcal{L}|=K+1} \text{proj} \left(\text{conv}(\{\beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} : \ell \in \mathcal{L}\}) \right) \\ &= \bigcup_{\mathcal{L} \subseteq \{1, \dots, L\}: |\mathcal{L}|=K+1} \text{conv} \left(\text{proj}(\{\beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} : \ell \in \mathcal{L}\}) \right). \end{aligned}$$

The first line follows by definition of \mathcal{P} . The second since projections and unions commute. The third from lemma 14. Hence

$$[\text{proj}(\mathcal{P})]_1 = \bigcup_{\mathcal{L} \subseteq \{1, \dots, L\}: |\mathcal{L}|=K+1} \left[\min_{\ell \in \mathcal{L}} \alpha' \beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}}, \max_{\ell \in \mathcal{L}} \alpha' \beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} \right]$$

Thus we see that the first component of $\text{proj}(\mathcal{P})$ is a union of closed intervals.

If \mathcal{L} and \mathcal{L}' differ by at most one element, then the intervals

$$\left[\min_{\ell \in \mathcal{L}} \alpha' \beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}}, \max_{\ell \in \mathcal{L}} \alpha' \beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} \right]$$

and

$$\left[\min_{\ell \in \mathcal{L}'} \alpha' \beta_{\mathcal{L}' \setminus \{\ell\}}^{2\text{SLS}}, \max_{\ell \in \mathcal{L}'} \alpha' \beta_{\mathcal{L}' \setminus \{\ell\}}^{2\text{SLS}} \right]$$

will overlap (have non-empty intersection). Their union is therefore also a closed interval. Because we take the union over all $\mathcal{L} \subseteq \{1, \dots, L\} : |\mathcal{L}| = K + 1$, we can find a sequence $(\mathcal{L}_1, \dots, \mathcal{L}_N)$ such that

$$\left[\min_{\ell \in \mathcal{L}_n} \alpha' \beta_{\mathcal{L}_n \setminus \{\ell\}}^{2\text{SLS}}, \max_{\ell \in \mathcal{L}_n} \alpha' \beta_{\mathcal{L}_n \setminus \{\ell\}}^{2\text{SLS}} \right]$$

and

$$\left[\min_{\ell \in \mathcal{L}_{n+1}} \alpha' \beta_{\mathcal{L}_{n+1} \setminus \{\ell\}}^{2\text{SLS}}, \max_{\ell \in \mathcal{L}_{n+1}} \alpha' \beta_{\mathcal{L}_{n+1} \setminus \{\ell\}}^{2\text{SLS}} \right]$$

overlap and such that

$$\bigcup_{n=1}^N \mathcal{L}_n = \{1, \dots, L\}$$

Thus

$$\bigcup_{\mathcal{L} \subseteq \{1, \dots, L\}: |\mathcal{L}|=K+1} \left[\min_{\ell \in \mathcal{L}} \alpha' \beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}}, \max_{\ell \in \mathcal{L}} \alpha' \beta_{\mathcal{L} \setminus \{\ell\}}^{2\text{SLS}} \right] = \left[\min_{\mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}|=K} \alpha' \beta_{\mathcal{L}}^{2\text{SLS}}, \max_{\mathcal{L} \subseteq \{1, \dots, L\}, |\mathcal{L}|=K} \alpha' \beta_{\mathcal{L}}^{2\text{SLS}} \right].$$

Putting everything together yields $[\text{proj}(\mathcal{P})]_1 = [\text{proj}(\text{FAS}^*)]_1$ as desired. \square

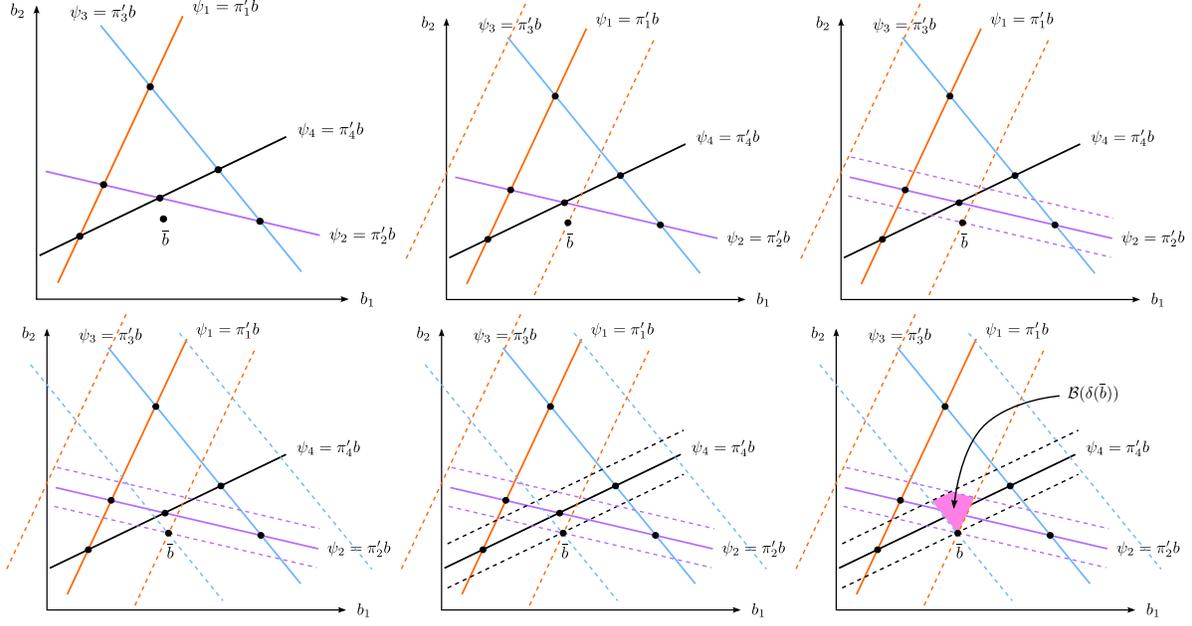


Figure 14: Geometric intuition for a key step in the proof of theorem 4. This figure is analogous to figure 13, except now we have $K = 2$ and $L = 4$. Moreover, here we pick \bar{b} to be outside of \mathcal{P} but inside FAS^* . Thus here we illustrate why FAS^* is not the falsification adaptive set when $L > K + 1$. The argument is the same as before: We relax each exclusion restriction just enough to include \bar{b} in the tubes. But when we do this their intersection is non-singleton. Thus we've relaxed exclusion too much. The point \bar{b} is not in the falsification adaptive set.

L Proofs for section 5

Proofs for section 5.1: Single binary instrument

Proof of theorem 5. This result follows by applying theorem 6 to the case with $J = 2$ and $L = 1$. Here we will show that the general expressions for $\mathcal{D}(c)$ and \mathcal{H}_x used in theorem 6 simplify to the ones used in theorem 5.

We start with $\mathcal{D}(c)$. From theorem 6, it is the set of $(a_0, a_1) \in [0, 1]^2$ such that the following eight inequalities hold

$$\begin{aligned}
(1 - p_Z)a_0 - \min\{1 - p_Z + c, 1\}(p_Z a_1 + (1 - p_Z)a_0) &\leq 0 \\
(1 - p_Z)a_0 - \max\{1 - p_Z - c, 0\}(p_Z a_1 + (1 - p_Z)a_0) &\geq 0 \\
(1 - p_Z)(1 - a_0) - \min\{1 - p_Z + c, 1\}[1 - (p_Z a_1 + (1 - p_Z)a_0)] &\leq 0 \\
(1 - p_Z)(1 - a_0) - \max\{1 - p_Z - c, 0\}[1 - (p_Z a_1 + (1 - p_Z)a_0)] &\geq 0 \\
p_Z a_1 - \min\{p_Z + c, 1\}(p_Z a_1 + (1 - p_Z)a_0) &\leq 0 \\
p_Z a_1 - \max\{p_Z - c, 0\}(p_Z a_1 + (1 - p_Z)a_0) &\geq 0 \\
p_Z(1 - a_1) - \min\{p_Z + c, 1\}[1 - (p_Z a_1 + (1 - p_Z)a_0)] &\leq 0 \\
p_Z(1 - a_1) - \max\{p_Z - c, 0\}[1 - (p_Z a_1 + (1 - p_Z)a_0)] &\geq 0.
\end{aligned}$$

Note that we indexed the elements of $\mathcal{D}(c)$ by $j \in \{0, 1\}$ rather than $j \in \{1, 2\}$. Also recall that $p_Z = \mathbb{P}(Z = 1)$. Using the function $k_z(c)$ and removing redundant equations, we can rearrange

these 8 inequalities to obtain

$$\begin{aligned} a_1 &\geq a_0 k_0(c) \\ a_0 &\geq a_1 k_1(c) \\ (1 - a_1) &\geq (1 - a_0) k_0(c) \\ (1 - a_0) &\geq (1 - a_1) k_1(c). \end{aligned}$$

This is the set of inequalities described in equation (14) and used in theorem 5.

Next consider \mathcal{H}_x . From theorem 6, we have

$$\begin{aligned} \mathcal{H}_x &= \{(a_0, a_1) \in [0, 1]^2 : \text{For some } (q_0, q_1) \in [0, 1]^2, \\ &\quad a_0 = \mathbb{P}(Y = 1, X = x \mid Z = 0) + \mathbb{P}(X = 1 - x \mid Z = 0)q_0, \\ &\quad a_1 = \mathbb{P}(Y = 1, X = x \mid Z = 1) + \mathbb{P}(X = 1 - x \mid Z = 1)q_1\}. \end{aligned}$$

This is simply

$$\begin{aligned} \mathcal{H}_x &= [\mathbb{P}(Y = 1, X = x \mid Z = 0), \mathbb{P}(Y = 1, X = x \mid Z = 0) + \mathbb{P}(X = 1 - x \mid Z = 0)] \\ &\quad \times [\mathbb{P}(Y = 1, X = x \mid Z = 1), \mathbb{P}(Y = 1, X = x \mid Z = 1) + \mathbb{P}(X = 1 - x \mid Z = 1)] \end{aligned}$$

as in theorem 5. □

Proof of proposition 6. Parts 1 and 2 are immediate corollaries of proposition 7. Consider part 3. By proposition 7, the correspondence ϕ is continuous on

$$(c^*, 1] = \{c \in [c^*, 1] : (\text{int}(\mathcal{D}(c)) \cap \text{int}(\mathcal{H}_x)) \neq \emptyset, x \in \{0, 1\}\}.$$

Here we show that ϕ is also continuous at $c = c^*$. Again by proposition 7, this correspondence is uhc at $c = c^*$. To show it is also lhc at $c = c^*$, let $U = [c^*, 1]$. U is an open neighborhood of c^* relative to the domain $[c^*, 1]$. Let E be any open set such that

$$(\Theta_0(c^*) \times \Theta_1(c^*)) \cap E \neq \emptyset.$$

Since the set $\Theta_0(c) \times \Theta_1(c)$ is weakly increasing in c (with respect to the set inclusion order \subseteq),

$$(\Theta_0(c^*) \times \Theta_1(c^*)) \subseteq (\Theta_0(c) \times \Theta_1(c))$$

for any $c \in U$. Hence

$$[(\Theta_0(c^*) \times \Theta_1(c^*)) \cap E] \subseteq [(\Theta_0(c) \times \Theta_1(c)) \cap E]$$

for any $c \in U$. Since the set on the left is not empty, the set on the right is not empty:

$$(\Theta_0(c) \times \Theta_1(c)) \cap E \neq \emptyset$$

for any $c \in U$. Therefore ϕ is continuous at $c = c^*$. □

Proof of corollary 5. This follows immediately from the more general result, corollary 6. □

Proofs for section 5.1: Multiple discrete instruments

Proof of theorem 6. There are two steps. First we show that $\Theta_0(c) \times \Theta_1(c)$ is an outer identified set, in the sense that it contains the true parameters. Then we show that it is sharp, in the sense that any element of this set is consistent with the data and the assumptions. Finally, the fact that the model is refuted if and only if this set is empty follows by the definition of a (sharp) identified set.

Step 1. We first show that the true conditional probabilities (p_0, p_1) are in $\Theta_0(c) \times \Theta_1(c)$. Fix $x \in \{0, 1\}$. By c_ℓ -dependence,

$$\begin{aligned}\mathbb{P}(Z_\ell = z_\ell^j | Y_x = 1) &\leq \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell, 1\} \\ \mathbb{P}(Z_\ell = z_\ell^j | Y_x = 1) &\geq \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c_\ell, 0\} \\ \mathbb{P}(Z_\ell = z_\ell^j | Y_x = 0) &\leq \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell, 1\} \\ \mathbb{P}(Z_\ell = z_\ell^j | Y_x = 0) &\geq \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c_\ell, 0\}\end{aligned}$$

for all $j \in \{1, \dots, J\}$. Consider the first inequality. Use Bayes' rule to rewrite the left hand side. This yields

$$\begin{aligned}\mathbb{P}(Y_x = 1 | Z_\ell = z_\ell^j)\mathbb{P}(Z_\ell = z_\ell^j) &\leq \min\{\mathbb{P}(Z = z_\ell^j) + c_\ell, 1\}\mathbb{P}(Y_x = 1) \\ &= \min\{\mathbb{P}(Z = z_\ell^j) + c_\ell, 1\} \sum_{k=1}^J \mathbb{P}(Y_x = 1 | Z_\ell = z_\ell^k)\mathbb{P}(Z_\ell = z_\ell^k).\end{aligned}$$

In the last line we applied the law of total probability. Hence

$$\mathbb{P}(Y_x = 1 | Z_\ell = z_\ell^j)\mathbb{P}(Z_\ell = z_\ell^j) - \min\{\mathbb{P}(Z = z_\ell^j) + c_\ell, 1\} \sum_{k=1}^J \mathbb{P}(Y_x = 1 | Z_\ell = z_\ell^k)\mathbb{P}(Z_\ell = z_\ell^k) \leq 0.$$

Repeat these derivations for the other three inequalities. This yields the definition of $\mathcal{D}_\ell(c_\ell)$. Thus we have shown that the true conditional probabilities $\mathbb{P}(Y_x = 1 | Z_\ell = z_\ell^j)$ are in $\mathcal{D}_\ell(c_\ell)$ for each $\ell \in \{1, \dots, L\}$. p_x is simply the vector of these conditional probabilities. Thus we have shown $p_x \in \mathcal{D}(c)$.

Next we show that $p_x \in \mathcal{H}_x$. Element $(\ell - 1)J + j$ of p_x can be written as

$$\begin{aligned}[p_x]_{(\ell-1)J+j} &= \mathbb{P}(Y_x = 1 | Z_\ell = z_\ell^j) \\ &= \mathbb{P}(Y_x = 1, X = x | Z_\ell = z_\ell^j) + \mathbb{P}(Y_x = 1, X = 1 - x | Z_\ell = z_\ell^j) \\ &= \mathbb{P}(Y = 1, X = x | Z_\ell = z_\ell^j) + \mathbb{P}(Y_x = 1, X = 1 - x | Z_\ell = z_\ell^j) \\ &= \mathbb{P}(Y = 1, X = x | Z_\ell = z_\ell^j) \\ &\quad + \sum_{k=1}^{J^{L-1}} \mathbb{P}(Y_x = 1 | X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k)\mathbb{P}(X = 1 - x, Z_{-\ell} = z_{-\ell}^k | Z_\ell = z_\ell^j).\end{aligned}$$

This is simply equation (16). Since

$$\mathbb{P}(Y_x = 1 | X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k) \in [0, 1],$$

these derivations imply that

$$p_x = \mathbf{b}_x + \mathbf{A}_x q_x$$

where $q_x \in [0, 1]^{J^L}$ is a vector with elements

$$\mathbb{P}(Y_x = 1 \mid X = 1 - x, Z_\ell = z_\ell^j, Z_{-\ell} = z_{-\ell}^k).$$

Thus $p_x \in \mathcal{H}_x$.

Thus we have shown that $p_x \in \mathcal{D}(c) \cap \mathcal{H}_x$. Since this is true for each $x \in \{0, 1\}$, we have $p \in \Theta_1(c) \times \Theta_0(c)$.

Step 2. Next we show sharpness. Let $(p_0, p_1) \in \Theta_1(c) \times \Theta_0(c)$. First, $p_x \in \mathcal{D}(c)$ implies

$$\mathbb{P}(Z = z_\ell^j \mid Y_x = y) \in [\max\{\mathbb{P}(Z_\ell = z_\ell^j) - c_\ell, 0\}, \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell, 1\}]$$

for all $\ell \in \{1, \dots, L\}$ and all $j \in \{1, \dots, J\}$. This follows by reversing the arguments at the beginning of step 1. Hence c_ℓ -dependence is satisfied for all $\ell \in \{1, \dots, L\}$.

Next, since $p_x \in \mathcal{H}_x$ there is a vector $q_x \in [0, 1]^{J^L}$ such that

$$p_x = \mathbf{b}_x + \mathbf{A}_x q_x.$$

This vector q_x consists of elements of the form

$$\mathbb{P}(Y_x = 1 \mid X = 1 - x, Z = z).$$

Take these elements and combine them with the observed probabilities

$$\mathbb{P}(Y = 1 \mid X = x, Z = z) = \mathbb{P}(Y_x = 1 \mid X = x, Z = z)$$

to get a distribution of $Y_x \mid (X, Z)$. Do this for both $x \in \{0, 1\}$. Combine these conditional marginal distributions into a joint distribution $(Y_0, Y_1) \mid (X, Z)$ using any copula, and finally combine these with the known marginals (X, Z) to get a joint distribution of (Y_0, Y_1, X, Z) . By construction, this joint distribution is consistent with c -dependence, the distribution of the observed data (Y, X, Z) , and yields the point (p_0, p_1) that we started with. \square

Lemma 15. Suppose B2' holds. Then the correspondence $\mathcal{D} : [0, 1]^L \rightrightarrows [0, 1]^{LJ}$ defined by $\mathcal{D}(c)$ in equation (18) is continuous.

Proof of lemma 15. To prove this result, we show that \mathcal{D}_ℓ is both upper hemicontinuous (uhc) and lower hemicontinuous (lhc).

Step 1: Upper hemicontinuity. Pick any $c_\ell \in [0, 1]$ and $a \in [0, 1]^J$. Consider a sequence $\{c_\ell^n\}$ converging to c_ℓ and a sequence $\{a^n\}$ converging to a such that $a^n \in \mathcal{D}(c^n)$ for all n . From the first inequality in equation (18),

$$\mathbb{P}(Z_\ell = z_\ell^j) a_j^n - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell^n, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k^n \leq 0.$$

Taking limits as $n \rightarrow \infty$ yields

$$\mathbb{P}(Z_\ell = z_\ell^j) a_j - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c_\ell, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k \leq 0.$$

Repeating this for the other three inequalities shows that $a \in \mathcal{D}_\ell(c_\ell)$. Thus \mathcal{D}_ℓ has a closed graph. Hence \mathcal{D}_ℓ is uhc by theorem 17.11 in Aliprantis and Border (2006), since the range $[0, 1]^J$ is closed.

Step 2: Lower hemicontinuity. We next show that \mathcal{D}_ℓ is lhc. We do this in two parts: First for $\bar{c} \in (0, 1)$ and second for $\bar{c} = 0$.

Part A. Let $\bar{c} \in (0, 1]$. We will show \mathcal{D}_ℓ is lhc at \bar{c} . Let E be an open set in $[0, 1]^J$ such that

$$\mathcal{D}_\ell(\bar{c}) \cap E \neq \emptyset.$$

By the definition of lhc, we must find an open set U (open relative to the domain $[0, 1]$) such that

1. $\bar{c} \in U$, and
2. $c \in U$ implies that $\mathcal{D}_\ell(c) \cap E \neq \emptyset$.

We consider two cases.

Case 1. Suppose there is an element $\bar{a} = (\bar{a}_1, \dots, \bar{a}_J) \in \mathcal{D}_\ell(\bar{c}) \cap E$ such that $\bar{a}_1 = \dots = \bar{a}_J$. In this case, we can simply let $U = [0, 1]$. Why? U is open (relative to $[0, 1]$). It contains \bar{c} since it has to be in $[0, 1]$. Moreover, we know that $\bar{a} \in E$. But since $\bar{a}_1 = \dots = \bar{a}_J$, we also know that $\bar{a} \in \mathcal{D}_\ell(0)$. (This is simply our baseline model where all these conditional probabilities equal each other, by statistical independence.) Hence $\bar{a} \in \mathcal{D}_\ell(c)$ for any $c \in [0, 1]$, since $\mathcal{D}_\ell(c)$ is weakly increasing in c . Thus $\mathcal{D}_\ell(c) \cap E \neq \emptyset$ for all $c \in [0, 1] = U$. This is the second requirement we needed of U .

Case 2. Suppose none of the elements in $\mathcal{D}_\ell(\bar{c}) \cap E$ have all components equal. Let \bar{a} be an arbitrary element of this intersection. Consider

$$a(\epsilon) = (1 - \epsilon)\bar{a} + \epsilon \frac{1}{2} \iota_J$$

where ι_J is a J -vector of ones. $\frac{1}{2} \iota_J \in \mathcal{D}_\ell(c)$ for all c since $\mathcal{D}_\ell(c)$ always contains the diagonal. This combined with convexity of the set $\mathcal{D}_\ell(c)$ implies that $a(\epsilon) \in \mathcal{D}_\ell(\bar{c})$ for all $\epsilon \in [0, 1]$. Finally, since E is open, there exists an $\epsilon^* \in [0, 1]$ such that $a(\epsilon^*) \in \mathcal{D}_\ell(\bar{c}) \cap E$. Let $a^* = a(\epsilon^*)$. Let

$$U = \{c \in [0, 1] : a^* \in \text{int}(\mathcal{D}_\ell(c))\}.$$

Equivalently,

$$U = \left\{ c \in [0, 1] : \text{For all } j = 1, \dots, J, \right. \\
\mathbb{P}(Z_\ell = z_\ell^j) a_j^* - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k^* < 0, \\
\mathbb{P}(Z_\ell = z_\ell^j) a_j^* - \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c, 0\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k^* > 0, \\
\mathbb{P}(Z_\ell = z_\ell^j) (1 - a_j^*) - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) (1 - a_k^*) < 0, \\
\left. \mathbb{P}(Z_\ell = z_\ell^j) (1 - a_j^*) - \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c, 0\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) (1 - a_k^*) > 0 \right\}.$$

This set U can be written as

$$U = \bigcap_{j=1}^J \left[\left\{ c \in [0, 1] : \mathbb{P}(Z_\ell = z_\ell^j) a_j^* - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k^* < 0 \right\} \right. \\
\cap \left\{ c \in [0, 1] : \mathbb{P}(Z_\ell = z_\ell^j) a_j^* - \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c, 0\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) a_k^* > 0 \right\} \\
\cap \left\{ c \in [0, 1] : \mathbb{P}(Z_\ell = z_\ell^j) (1 - a_j^*) - \min\{\mathbb{P}(Z_\ell = z_\ell^j) + c, 1\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) (1 - a_k^*) < 0 \right\} \\
\left. \cap \left\{ c \in [0, 1] : \mathbb{P}(Z_\ell = z_\ell^j) (1 - a_j^*) - \max\{\mathbb{P}(Z_\ell = z_\ell^j) - c, 0\} \sum_{k=1}^J \mathbb{P}(Z_\ell = z_\ell^k) (1 - a_k^*) > 0 \right\} \right].$$

By continuity in c of the inequality components, this is a finite intersection of open sets. Hence U is open. U is not empty since $\bar{c} \in U$, which follows since $a^* \in \text{int}(\mathcal{D}_\ell(\bar{c}))$.

Finally, we show that U satisfies the second property we need: Suppose $c \in U$. By definition of U , $a^* \in \text{int}(\mathcal{D}_\ell(c))$. Moreover, $a^* \in E$ by construction. Thus $a^* \in (\mathcal{D}_\ell(c) \cap E)$ for any $c \in U$. Hence

$$\mathcal{D}_\ell(c) \cap E \neq \emptyset$$

for any $c \in U$.

Part B. Let $\bar{c} = 0$. Let E be an open set in $[0, 1]^J$ such that $\mathcal{D}_\ell(0) \cap E \neq \emptyset$. Let $U = [0, 1]$. U is open (relative to $[0, 1]$). Moreover,

$$(\mathcal{D}_\ell(0) \cap E) \subseteq (\mathcal{D}_\ell(c) \cap E)$$

since \mathcal{D}_ℓ is weakly increasing. Since the set on the left is not empty, the set on the right is not empty:

$$\mathcal{D}_\ell(c) \cap E \neq \emptyset$$

for any $c \in U$.

Step 3: Take products. Finally, theorem 17.28 in Aliprantis and Border (2006) shows that the product of continuous correspondences is continuous. Hence the correspondence $\mathcal{D}(\cdot)$ defined by

$$\mathcal{D}(c) = \prod_{\ell=1}^L \mathcal{D}_\ell(c_\ell)$$

is continuous. □

Proof of proposition 7.

1. By lemma 15, the correspondence $\mathcal{D}(\cdot)$ is continuous. Moreover $\mathcal{D}(\cdot)$ is closed-valued since it is defined by a finite number of linear weak inequalities. The correspondence that maps c into \mathcal{H}_x is continuous since it is constant-valued. It is also closed- and compact-valued since $[0, 1]^{J^L}$ is compact, and \mathcal{H}_x is the image of $[0, 1]^{J^L}$ under a continuous (affine) mapping. Hence the set

$$\mathcal{C}_x = \{c \in [0, 1]^L : \mathcal{D}(c) \cap \mathcal{H}_x \neq \emptyset\}$$

is closed, by exercise 11.18(b) on page 58 of Border (1985). Finally, let

$$\mathcal{C} = \mathcal{C}_0 \cap \mathcal{C}_1.$$

This set is the intersection of two closed sets and hence is closed. It depends only on $\mathcal{D}(\cdot)$ and \mathcal{H}_x , which are point identified. Hence \mathcal{C} is point identified.

2. By definition of $\Theta_x(c)$,

$$\Theta_0(c) \times \Theta_1(c) = (\mathcal{D}(c) \times \mathcal{H}_0) \times (\mathcal{D}(c) \times \mathcal{H}_1).$$

Each of these sets is a closed, convex polytope. To see this, consider $\mathcal{D}(c)$ and \mathcal{H}_x separately.

- (a) $\mathcal{D}(c)$ is closed, as shown in part 1. It is convex since it is defined by a collection of linear inequalities.
- (b) \mathcal{H}_x is compact, as shown in part 1. Therefore it is closed and bounded. It is convex valued since it is an affine transformation of the convex set $[0, 1]^{J^L}$, and convexity is preserved by affine transformations.

Since both $\mathcal{D}(c)$ and \mathcal{H}_x are closed, convex, bounded by $[0, 1]^{LJ}$, and defined by linear inequalities, they are both polytopes. Finally, all three properties—closed, convex, polytope—are preserved by taking finite Cartesian products.

3. Next we show continuity of the correspondence $\phi : \mathcal{C} \rightrightarrows [0, 1]^{2LJ}$ defined by $\phi(c) = \Theta_0(c) \times \Theta_1(c)$. First consider the correspondence $\phi_x : \mathcal{C} \rightrightarrows [0, 1]^{LJ}$ defined by

$$\phi_x(c) = \Theta_x(c) = \mathcal{D}(c) \cap \mathcal{H}_x.$$

It is convex-valued and closed-valued since its values are the intersection of two closed and convex sets. It is uhc by theorem 17.25 in Aliprantis and Border (2006), since $\mathcal{D}(c)$ and \mathcal{H}_x are both uhc. It is lhc for all c such that $\text{int}(\mathcal{D}(c)) \cap \text{int}(\mathcal{H}_x) \neq \emptyset$, by applying theorem B in Lechicki and Spakowski (1985). Finally, ϕ is just the product of the correspondences ϕ_0 and ϕ_1 . The product of continuous correspondences is continuous by theorem 17.28 in Aliprantis and Border (2006).

□

Proof of corollary 6.

1. We can write the identified set for $(\mathbb{P}(Y_0 = 1), \mathbb{P}(Y_1 = 1))$ as

$$\begin{aligned} & \left\{ \left(\sum_{j=1}^J \sum_{\ell=1}^L a_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j), \sum_{j=1}^J \sum_{\ell=1}^L \tilde{a}_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j) \right) : \mathbf{a} \in \Theta_0(c), \tilde{\mathbf{a}} \in \Theta_1(c) \right\} \\ &= \left\{ \sum_{j=1}^J \sum_{\ell=1}^L a_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j) : \mathbf{a} \in \Theta_0(c) \right\} \times \left\{ \sum_{j=1}^J \sum_{\ell=1}^L \tilde{a}_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j) : \tilde{\mathbf{a}} \in \Theta_1(c) \right\}. \end{aligned}$$

By proposition 7, $\Theta_x(c)$ is closed, non-empty, and bounded for $x \in \{0, 1\}$. Therefore, suprema and infima of these sets are attained on $\Theta_x(c)$. Moreover, since $\Theta_x(c)$ is convex, the set

$$\left\{ \sum_{j=1}^J \sum_{\ell=1}^L a_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j) : \mathbf{a} \in \Theta_x(c) \right\}$$

equals a closed interval ranging from the infimum to the supremum.

2. The mapping $f(\cdot)$ defined by

$$f(\mathbf{a}) = \sum_{j=1}^J \sum_{\ell=1}^L a_{(\ell-1)J+j} \mathbb{P}(Z_\ell = z_\ell^j)$$

is continuous at all $\mathbf{a} \in \mathbb{R}^{LJ}$. The correspondence $\Theta_x(\cdot)$ is continuous and compact-valued for all c such that $\text{int}(\mathcal{D}(c)) \cap \text{int}(\mathcal{H}_x) \neq \emptyset$ by proposition 7. Therefore

$$\bar{P}_x(c) = \max\{f(\mathbf{a}) : \mathbf{a} \in \Theta_x(c)\}$$

is a continuous function of c on $\{c \in \mathcal{C} : \text{int}(\mathcal{D}(c)) \cap \text{int}(\mathcal{H}_x) \neq \emptyset\}$ by the Maximum Theorem (for example, see theorem 9.14 in Sundaram 1996). By continuity of $-f(\mathbf{a})$, $\underline{P}_x(c)$ is continuous over the same domain.

3. This follows immediately from the identified set $\Theta_0(c) \times \Theta_1(c)$ being a Cartesian product, which implies that maxima and minima for each component can be simultaneously attained.

□

Proofs for section 5.2

Proof of proposition 8. See proposition 3.1 of Kitagawa (2009). □

Proof of theorem 7. This proof follows the structure of the proof of theorem 6.

Step 1. We first show that the true densities $((\mathbf{f}_{Y_0|Z_\ell})_{\text{all } \ell}, (\mathbf{f}_{Y_1|Z_\ell})_{\text{all } \ell})$ are in $\Theta_0(c) \times \Theta_1(c)$. Fix $x \in \{0, 1\}$. By c_ℓ -dependence,

$$\begin{aligned} \mathbb{P}(Z_\ell = z | Y_x = y) &\leq \min\{\mathbb{P}(Z_\ell = z) + c_\ell, 1\} \\ \mathbb{P}(Z_\ell = z | Y_x = y) &\geq \max\{\mathbb{P}(Z_\ell = z) - c_\ell, 0\}. \end{aligned}$$

all $y \in \text{supp}(Y_x)$ and each $z \in \{0, 1\}$. These inequalities for $z = 0$ are equivalent to the inequalities for $z = 1$, so there are only two non-redundant inequalities here. Use Bayes' rule to rewrite the left hand side. This yields

$$\begin{aligned} f_{Y_x|Z_\ell}(y|z)\mathbb{P}(Z_\ell = z) &\leq \min\{\mathbb{P}(Z_\ell = z) + c_\ell, 1\}f_{Y_x}(y) \\ &= \min\{\mathbb{P}(Z_\ell = z) + c_\ell, 1\} \sum_{\tilde{z}=0}^1 f_{Y_x|Z_\ell}(y|\tilde{z})\mathbb{P}(Z_\ell = \tilde{z}). \end{aligned}$$

Hence

$$f_{Y_x|Z_\ell}(y|z)\mathbb{P}(Z_\ell = z) - \min\{\mathbb{P}(Z_\ell = z) + c_\ell, 1\} \sum_{\tilde{z}=0}^1 f_{Y_x|Z_\ell}(y|\tilde{z})\mathbb{P}(Z_\ell = \tilde{z}) \leq 0.$$

Evaluating at $z = 0$ yields

$$f_{Y_x|Z_\ell}(y|0)\mathbb{P}(Z_\ell = 0) - \min\{\mathbb{P}(Z_\ell = 0) + c_\ell, 1\}(f_{Y_x|Z_\ell}(y|0)\mathbb{P}(Z_\ell = 0) + f_{Y_x|Z_\ell}(y|1)\mathbb{P}(Z_\ell = 1)) \leq 0.$$

Rearranging yields

$$f_{Y_x|Z_\ell}(y|0)\mathbb{P}(Z_\ell = 0)(1 - \min\{\mathbb{P}(Z_\ell = 0) + c_\ell, 1\}) \leq \mathbb{P}(Z_\ell = 1) \min\{\mathbb{P}(Z_\ell = 0) + c_\ell, 1\}f_{Y_x|Z_\ell}(y|1)$$

or

$$f_{Y_x|Z_\ell}(y|0) \frac{\mathbb{P}(Z_\ell = 0)(1 - \min\{\mathbb{P}(Z_\ell = 0) + c_\ell, 1\})}{\mathbb{P}(Z_\ell = 1) \min\{\mathbb{P}(Z_\ell = 0) + c_\ell, 1\}} \leq f_{Y_x|Z_\ell}(y|1).$$

In the numerator,

$$1 - \min\{\mathbb{P}(Z_\ell = 0) + c_\ell, 1\} = \max\{\mathbb{P}(Z_\ell = 1) - c_\ell, 0\}.$$

Hence we have

$$f_{Y_x|Z_\ell}(y|0)k_0^\ell(c_\ell) \leq f_{Y_x|Z_\ell}(y|1).$$

Repeat these derivations for the other inequality evaluated at $z = 0$ to get the second inequality in the definition of $\mathcal{D}_{x,\ell}(c_\ell)$. Thus we have shown that the true conditional densities $\mathbf{f}_{Y_x|Z_\ell}$ are in $\mathcal{D}_{x,\ell}(c_\ell)$ for all ℓ . Since $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell}$ is just the vector of these conditional densities, we have shown that it is in $\mathcal{D}_x(c)$.

Next we show that $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell} \in \mathcal{H}_x$. For $z \in \{0, 1\}$, element $(\ell - 1)2 + 1 + z$ of $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell}$ can be written as

$$\begin{aligned} &[(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell}]_{(\ell-1)2+z} \\ &= f_{Y_x|Z_\ell}(y|z) \\ &= f_{Y_x,X|Z_\ell}(y,x|z) + f_{Y_x,X|Z_\ell}(y,1-x|z) \\ &= f_{Y,X|Z_\ell}(y,x|z) + f_{Y_x,X|Z_\ell}(y,1-x|z) \\ &= f_{Y,X|Z_\ell}(y,x|z) + \sum_{k=1}^{2^L} f_{Y_x|X,Z_\ell,Z_{-\ell}}(y|1-x,z,z_{-\ell}^k)\mathbb{P}(X=1-x, Z_{-\ell}=z_{-\ell}^k | Z_\ell=z). \end{aligned}$$

Since

$$f_{Y_x|X,Z_\ell,Z_{-\ell}}(y,1-x|z,z_{-\ell}^k) \in \mathcal{P}(\mathcal{Y}_x)$$

these derivations imply that

$$(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell} = (\mathbf{f}_{Y,X|Z_\ell}(\cdot, x))_{\text{all } \ell} + \mathbf{A}_x \mathbf{q}$$

where $\mathbf{q} \in \mathcal{P}(\mathcal{Y}_x)^{2^L}$ is a vector with elements $f_{Y_x|X,Z_\ell,Z_{-\ell}}(y | 1-x, z, z_{-\ell}^k)$. Thus $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell} \in \mathcal{H}_x$.

Thus we have shown that $(\mathbf{f}_{Y_x|Z_\ell})_{\text{all } \ell} \in \mathcal{D}_x(c) \cap \mathcal{H}_x$. Since this is true for each $x \in \{0, 1\}$, we have $((\mathbf{f}_{Y_0|Z_\ell})_{\text{all } \ell}, (\mathbf{f}_{Y_1|Z_\ell})_{\text{all } \ell}) \in \Theta_0(c) \times \Theta_1(c)$.

Step 2. Next we show sharpness. Let $(\mathbf{f}_0, \mathbf{f}_1) \in \Theta_1(c) \times \Theta_0(c)$. First, $\mathbf{f}_x \in \mathcal{D}_x(c)$ implies

$$\mathbb{P}(Z_\ell = z | Y_x = y) \in [\max\{\mathbb{P}(Z_\ell = z) - c_\ell, 0\}, \min\{\mathbb{P}(Z_\ell = z) + c_\ell, 1\}]$$

for all $\ell \in \{1, \dots, L\}$ and all $z \in \{0, 1\}$. This follows by reversing the arguments at the beginning of step 1. Hence c_ℓ -dependence is satisfied for all $\ell \in \{1, \dots, L\}$.

Next, since $\mathbf{f}_x \in \mathcal{H}_x$, there is a vector of counterfactual densities $\mathbf{q}_x \in \mathcal{P}(\mathcal{Y}_x)^{2^L}$ such that

$$\mathbf{f}_x = (\mathbf{f}_{Y,X|Z_\ell}(\cdot, x))_{\text{all } \ell} + \mathbf{A}_x \mathbf{q}_x.$$

This vector \mathbf{q}_x consists of elements of the form $f_{Y_x|X,Z}(\cdot | 1-x, z)$. Take these densities and combine them with the observed densities

$$f_{Y|X,Z}(\cdot | x, z) = f_{Y_x|X,Z}(\cdot | x, z)$$

to get a distribution of $Y_x | (X, Z)$. Do this for both $x \in \{0, 1\}$. Combine these conditional marginal distributions into a joint distribution $(Y_0, Y_1) | (X, Z)$ using any copula, and finally combine these with the known marginals (X, Z) to get a joint distribution of (Y_0, Y_1, X, Z) . By construction, this joint distribution is consistent with c -dependence, the distribution of the observed data (Y, X, Z) , and yields the point $(\mathbf{f}_0, \mathbf{f}_1)$ that we started with. \square