

Design-based identification with formula instruments: A review

KIRILL BORUSYAK[†], PETER HULL[‡] AND XAVIER JARAVEL[§]

[†]UC Berkeley, 222 Giannini Hall, Berkeley CA, USA 94720.

Email: k.borusyak@berkeley.edu

[‡]Brown University, 64 Waterman Street, Providence RI, USA 02912.

Email: peter.hull@brown.edu

[§]London School of Economics, Houghton Street, London, United Kingdom WC2A 2AE.

Email: x.jaravel@lse.ac.uk

First version received: 20 June 2023; final version accepted: 12 January 2024.

Summary: Many studies in economics use instruments or treatments that combine a set of exogenous shocks with other predetermined variables via a known formula. Examples include shift-share instruments and measures of social or spatial spillovers. We review recent econometric tools for this setting, which leverage the assignment process of the exogenous shocks and the structure of the formula for identification. We compare this design-based approach with conventional estimation strategies based on conditional unconfoundedness, and contrast it with alternative strategies that leverage a model for unobservables.

Keywords: *formula instruments, shift-share instruments, network spillovers, design-based identification.*

JEL codes: C21, C26.

1. INTRODUCTION

Many studies in economics use instruments or treatments that combine a set of observed shocks with other predetermined variables via a known formula. A leading example is shift-share instruments, which sum or average a common set of shocks $(g_k)_{k=1}^K$, varying at a different ‘level’ than observations i , with weights s_{ik} (‘shares’) reflecting heterogenous shock exposure: $z_i = \sum_k s_{ik} g_k$.¹ Other examples include treatments capturing the transmission of shocks in social or geographic networks, and variables based on formulas for policy eligibility.²

This paper reviews a recent econometric literature that shows how causal effects or structural parameters of linear models can be estimated in such settings when the shocks in the formula are exogenous with a known ‘design’. Exogeneity here means that the shocks are conditionally

¹ Shift-share (or ‘Bartik’) instruments were originally developed by Bartik (1991) and Blanchard and Katz (1992) to estimate labour demand elasticities. More recent applications study topics in trade (e.g., Autor et al., 2013; Hummels et al., 2014), immigration (e.g., Card, 2009; Peri et al., 2016), finance (e.g., Greenstone et al., 2020; Xu, 2022), public economics (e.g., Saiz, 2010; Diamond, 2016), and macroeconomics (e.g., Nakamura and Steinsson, 2014; Jaravel, 2019; Oberfield and Raval, 2021).

² Settings with such treatments and instruments include Miguel and Kremer (2004), who study spillovers from deworming shocks across students; Donaldson and Hornbeck (2016), who study market access effects from new railway construction; and Currie and Gruber (1996), who study the effects of Medicaid eligibility.

independent of the outcome error, allowing all other components of the formula to be endogenous. A known design means a specification of the assignment process from which the vector of observed shocks is drawn. These assumptions jointly formalize the notion of the shocks being ‘as-good-as-randomly assigned’. Borusyak et al. (2022, henceforth BHI) first develop this approach to shift-share identification and consistency, while Adão et al. (2019) study inference in this setting. Borusyak and Hull (2023, henceforth BH) extend the approach to general formula instruments. Our review unifies the BHI and BH frameworks and draws connections to alternative approaches to causal inference with formula instruments, such as by Goldsmith-Pinkham et al. (2020).

The key problem addressed by BHI and BH is that the exogeneity of the shocks is not by itself enough to avoid omitted variable bias (OVB). This is easily seen with shift-share instruments: suppose the shocks are drawn completely at random (with, say, a positive expectation), but observations vary in the sum of exposure shares to all shocks, $S_i = \sum_k s_{ik}$. The instrument will then be positively correlated with the sum of shares, since $\mathbb{E}[z_i | s_{i1}, \dots, s_{iK}] = \mathbb{E}[g_k] S_i$, and thus may also correlate with the errors. Randomizing the shocks does not imply instrument validity when shock exposure is not exogenous.

BHI and BH show how leveraging the formula allows one to avoid OVB when the shock design is known. In the previous example, the shift-share formula along with the random assignment of shocks imply that all systematic variation in the instrument is captured by the sum of shares; controlling for this sum eliminates OVB. Identification in the general case, for arbitrary formulas and designs, follows from simple adjustments based on the *expected instrument*: the average value of the formula across counterfactual sets of shocks, drawn from the specified assignment process. Specifically, OVB is avoided by either adding the expected instrument as a control or by using a *recentred instrument* that subtracts the expected instrument from the original formula. Controlling for or recentring by the expected instrument is generally necessary for identification with formula instruments, absent additional assumptions on the exogeneity of shock exposure.

How can the shock assignment process be specified in practice? This is trivial when the shocks are generated from a true experiment as shock counterfactuals then follow from the randomization protocol.³ In observational data, it is useful to distinguish two cases. With shift-share instruments, only the conditional mean of the shocks needs to be specified. BHI note that this task, while nontrivial, is no harder than the usual one of selecting controls in a more conventional shock-level analysis. In the general case, shock design can follow from the exchangeability of shocks conditional on observables, such that appropriate permutations of shocks can serve as valid shock counterfactuals. Exchangeability of the running variable around a policy threshold, as in the local randomization approach to regression discontinuity designs (e.g., Lee, 2008), scientific knowledge (e.g., geological models when using earthquakes as shocks), and other forms of institutional knowledge can also serve as sources of shock design. Such strategies build on a long tradition in the analysis of randomized experiments, going back to Neyman (1923).

We next review results on the consistency of IV estimators based on (recentred) formula instruments, and on valid statistical inference using them. BHI and BH show that shift-share IV estimators converge to the true parameter given a large number of exogenous shocks with sufficiently dispersed average exposure across observations. With shift-share instruments, valid asymptotic inference can be conducted under similar assumptions using results in Adão et al. (2019) or using an equivalent shock-level IV regression proposed by BHI. In settings with a

³ Note, however, that even in randomized experiments the assumption of exogenous shocks is not trivial: it requires the included treatment variable to capture all mechanisms through which the common shocks affect observations (an implicit exclusion restriction).

nonlinear formula or few shocks, randomization inference (RI) can be a natural alternative with valid finite-sample coverage.

The baseline setting we consider involves a constant-effect causal or structural model with a scalar treatment in a cross-section of observations and a known shock design. We discuss, however, various extensions of the design-based approach: to settings with heterogeneous treatment effects, multiple treatments and instruments, shock designs with estimated parameters, and panel data. We also review the results of Borusyak and Hull (2021a) on the most efficient formula instrument constructions for a given treatment variable.

We compare design-based identification with two alternative strategies that do not leverage the structure of the formula for identification. First, we draw a connection to conventional identification methods based on a conditional unconfoundedness assumption. Most closely related is the E-estimator of Robins et al. (1992), which uses as an instrument the difference between a treatment and an estimate of its expectation conditional on the covariates. The formula instrument setting is more general, involving instruments constructed from a common set of exogenous shocks that may vary at a different ‘level’ than the observations. Unlike conventional settings, the potentially confounding relationship between the instrument and shock exposure cannot be learned from the data since only one set of shocks is observed. Instead, OVB can be purged by specifying or estimating the assignment process for shocks and using the formula. Second, we contrast the design-based approach to formula instruments with strategies that do not directly use the formula, but instead model the unobservables (via, e.g., a parallel trends assumption) as in the Goldsmith-Pinkham et al. (2020) approach to shift-share instruments. Such approaches may be especially useful when there are too few shocks or when it is difficult to credibly specify the shock design. We note, however, that the underlying assumptions of these approaches can be inappropriate when observations are exposed to common unobserved shocks in the same way as they are exposed to the observed shocks in the formula instrument.

The rest of this paper is structured as follows. The next section outlines the basic setting and assumptions. Section 3 reviews results on identification, both in general and in special cases such as shift-share instruments. Section 4 reviews results on consistency and inference. Section 5 discusses extensions of the main results. Section 6 compares the design-based approach with alternative methods. Section 7 concludes with open questions in this recent literature.

2. SETTING

Consider a simple causal model for an outcome y_i and treatment x_i observed across units $i = 1, \dots, N$:

$$y_i = \beta x_i + \varepsilon_i. \quad (2.1)$$

Here β is the parameter of interest (i.e., the treatment effect) and ε_i is the unobserved error (i.e., unit i ’s potential outcome when x_i is set to zero). For initial notational simplicity, we assume y_i and x_i (and thus ε_i) have been de-measured across the units. To estimate β , we suppose a researcher has constructed a candidate ‘formula’ instrument

$$z_i = f_i(\mathbf{s}, \mathbf{g}), \quad (2.2)$$

for a set of known functions $f_1(\cdot), \dots, f_N(\cdot)$, a vector of observed shocks $\mathbf{g} \in \mathbb{R}^K$, and some observed data \mathbf{s} . The shocks are assumed to be exogenous conditional on \mathbf{s} and possibly some other observed \mathbf{q} ; formally, with $\boldsymbol{\varepsilon} = (\varepsilon_i)_{i=1}^N$ and $\mathbf{w} = (\mathbf{s}, \mathbf{q})$, we consider:

ASSUMPTION 2.1. (*Shock exogeneity*) $\boldsymbol{\varepsilon} \perp\!\!\!\perp \mathbf{g} \mid \mathbf{w}$.

Exogeneity, here formalized by the conditional independence of $\boldsymbol{\varepsilon}$ and \mathbf{g} given \mathbf{w} , is satisfied in a four-stage data-generating process where first \mathbf{w} is determined, second \mathbf{g} and $\boldsymbol{\varepsilon}$ are drawn independently (but possibly in a way that depends on \mathbf{w}), third $\mathbf{x} = (x_i)_{i=1}^N$ is determined in an unspecified way, and fourth $\mathbf{y} = (y_i)_{i=1}^N$ is determined by equation (2.1). We discuss the mapping of Assumption 2.1 onto several empirical contexts in Section 3.2; below we discuss a weaker mean-independence condition that suffices in some important settings. Exogeneity can also be understood as combining two distinct assumptions: exclusion of the shocks from the causal model and conditionally independent shock assignment. The working paper of BH (Borusyak and Hull, 2021b, appendix C.1) presents a more general potential outcomes model that isolates these two assumptions.

Three examples of candidate instrument constructions illustrate the generality of this setup. In the first example of linear shift-share IV, the instrument is an average (or sum) of the shocks g_k with weights s_{ik} reflecting heterogeneous shock exposure:

EXAMPLE 2.1. (*Linear shift-share*) \mathbf{s} is an $N \times K$ matrix of $s_{ik} \geq 0$ and $z_i = \sum_{k=1}^K s_{ik} g_k$.

For instance, supply shocks in China g_k across industries k could be combined with industry employment shares s_{ik} across US commuting zones i to study local employment outcomes, as in Autor et al. (2013). A key feature of such z_i is the linearity of $f_i(\mathbf{s}, \mathbf{g})$ in \mathbf{g} ; as we discuss below, this simplifies identification, consistency, and inference. This instrument construction is also ‘anonymous’, in that $f_i(\mathbf{s}, \mathbf{g}) = f(s_i, \mathbf{g})$ for a common $f(\cdot)$ across i and $s_i = (s_{ik})_{k=1}^K$. Often (though not always) $\sum_k s_{ik} = 1$, in which case z_i is a convex average of the shocks.⁴

A second example relaxes linearity of $f_i(\mathbf{s}, \mathbf{g})$ while maintaining the anonymous construction:

EXAMPLE 2.2. (*Nonlinear shift-share*) $z_i = f(s_i, \mathbf{g})$ for nonlinear $f(\cdot)$ and vectors s_i .

Here a concrete setting comes from Boustan et al. (2013), who instrument changes in the income Gini coefficient of municipalities i with a predicted change in the area’s Gini coefficient $f(s_i, \mathbf{g})$ based on national shocks g_k to the average income of worker groups (defined by percentiles of initial income) k with baseline regional shares of these groups s_{ik} . Another popular construction in this class, seen for instance in Berman et al. (2015), comes from taking the natural log of a linear shift-share variable: $z_i = \log(\sum_k s_{ik} g_k)$.

A third example of network spillovers illustrates the generality of the setting by relaxing both linearity and anonymity:

EXAMPLE 2.3. (*Network spillovers*) $K = N$, $\mathbf{s} = (s_{ik})_{i,k=1}^N$ is a network adjacency matrix, and $z_i = f_i(\mathbf{s}, \mathbf{g})$.

In Carvalho et al. (2021), for instance, a natural disaster generates shocks g_k to firms k that propagate across the firm-to-firm supplier network. Here $f_i(\mathbf{s}, \mathbf{g})$ returns the network distance between i and the nearest shocked firm, which depends on the full adjacency matrix \mathbf{s} . Of course, there are also network spillover applications that satisfy linearity and anonymity. For example, in Miguel and Kremer (2004), health shocks from a randomized deworming intervention propagate across a network of students captured by \mathbf{s} . Here $f_i(\mathbf{s}, \mathbf{g}) = s_i' \mathbf{g}$ simply counts the number of

⁴ The restriction that $s_{ik} \geq 0$ is not essential. We impose it because it is satisfied in most applications while it simplifies the results below. In particular, the conditions for consistent estimation in Proposition 4.1 need to be adjusted to accommodate negative exposure shares.

dormed neighbours of student i , showing this example of network spillovers can also be viewed as a linear shift-share variable with varying sum of exposure shares $\sum_k s_{ik}$.

BH discuss other examples of formula instruments, including market access instruments used to study the effects of transportation upgrades (as in Donaldson and Hornbeck, 2016) and policy eligibility instruments used to study the effects of complex eligibility changes (as in Currie and Gruber, 1996). The class of formula instruments is quite broad, containing any z_i which can be computed from a set of observed shocks \mathbf{g} and other observed data \mathbf{s} .

REMARK 2.1. The setup allows $x_i = z_i$, in which case β captures the reduced-form effect of the formula instrument itself. The exclusion restriction in this case requires the formula $f_i(\cdot)$ to be correct, in the sense of capturing all channels by which the shocks affect the outcome. In the more general IV case, Assumption 2.1 allows the consideration of any formula $f_i(\cdot)$ as long as x_i captures all shock effects. While strong, this exclusion restriction is standard.

REMARK 2.2. Formula instruments are often derived from the structure of x_i by removing or replacing some endogenous components. For example, canonical shift-share instruments were derived in settings where $x_i = \sum_k s_{ik} \tilde{x}_{ik}$ for observation-specific shocks \tilde{x}_{ik} (e.g., local industry growth rates), which were replaced with national averages when constructing z_i . Similarly, in the Boustan et al. (2013) example, z_i replaces the local income shocks at national income percentiles with the corresponding national changes in their predictions of local Gini coefficients. We return to this relationship between x_i and z_i when discussing instrument relevance in Section 4.1 and asymptotic efficiency in Section 5.

REMARK 2.3. The setup does not assume the (y_i, x_i) are independently and identically distributed (*iid*), as when sampled from some population, or put other restrictions on the dependence of ε_i across units. This permits the units to have common exposure to shocks, both observed and unobserved. The approach is consistent with the design-based tradition of conditioning on the set of potential outcomes, though we do not require such conditioning. It further allows the N units to represent a population—for example, all regions of a country (Abadie et al., 2020)—where random sampling assumptions are inappropriate.

In addition to shock exogeneity, we assume knowledge about the shock ‘assignment process’—i.e., restrictions on the distribution of \mathbf{g} given \mathbf{w} . The most demanding version of such design knowledge is complete specification of this distribution:

ASSUMPTION 2.2. (*Known design*) The distribution of \mathbf{g} given \mathbf{w} , denoted $G(\mathbf{g} \mid \mathbf{w})$, is known.

In randomized controlled trials (RCTs), where \mathbf{g} is drawn from according to an experimental protocol that may depend on \mathbf{q} (e.g., Miguel and Kremer, 2004), this assumption holds trivially as $G(\mathbf{g} \mid \mathbf{w})$ is then given by the protocol. For instance, the g_k may be independent Bernoulli random variables with known strata-specific means, with strata indicators included in \mathbf{q} .

Outside of true experiments, Assumption 2.2 may be satisfied by appropriate exchangeability assumptions which specify the permutations of shocks that were as likely to have occurred. For example, assuming the g_k are *iid* across k conditional on \mathbf{s} implies permutations of \mathbf{g} are equally likely to arise. Hence $G(\mathbf{g} \mid \mathbf{w})$ is known to be uniform over permutations $\Pi(\mathbf{g}) = \{\pi(\mathbf{g}) \mid \pi \in \Pi_K\}$, with Π_K denoting the set of permutation operators $\pi(\cdot)$ on vectors of length K , and with $\mathbf{q} = \Pi(\mathbf{g})$. Similar design knowledge follows under weaker shock exchangeability conditions, such as when the g_k are *iid* within, but not across, a set of known clusters and \mathbf{q} contains the class of within-cluster permutations of \mathbf{g} . BH, for example, consider a setting where the shocks

are indicators for whether each planned interregional railway line has opened before a certain date; the treatment is regional market access created by those lines. The timing of line opening is assumed random within groups of planned lines with similar characteristics—specifically, the number of regions a line connects. Thus, expected market access is computed by permuting the opening dates of planned lines that connect the same number of regions. Exchangeability could also be assumed local to some policy threshold, as in the local randomization view of regression discontinuity in Lee (2008) and Cattaneo et al. (2015). Scientific or other institutional knowledge (e.g., geological models when using earthquakes as shocks) may also yield specifications of $G(\mathbf{g} \mid \mathbf{w})$, as we discuss below.

An important relaxation of Assumptions 2.1 and 2.2 follows when $f_i(s, \mathbf{g})$ is linear in the shocks, as in Example 2.1. Specifically, they can be jointly weakened to a restriction on the conditional mean of the shocks:

ASSUMPTION 2.3. (*Conditionally linear shock means*) *There exists unknown θ such that $\mathbb{E}[g_k \mid \boldsymbol{\varepsilon}, \mathbf{w}] = \mathbf{q}'_k \theta$ for known $\mathbf{q} = (\mathbf{q}_k)_{k=1}^K$ included in \mathbf{w} and for all k .*

This restriction can be seen as combining a shock mean-independence condition of $\mathbb{E}[g_k \mid \boldsymbol{\varepsilon}, \mathbf{w}] = \mathbb{E}[g_k \mid \mathbf{w}]$ with a linearity restriction of $\mathbb{E}[g_k \mid \mathbf{w}] = \mathbf{q}'_k \theta$ (for a θ that can implicitly depend on \mathbf{w}). The former condition weakens Assumption 2.1 to allow higher moments of the shocks to depend on the unobserved $\boldsymbol{\varepsilon}$, while the latter restriction weakens Assumption 2.2 to a partial parameterization of $G(\mathbf{g} \mid \mathbf{w})$. For example, in the shift-share application of Autor et al. (2013), the latter condition allows the mean of industry-specific supply shocks in China to systematically vary across broader sectors (e.g., consumer electronics, apparel, food, etc.). This would be captured by a set of sectoral dummies \mathbf{q}_k , with the aim of leveraging residual shock variation within each sector (e.g., between TVs and computers). Below we further discuss and illustrate the usefulness of Assumption 2.3 in the linear shift-share setting, while also showing how it can be used with linear approximations of nonlinear shift-share instruments. In Section 5, we discuss other incomplete specifications for $G(\mathbf{g} \mid \mathbf{w})$.

3. IDENTIFICATION

3.1. Expected and recentered instruments

We first consider IV identification of β , or whether the parameter can be learned from first-stage and reduced-form moments.⁵ Formally, we consider the key orthogonality condition of $\mathbb{E}[\frac{1}{N} \sum_i z_i \varepsilon_i] = 0$. Identification follows under this condition with $\beta = \mathbb{E}[\frac{1}{N} \sum_i z_i y_i] / \mathbb{E}[\frac{1}{N} \sum_i z_i x_i]$ so long as $\mathbb{E}[\frac{1}{N} \sum_i z_i x_i] \neq 0$: a relevance condition we return to below.⁶

A key insight of BH is that shock exogeneity is not, by itself, enough for formula instrument orthogonality. To see this, define $\mu_i(\mathbf{w}) \equiv \mathbb{E}[f_i(s, \mathbf{g}) \mid \mathbf{w}]$ as the expectation of the candidate

⁵ Identification in this setting is less straightforwardly defined than it would be under the usual assumption of *iid* data (see Remark 2.3). Following Goldsmith-Pinkham and Imbens (2013), we approach the problem in two steps: first by considering whether certain moment conditions hold (in this section) and then considering whether the moments can be asymptotically learned from the data (in our consistency results in Section 4.1). Appropriately for non-*iid* data, we consider full-sample moments like $\mathbb{E}[\frac{1}{N} \sum_i z_i \varepsilon_i]$ that generalize moments like $\mathbb{E}[z_i \varepsilon_i]$ studied in *iid* data.

⁶ Expectations here are taken with respect to the joint distribution of $(\mathbf{g}, \mathbf{x}, \mathbf{w}, \boldsymbol{\varepsilon})$. The framework allows, but does not require, $(\mathbf{w}, \boldsymbol{\varepsilon})$ to be considered fixed, in line with the design-based tradition of focusing on randomness in the shocks and induced variation in the treatment.

instrument of unit i across the distribution of shocks while conditioning on the predetermined data in s and q . BH call this object, which we write as μ_i for brevity, the *expected instrument*. They show that under Assumption 2.1:

$$\mathbb{E} \left[\frac{1}{N} \sum_i z_i \varepsilon_i \right] = \mathbb{E} \left[\frac{1}{N} \sum_i \mu_i \varepsilon_i \right]. \quad (3.1)$$

In other words, shock exogeneity does not make the cross-sectional covariance between z_i and ε_i zero, unless the expected instrument is itself orthogonal to ε_i . Adjusting for μ_i —the key confounder in this setting—is therefore necessary for IV identification, absent additional restrictions on the exogeneity of shock exposure. The proof of (3.1) follows by repeated use of the law of iterated expectations: for any i , $\mathbb{E} [z_i \varepsilon_i] = \mathbb{E} [\mathbb{E} [f_i(\mathbf{g}, \mathbf{w}) \varepsilon_i \mid \mathbf{w}]] = \mathbb{E} [\mu_i \mathbb{E} [\varepsilon_i \mid \mathbf{w}]] = \mathbb{E} [\mu_i \varepsilon_i]$, using Assumption 2.1 in the second equality. When z_i is linear in the shocks, as with linear shift-share instruments, the same result holds under Assumption 2.3.

Two simple adjustments for μ_i follow: recentring and controlling. Note first that equation (3.1) immediately implies that the *recentred instrument* $\tilde{z}_i = z_i - \mu_i$ satisfies orthogonality:

$$\mathbb{E} \left[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i \right] = 0. \quad (3.2)$$

Thus, as long as it also satisfies relevance, $\mathbb{E} [\frac{1}{N} \sum_i \tilde{z}_i x_i] \neq 0$, using \tilde{z}_i as an instrument identifies β . This follows because $\mathbb{E} [\tilde{z}_i \mid \mathbf{w}] = 0$ by construction. Hence \tilde{z}_i cannot be correlated with any function of \mathbf{w} which, through its unrestricted dependence with ε , could violate orthogonality.

The second adjustment is perhaps more familiar: controlling for the confounder μ_i —or, more generally, for a vector of functions $\mathbf{r}_i \equiv \mathbf{r}_i(\mathbf{w})$ that linearly span μ_i —in an IV regression that uses the unadjusted z_i as the instrument. This approach works because controlling for such a \mathbf{r}_i implicitly recentres z_i by residualizing it on μ_i . Formally, by the Frisch–Waugh–Lovell theorem, the controlled IV regression is equivalent to an uncontrolled IV regression of y_i^\perp on x_i^\perp where \perp denotes the residuals from an in-sample projection on \mathbf{r}_i . The orthogonality condition of this regression,

$$\mathbb{E} \left[\frac{1}{N} \sum_i z_i \varepsilon_i^\perp \right] = 0, \quad (3.3)$$

holds because $\mathbb{E} [\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i^\perp] = 0$ by analogy to equation (3.1) and because $\mathbb{E} [\frac{1}{N} \sum_i (z_i - \tilde{z}_i) \varepsilon_i^\perp] = 0$ by the orthogonality between fitted values and residuals of the in-sample projection.⁷ The \mathbf{r}_i controls remove some variation from ε_i , which may improve efficiency in large samples.

Under Assumption 2.2, the recentring and controlling procedures are both straightforward to implement because μ_i is easily computed. In some cases the expectation $\mathbb{E} [z_i \mid \mathbf{w}]$ can be computed analytically or (as we show in the next section) is known to be linear in some observed \mathbf{r}_i . Otherwise, μ_i can be approximated by a simple simulation procedure:

⁷ Similarly, if the recentred instrument is used, any set of functions of \mathbf{w} can be controlled for, whether or not this set includes μ_i . For this reason, BH argue recentring is the key step for identification, while controlling for \mathbf{r}_i is a convenient implementation. We note that these results go through even if the number of controls is large. As is well known (e.g., Freedman, 2008), including many controls can create bias with heterogeneous effects (which we consider in Section 5) or otherwise reduce estimation efficiency or affect the validity of inference techniques.

- (1) For simulations $j = 1, \dots, J$, redraw counterfactual shocks $\mathbf{g}^{(j)}$ from $G(\cdot | \mathbf{w})$ (e.g., by following the RCT protocol or taking appropriate permutations of observed shocks).
- (2) Compute $z_i^{(j)} = f_i(\mathbf{s}, \mathbf{g}^{(j)})$.
- (3) For each observation, take the average across simulations: $\mu_i \approx \frac{1}{J} \sum_{j=1}^J z_i^{(j)}$.

For identification, it is enough to approximate μ_i by any number of simulations J so long as draws are independent of \mathbf{g} and $\boldsymbol{\varepsilon}$ conditionally on \mathbf{w} : $\mathbb{E} \left[\frac{1}{N} \sum_i (z_i - \frac{1}{J} \sum_j f_i(\mathbf{s}, \mathbf{g}^{(j)})) \varepsilon_i \right] = 0$ by the law of iterated expectations, since $\mathbb{E} [z_i | \mathbf{w}, \boldsymbol{\varepsilon}] = \mathbb{E} [f_i(\mathbf{s}, \mathbf{g}^{(j)}) | \mathbf{w}, \boldsymbol{\varepsilon}]$.⁸

REMARK 3.1. Should recentring or controlling be used in practice? BH note that the answer may depend on whether the shocks arise from an RCT or a natural experiment. In RCTs, Assumption 2.2 is satisfied trivially. Thus, the researcher can first recentre the instrument to avoid OVB and then choose the controls solely by efficiency considerations. Controlling for μ_i , as opposed to any other \mathbf{r}_i , is only preferable if μ_i is believed to be strongly correlated with ε_i . In a natural experiment, however, the controlling strategy has an additional appeal relative to recentring, as it allows the researcher to incorporate multiple guesses for μ_i arising from different specifications of $G(\mathbf{g} | \mathbf{w})$. If these candidate μ_i are included as controls while using z_i as the instrument, identification follows if any one candidate is correctly specified.

REMARK 3.2. A related approach to avoid bias from nonrandom exposure is developed by Aronow and Samii (2017) in the context of discrete network treatments x_i and discrete shocks g_k . While their estimator is based on inverse probability weighting, recentring is a regression adjustment and hence does not require an additional overlap condition.

3.2. Linear shift-share IV

A particularly important case of formula instruments is linear shift-share instruments, with $z_i = \sum_k s_{ik} g_k$, where identification can be achieved under weaker conditions than in the general case. We consider three cases of increasing complexity.

Case 1: Complete shares, no controls We first suppose the exposure shares are ‘complete’, in the sense that $\sum_k s_{ik} = 1$ for all i , and that Assumption 2.3 holds with no additional controls such that $\mathbf{q}_k = 1$ and $\mathbb{E} [g_k] = \theta$ for all k . The latter condition requires each shock g_k to have the same mean θ , regardless of the realizations of the unobservables $\boldsymbol{\varepsilon}$ and exposure shares \mathbf{s} .

In this case, instrument orthogonality holds without any recentring as long as the estimating equation includes an intercept. This follows because the expected instrument is constant, $\mu_i = \sum_k s_{ik} \mathbb{E} [g_k | \mathbf{w}] = \theta \sum_k s_{ik} = \theta$, and is thus absorbed by the intercept. The weaker mean-independence restriction of Assumption 2.3 is furthermore enough because of the linear construction of z_i . Intuitively, when all the shocks are exogenous and have the same expectation, there is no reason for a weighted average of those shocks z_i to be systematically higher or lower for units with any particular shares \mathbf{s}_i and thus with higher or lower ε_i .

For further intuition, it is instructive to rewrite the (\mathbf{r}_i -controlled) orthogonality condition at the shock ‘level’, following BHHJ:

$$0 = \mathbb{E} \left[\frac{1}{N} \sum_i z_i \varepsilon_i^\perp \right] = \mathbb{E} \left[\frac{1}{N} \sum_i \sum_k s_{ik} g_k \varepsilon_i^\perp \right] = \mathbb{E} \left[\sum_k s_k g_k \bar{\varepsilon}_k^\perp \right], \quad (3.4)$$

⁸ BH further show that the number of simulations is not important for consistency of the estimator. The proof of their proposition 1 shows the number of simulations generally affects the estimator’s efficiency, however.

where $s_k = \frac{1}{N} \sum_i s_{ik}$ and $\bar{\varepsilon}_k^\perp = \frac{\sum_i s_{ik} \varepsilon_i^\perp}{\sum_i s_{ik}}$ with ε_i^\perp again denoting the residualized error. This shows that the orthogonality of z_i and $\bar{\varepsilon}_k^\perp$ is equivalent to a weighted ‘shock-level’ orthogonality condition, which considers the relationship between g_k and unobserved $\bar{\varepsilon}_k^\perp$. The s_k weights capture the ‘importance’ of each shock, in terms of its average exposure s_{ik} , while $\bar{\varepsilon}_k^\perp$ are exposure-weighted averages of ε_i^\perp . Since $\sum_k s_k \bar{\varepsilon}_k^\perp = \frac{1}{N} \sum_i \sum_k s_{ik} \varepsilon_i^\perp = \frac{1}{N} \sum_i \varepsilon_i^\perp = 0$, with r_i including a constant, equation (3.4) shows that orthogonality holds when the weighted covariance of g_k and $\bar{\varepsilon}_k^\perp$ is zero.

As a concrete example, consider the setting of Aghion et al. (2022, henceforth AABJ), who study the impact of lowering the cost of automation technologies on labour demand in France in the 2000s. Here y_i is the change in employment of firm i , x_i is the change in the firm’s stock of automation technologies, and $z_i = \sum_k s_{ik} g_k$ is the predicted change in the cost of importing these technologies. This prediction is based on shocks g_k to the supply of imported technologies across source countries and technology categories, such that k indexes country-technology pairs (e.g., robots from Italy). Exposure shares s_{ik} are computed as the proportion of imports from cell k in firm i ’s total imports of automation technologies in a previous period (which sum to one).

In this example, Assumption 2.3 requires cell-level supply shocks to have the same expectation regardless of all firm-level unobservables. The representation (3.4) provides further intuition: the shift-share instrument would violate orthogonality if the country-technology pairs that have pre-existing relationships with French firms with growing employment (for reasons other than automation) receive systematically different supply shocks. For instance, orthogonality fails if better managed French firms (with better employment trends) established ties with Chinese suppliers in the 1990s in anticipation of the growing supply from China in the 2000s. We next relax the assumption of equal expected shocks to accommodate situations like this.

Case 2: Complete shares with controls We now turn to linear shift-share instruments with general shock-level controls q_k . In the AABJ application, for example, it appears important to allow for productivity shocks that differ systematically across supplier countries (e.g., China experiences positive productivity shocks in the 2000s) and for firms connected to certain countries to have systematically different unobservables. Similarly, one may allow productivity shocks to vary across technology categories (e.g., robots experience faster productivity growth than textile machines), with firms connected to those technologies being systematically different. The shock-level controls q_k could then include supplier country and technology fixed effects.

With shock-level controls, the expected instrument is linear in exposure-weighted averages of these controls:

$$\mu_i = \sum_k s_{ik} q_k' \theta = Q_i' \theta, \quad Q_i = \sum_k s_{ik} q_k. \tag{3.5}$$

Hence, we can view z_i as combining systematic (and potentially confounding) variation in $\mu_i = Q_i' \theta$ with idiosyncratic variation in shocks drawn above or below their expectation:

$$z_i = Q_i' \theta + \sum_k s_{ik} (g_k - q_k' \theta). \tag{3.6}$$

Including Q_i in the control vector r_i removes the systematic variation, ensuring instrument orthogonality. In the AABJ example this would mean controlling for the total exposure share of a firm to each supplying country, as well as the total exposure share to each technology.

An alternative approach in this case would be to first residualize the g_k shocks on the q_k controls then use a modified shift-share instrument $\sum_k s_{ik} (g_k - q_k' \hat{\theta})$, possibly with no additional controls. This can be viewed as a recentring approach, although with an estimated parameter $\hat{\theta}$

affecting shock expectations (we allow for shock assignment process specified with parameters in Section 5). In the AABJ case this can be implemented by residualizing cell-level supply shocks on country and technology fixed effects and constructing the shift-share instruments from the estimated residuals.

Case 3: Incomplete shares In certain shift-share applications, the exposure shares are ‘incomplete’ in that they do not sum to one: i.e., $S_i = \sum_k s_{ik}$ varies across observations i . This is the case in the ‘China shock’ study of Autor et al. (2013, henceforth ADH) where i denotes commuting zones and k indexes manufacturing industries. ADH estimate the response of labour market outcomes (e.g., manufacturing employment) across commuting zones to exposure to industry-specific supply shocks in China. Here g_k captures the supply shock in China, while the shares s_{ik} are lagged industry employment shares in each commuting zone. The shares are measured relative to total employment (i.e., both in manufacturing and in other sectors) and so they add up to the lagged total manufacturing share of region i : $S_i < 1$.

Even when Assumption 2.3 holds with no controls ($\mathbf{q}_k = 1$), the incompleteness of shares can make instrument orthogonality fail. Intuitively, z_i leverages variation in S_i which may be correlated with the error term even when shocks are fully random. More formally, we again consider the expected instrument:

$$\mu_i = \sum_k s_{ik} \mathbb{E}[g_k] = \theta \sum_k s_{ik} = \theta S_i.$$

Instrument orthogonality generally fails when $\theta \neq 0$ and when S_i is cross-sectionally correlated with ε_i . For example, in the ADH setting, industry-level China shocks g_k are positive in expectation such that z_i is systematically higher in regions where manufacturing as a whole is a bigger share of the local economy. Because manufacturing employment is on a downward trend for reasons other than trade with China (e.g., structural transformation), this can be a source of negative bias.⁹

Controlling for the sum of shares S_i isolates the idiosyncratic variation in z_i which yields instrument orthogonality.¹⁰ Constructing the shift-share instrument from de-measured shocks is again an alternative recentring solution. The more general case with incomplete shares and other controls \mathbf{q}_k follows similarly. Here $\mu_i = \mathbf{Q}_i' \theta$ for $\mathbf{Q}_i = \sum_k s_{ik} \mathbf{q}_k$, but this μ_i is not an exposure-weighted average of \mathbf{q}_k since the weights do not add up to one. Still, as before, including \mathbf{Q}_i in \mathbf{r}_i would isolate the idiosyncratic variation in the instrument.¹¹

⁹ In practice, ADH include a control that is highly correlated with S_i : total manufacturing employment share from a later period. BHH show their estimates of manufacturing employment effects are similar when controlling for S_i itself.

¹⁰ As BHH note, another way to arrive at this solution is to imagine a ‘missing’ shock $g_0 = 0$ with exposure share $s_{i0} = 1 - S_i$: e.g., the nonmanufacturing sector in ADH that is not subject to the China shock. With this shock added, there are no incomplete shares. However, $\mathbb{E}[g_0] = 0$ while $\mathbb{E}[g_k] = \theta$ for $k \neq 0$. Thus, unless $\theta = 0$ (e.g., the expected supply shock in China in each manufacturing industry is zero), the indicator $\mathbf{1}[k > 0]$ should be included in \mathbf{q}_k , and correspondingly the share-weighted exposure to it, which equals S_i , should be controlled for.

¹¹ For example, ADH study two periods (the 1990s and 2000s) with systematically higher shocks to manufacturing industries in the 2000s. With period fixed effects in \mathbf{q}_k , the corresponding \mathbf{Q}_i includes the total manufacturing share interacted with period fixed effects. BHH show that adding this control is empirically relevant in the ADH context. The intuition is that regions with a higher total manufacturing share always have systematically higher z_i , but especially in the 2000s. If manufacturing-heavy regions face stronger declines in employment in the 2000s for reasons other than trade with China (again, perhaps because of structural change), the shift-share IV estimator is biased even if period dummies and S_i are separately controlled for.

3.3. Linear vs. nonlinear formulas

We now contrast identification with linear and nonlinear formula instruments, using the nonlinear shift-share example of $z_i = f(s_i, \mathbf{g})$ for concreteness. As with the linear case, a common motivation for such instruments is that the endogenous treatment of interest can be written as $x_i = f(s_i, \tilde{\mathbf{x}}_i)$ for some (potentially unobserved) $\tilde{\mathbf{x}}_i = (\tilde{x}_{ik})_{k=1}^K$. An intuitive instrument in this case replaces the \tilde{x}_{ik} with an exogenous shifter g_k . Concrete examples include the predicted change in a regional Gini coefficient in Boustan et al. (2013), the predicted share of migrants in Basso and Peri (2015), and the predicted foreign demand instrument of Berman et al. (2015).

Nonlinearity of $f(\cdot)$ complicates the structure of the expected instrument relative to the linear case. Even with complete shares and fully random shocks (i.e., *iid* g_k conditionally on \mathbf{w})—a case where μ_i is constant for linear shift-share instruments—a nonlinear shift-share instrument generally has a μ_i that depends on s_i in complex ways. Outside this case, μ_i further depends on the heteroskedasticity and mutual dependence between the shocks making the mean-independence condition (Assumption 2.3) insufficient.

Borusyak and Hull (2021b, appendix D.4) make these points concrete by considering an instrument that could be called a ‘shift-share in logs’—see, e.g., Berman et al. (2015), Berthou et al. (2019), and Costa et al. (2019). Suppose $x_i = \log(X_{i1}/X_{i0})$ measures the growth rate of a regional variable that can be represented as a sum of local industry components, $X_{it} = \sum_k \tilde{X}_{ikt}$ for $t \in \{0, 1\}$. Then x_i can be rewritten as a nonlinear function of initial shares $s_{ik} = \frac{\tilde{X}_{ik0}}{X_{i0}}$ and regional growth rates $\tilde{x}_{ik} = \frac{\tilde{X}_{ik1}}{\tilde{X}_{ik0}}$, as $x_i = \log(\sum_k s_{ik} \tilde{x}_{ik})$. Suppose the \tilde{x}_{ik} are endogenous, but we observe an industry characteristic G_{kt} with plausibly exogenous growth rates $g_k = \frac{G_{k1}}{G_{k0}}$ that predict the \tilde{x}_{ik} . A natural nonlinear shift-share instrument is then $z_i = \log(\sum_k s_{ik} g_k)$.

The log transformation makes μ_i heterogeneous across regions, even with fully random g_k and complete shares. In particular, by Jensen’s inequality, $\mathbb{E}[\log(\sum_k s_{ik} g_k) | s]$ is the lowest for regions where $\sum_k s_{ik} g_k$ has little variance conditional on s . Thus, regions with very dispersed industries (i.e., where $\sum_k s_{ik}^2 \approx 0$) will tend to have lower μ_i and z_i . Such regions may also have systematically different unobservables ε_i . For example, a more dispersed local economy may not reap the benefits from returns to scale leading to higher regional economic growth, which would generate a downward bias in an IV estimator taking growth as outcome.

One solution to this problem, following the general approach in Section 3.1, is to compute μ_i from a specification of shock counterfactuals and either centre by or control for it. An alternative solution is to take a first-order approximation of $f(\cdot)$ around some fixed vector of shocks to return to the linear case and obviate the need for a full specification of shock counterfactuals. For example, taking a log-linear approximation around $g_k = 1$ in the preceding example gives $\tilde{z}_i = \sum_k s_{ik} \log g_k$, with $s_{ik} = \frac{\partial f(s_i, 1)}{\partial \log g_k}$. This is a linear shift-share instrument, with logged shocks and exposure shares which need not sum to one. As an approximation to z_i , the linear instrument might predict x_i less well and thus be less efficient. However, unlike z_i , the orthogonality of \tilde{z}_i depends only on correct specification of $\mathbb{E}[\log g_k | \boldsymbol{\varepsilon}, \mathbf{w}]$ (i.e., Assumption 2.3), making identification more robust.

3.4. Nonanonymous constructions

We now consider cases where the formula used to build the instrument is nonanonymous: $z_i = f_i(s, \mathbf{g})$, i.e., where the exposure of unit i to the shocks cannot be naturally summarized by some

observation-specific vector s_i .¹² For example, one may study the propagation of shocks across a network where the units i represent nodes and \mathbf{g} captures shocks that are as-good-as-randomly assigned to nodes.

Avoiding bias may seem particularly challenging with such cases. Consider for example the setting of Carvalho et al. (2021), who study the effects of network distance between a focal firm i and the nearest firm located in the geographic area of the 2011 Tohoku earthquake. Exposure to the earthquake shock is given by the entire firm-to-firm supplier network, with network distance yielding a complex nonanonymous construction.

In this setting, it is not clear how to pick the right set of controls to span μ_i without leveraging the formula and some knowledge of the shock assignment process. However, the general simulation procedure from Section 3. yields μ_i given specification of shock counterfactuals: for example, by drawing on geological models of earthquake probabilities across regions to redraw earthquake realizations. This example demonstrates how the recentering logic and the specification of counterfactual shocks help address the specific challenges with nonanonymous (and nonlinear) constructions.

4. CONSISTENCY AND INFERENCE

4.1. Consistency with many shocks

We next consider consistency of IV estimators that recenter by or control for μ_i . Even if the orthogonality condition (3.2) holds and \tilde{z}_i is relevant, additional conditions are generally required for consistent estimation because of the potentially complex dependencies in the data. Indeed, \tilde{z}_i captures common exposure of the observations to the shocks in \mathbf{g} and may thus be correlated across i in nonstandard ways. Similarly, the ε_i may exhibit nonstandard dependencies from their common dependence on unobserved shocks. Consistency may nevertheless be guaranteed as long as \mathbf{g} includes a large number of sufficiently independent shocks regardless of the dependence structure of unobservables.

Formally, we consider the recentered IV estimator

$$\hat{\beta} = \frac{\sum_i y_i \tilde{z}_i}{\sum_i x_i \tilde{z}_i} = \beta + \frac{\sum_i \varepsilon_i \tilde{z}_i}{\sum_i x_i \tilde{z}_i}. \quad (4.1)$$

Here we assume that y_i and x_i (and correspondingly ε_i) have been residualized on some regression controls $\mathbf{r}_i \equiv \mathbf{r}_i(\mathbf{w})$, dropping the previous \perp notation. This definition of $\hat{\beta}$ therefore nests estimators that control for functions of \mathbf{w} that span μ_i (making recentering unnecessary) or include any predetermined controls after recentering z_i .

BHJ and BH study the convergence of $\hat{\beta}$ along a sequence of distributions P_N for the complete data $\{(y_i, x_i, z_i)_{i=1}^N, \mathbf{s}, \mathbf{g}, \mathbf{q}\}$ where the number of shocks $K_N = \dim(\mathbf{g})$ can vary with the number of observations N . In particular, they consider consistency: $\hat{\beta} \xrightarrow{P} \beta$ as $N \rightarrow \infty$. Here we suppress the N subscript for clarity.

As with the identification discussion in Section 3.2, it is instructive to first study consistency in the special case of shift-share instruments. Without loss of generality, assume $s_k = \frac{1}{N} \sum_i s_{ik} > 0$ for every k since, if $s_{ik} = 0$ for all i , shock g_k can be dropped. Consider:

¹² Note that it is technically always possible to write $z_i = f_i(\mathbf{s}, \mathbf{g})$ anonymously, for example as $z_i = \tilde{f}(\mathbf{a}_i, \mathbf{g})$ where \mathbf{a}_i includes copies of \mathbf{s} and an observation indicator such that $\tilde{f}(\cdot)$ returns $f_i(\mathbf{s}, \mathbf{g})$. Such constructions are unnatural in most settings with nonanonymous constructions, however.

ASSUMPTION 4.1. (a) *Many uncorrelated shocks*: $\mathbb{E} \left[\sum_{k=1}^K s_k^2 \right] \rightarrow 0$ and $\text{Cov} [g_k, g_{k'} \mid \boldsymbol{\varepsilon}, \mathbf{w}] = 0$ for all $k \neq k'$.

(b) *First stage*: $\frac{1}{N} \sum_i x_i \tilde{z}_i \xrightarrow{P} \pi \neq 0$.

(c) *Regularity*: $\text{Var} [g_k \mid \boldsymbol{\varepsilon}, \mathbf{w}] \leq B_g$ and $\mathbb{E} [\bar{\varepsilon}_k^2 \mid \mathbf{w}] \leq B_\varepsilon$ uniformly over N , for $\bar{\varepsilon}_k = \frac{1}{N} \sum_i s_{ik} \varepsilon_i / s_k$.¹³

We then have:

PROPOSITION 4.1. *In the shift-share case where $\tilde{z}_i = \sum_{k=1}^K s_{ik} \tilde{g}_k$ with $s_{ik} \geq 0$ and $\tilde{g}_k = g_k - \mathbf{q}'_k \theta$, Assumptions 2.3 and 4.1 imply $\hat{\beta} \xrightarrow{P} \beta$.*

Proof. Follows from proposition 3 of BHJ applied to the de-meaned shocks $\tilde{\mathbf{g}} = (\tilde{g}_k)_{k=1}^K$. \square

Assumption 4.1(a) ensures that the law of large numbers applies to $\frac{1}{N} \sum_i \varepsilon_i \tilde{z}_i$ regardless of the mutual correlations of ε_i . It requires that the number of shocks grows with the sample size (since $\sum_{k=1}^K s_k^2 \geq 1/K$), that shocks are mutually uncorrelated, and average exposure to them is dispersed in the sense of the Herfindahl index converging to zero.¹⁴ Just as in conventional regressions, mutual uncorrelatedness of the shocks can easily be relaxed: e.g., by having many uncorrelated clusters of shocks or other forms of weak dependence between them, with the Herfindahl condition appropriately strengthened (see assumptions 5 and 6 in BHJ). In a panel context, the large number of shocks can arise either from many shocks in each cross-section or from a long time series, a point we return to in Section 5.

Assumption 4.1(b) requires the first-stage covariance $\frac{1}{N} \sum_i \tilde{z}_i x_i$ to converge to some nonzero constant. BHJ provide low-level conditions sufficient for this using a model that often underlies the usage of shift-share instruments in the first place (and which holds trivially in reduced-form studies where $x_i = z_i$). Specifically, they assume that the treatment can naturally be decomposed into k -specific components each strongly affected by the respective shock: $x_i = \sum_k s_{ik} \tilde{x}_{ik}$ with $\tilde{x}_{ik} = \pi_{ik} g_k + u_{ik}$, $\pi_{ik} \geq \bar{\pi} > 0$ and $\mathbb{E} [\mathbf{g} \mid \mathbf{w}, \boldsymbol{\varepsilon}, \mathbf{u}] = \mathbb{E} [\mathbf{g} \mid \mathbf{w}]$ for $\mathbf{u} = (u_{ik})_{i,k}$. With this model and additional regularity conditions, Assumption 4.1(b) follows if $\mathbb{E} \left[\frac{1}{N} \sum_i (\sum_k s_{ik}^2) \right]$ is bounded from below. This last condition requires an average observation to be effectively exposed to only a small number of shocks, in the sense of a nonvanishing Herfindahl index of exposure shares, such that the law of large numbers does not apply to individual \tilde{z}_i and the variance of the instrument does not converge to zero. It is instructive to compare this condition with Assumption 4.1(a), which requires the *average* exposure to be dispersed. Both conditions can hold when most observations are exposed to a small set of shocks, but different ones for different observations. For example, they hold in the ADH setting when most local labour markets specialize in a small number of manufacturing industries—the identities of which vary across markets.

Propositions 1–6 in BH generalize Proposition 4.1 and the sufficient conditions for a nonvanishing first stage to nonlinear and nonanonymous formula instruments. While the assumptions become more technically restrictive, the main intuition extends: dispersed average exposure to sufficiently independent shocks makes $\frac{1}{N} \sum_i \varepsilon_i \tilde{z}_i$ converge to zero regardless of the mutual correlation in the errors, while concentrated individual shock exposure is key to the first stage. We

¹³ All conditions here and in Assumption 4.2 should be understood as applying almost surely with respect to \mathbf{w} .

¹⁴ The literal interpretation of $\sum_k s_k^2$ as a Herfindahl index and the result of $\sum_{k=1}^K s_k^2 > 1/K$ both technically require complete shares, $\sum_k s_{ik} = 1$, which implies $\sum_k s_k = 1$. However, in typical incomplete share cases where $\sum_k s_k$ is between zero and one and bounded away from zero, the same intuition applies.

reproduce one of the sets of sufficient conditions for convergence of $\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i$, applicable to continuous shocks:

ASSUMPTION 4.2. (a) *Many independent shocks: the components of \mathbf{g} are jointly independent conditionally on \mathbf{w} , and $\sum_k \mathbb{E} \left[\left(\frac{\partial \bar{f}(\mathbf{g}, \mathbf{w})}{\partial g_k} \right)^2 \right] \rightarrow 0$ for $\bar{f}(\mathbf{g}, \mathbf{w}) = \frac{1}{N} \sum_i (f_i(s, \mathbf{g}) - \mu_i(\mathbf{w}))$.*

(b) *Each $f_i(s, \mathbf{g})$ is weakly monotone in \mathbf{g} .*

(c) *Each g_k has a Gaussian distribution.*

(d) *Regularity: $\mathbb{E}[\varepsilon_i^2 | \mathbf{w}] \leq U_\varepsilon$ uniformly; $\mathbb{E} \left[\left| \frac{\partial \bar{f}(\mathbf{g}, \mathbf{w})}{\partial g_k} \right| | \mathbf{w} \right] < \infty$ for all k ; $\text{Var}[g_k | \mathbf{w}] \in [L_\sigma, U_\sigma]$ for fixed constants $0 < L_\sigma < U_\sigma < \infty$.*

PROPOSITION 4.2. *Assumptions 2.1 and 4.2 imply $\text{Var} \left[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i \right] \rightarrow 0$.*

Proof. Follows from propositions 1 and 4(a) of BH. □

BH also establish a similar result for Bernoulli-distributed shocks, where $\frac{\partial \bar{f}(\mathbf{g}, \mathbf{w})}{\partial g_k}$ is replaced by discrete derivatives, as well as analogous low-level conditions for convergence of the first stage.

4.2. Asymptotic inference with shift-share variables

We now consider statistical inference on β . Inference is challenging in this setting, even with a large number of shocks, again because the observations of z_i (ε_i) are jointly exposed to the observed (unobserved) shocks and may therefore exhibit complex mutual correlations. Adão et al. (2019, henceforth AKM) illustrate this problem in the linear shift-share setting, via a Monte Carlo simulation based on the specifications in Autor et al. (2013). They show that conventional heteroskedasticity-robust and state-clustered standard errors, with nominal 5% significance level, lead to rejections of the true null of $\beta = 0$ in around half of their placebo samples with randomly generated shocks.

To address this issue, AKM derive new design-based central limit theorems and variance estimators for $\hat{\beta}$ in the linear shift-share context, which are asymptotically valid regardless of the correlation structure of the errors—what BHJ label ‘exposure-robust’. The variance estimators can be motivated by first writing

$$\sqrt{N}(\hat{\beta} - \beta) = \frac{\frac{1}{N} \sum_i \varepsilon_i \tilde{z}_i}{\frac{1}{N} \sum_i x_i \tilde{z}_i} = \left(\frac{1}{\sqrt{N}} \sum_i \varepsilon_i \tilde{z}_i \right) \pi^{-1} (1 + o_p(1)), \quad (4.2)$$

using Assumption 4.1(b). For simplicity, suppose Assumption 2.3 holds with $\theta = 0$, such that $\mathbb{E}[g_k | \boldsymbol{\varepsilon}, \mathbf{w}] = 0$ and $\tilde{z}_i = z_i$. Then:

$$\begin{aligned} \text{Var}[\hat{\beta}] &\approx \frac{1}{\pi^2} \text{Var} \left[\frac{1}{N} \sum_i \varepsilon_i z_i \right] \\ &= \frac{1}{\pi^2} \text{Var} \left[\frac{1}{N} \sum_k g_k \left(\sum_i s_{ik} \varepsilon_i \right) \right] \\ &= \frac{1}{\pi^2} \sum_k \mathbb{E} \left[g_k^2 \left(\frac{1}{N} \sum_i s_{ik} \varepsilon_i \right)^2 \right], \end{aligned} \quad (4.3)$$

using Assumption 4.1(a) in the final line. This suggests a feasible estimator of $\frac{1}{\hat{\pi}^2} \sum_k g_k^2 \left(\frac{1}{N} \sum_i s_{ik} \hat{\varepsilon}_i \right)^2$ where $\hat{\pi} = \frac{1}{N} \sum_i \tilde{z}_i x_i$ and $\hat{\varepsilon}_i = y_i - \hat{\beta} x_i$. AKM prove the validity of similar variance estimators in the general case of $\theta \neq 0$ and when Assumption 4.1(a) is weakened to allow for clustered shocks.¹⁵

BHJ propose a convenient way to obtain exposure-robust standard errors based on this asymptotic approximation, which come from standard formulas applied to an equivalent IV regression estimated at the ‘level’ of the shocks. To see this result, it is first useful to note (following proposition 1 in BHJ) that $\hat{\beta}$ equals the second-stage coefficient from an s_k -weighted shock-level IV regression, which uses the shocks g_k as the instrument to estimate¹⁶

$$\bar{y}_k = \beta \bar{x}_k + \bar{\varepsilon}_k. \tag{4.4}$$

This result follows from the definition of z_i :

$$\hat{\beta} = \frac{\sum_i \frac{1}{N} \left(\sum_k s_{ik} g_k \right) y_i}{\sum_i \frac{1}{N} \left(\sum_k s_{ik} g_k \right) x_i} = \frac{\sum_k g_k \left(\frac{1}{N} \sum_i s_{ik} y_i \right)}{\sum_k g_k \left(\frac{1}{N} \sum_i s_{ik} x_i \right)} = \frac{\sum_k s_k g_k \bar{y}_k}{\sum_k s_k g_k \bar{x}_k}. \tag{4.5}$$

Moreover, if $\mathbf{Q}_i = \sum_k s_{ik} \mathbf{q}_k$ has been included in the controls \mathbf{r}_i , the same coefficient is obtained from a richer shock-level IV specification:

$$\bar{y}_k = \beta \bar{x}_k + \mathbf{q}'_k \delta + \bar{\varepsilon}_k. \tag{4.6}$$

The *ssaggregate* packages in Stata and R automate the translation of linear shift-share regressions to the shock level.

Under regularity conditions (see BHJ, proposition 5), the conventional heteroskedasticity-robust standard error for $\hat{\beta}$ in (4.6) yields asymptotically valid confidence intervals for β . While this derivation follows under the assumption of conditionally uncorrelated shocks (Assumption 4.1(a)), BHJ note the same shock-level regression can be used to conduct asymptotically valid inference when shocks are clustered, again using conventional clustered standard errors. Indeed, one convenient feature of their approach is the flexibility of accommodating other relaxations of Assumption 4.1(a), such as serial correlation (with conventional heteroskedastic-and-autocorrelation-consistent standard errors) or two-way clustering.

REMARK 4.1. While BHJ only consider the case where \mathbf{Q}_i has been initially controlled for, equation (4.6) produces valid estimates (and, we conjecture, standard errors) regardless of this, consistent with the recentring logic. Indeed, by the Frisch–Waugh–Lovell theorem, $\hat{\beta}$ from (4.6) can be obtained by residualizing the shocks on \mathbf{q}_k (with weights s_k) and constructing the shift-share instrument from the residualized shocks.

REMARK 4.2. A key assumption underlying the validity of the BHJ approach to exposure-robust inference is that only two types of controls are included in \mathbf{r}_i . First, all sources of shock-level confounding have to be captured by controls with a shift-share structure (i.e., \mathbf{Q}_i). Second, all other controls should not be asymptotically correlated with the instrument; these controls may, however, increase the asymptotic efficiency of the estimator. BHJ show that the standard errors are asymptotically conservative under a weaker assumption, which allows for controls in \mathbf{r}_i to be of the form $\sum_k s_{ik} \mathbf{p}_k + u_i$ where \mathbf{p}_k are unobserved shock-level confounders and u_i is noise. AKM

¹⁵ AKM consider other extensions, including confidence intervals which impose a null hypothesis on β and settings with heterogeneous treatment effects (but assuming that each observation is asymptotically exposed to only one shock).

¹⁶ Note that (4.4) does not include an intercept, but (4.6) shows that including it does not change the estimate so long as the shares are complete or the sum of shares is controlled for.

provide a variance estimator that is asymptotically exact when controls have this ‘approximate’ shift-share structure. Asymptotically exact shift-share inference with general control vectors (which are necessary for identification and do not have this form) remains an open problem.

REMARK 4.3. The asymptotic inference results of AKM and BHJ are also useful to conduct regression-based falsification tests of shock orthogonality, and for verifying the strength of the instrument. Falsification tests for orthogonality can be conducted by regressing any proxy for the unobserved error (such as a pre-trend in the outcome) on the shift-share instrument while controlling for \mathbf{Q}_i . For asymptotically valid tests, the coefficient on the instrument can be computed by the shock-level regression (4.6) with exposure-robust standard errors. Similarly, a valid first-stage F-statistic can be computed by translating the regression of x_i on z_i and \mathbf{Q}_i to the shock level.

REMARK 4.4. Note that reduced-form regressions on z_i are still IV regressions at the shock level, which instrument \bar{z}_k with g_k controlling for \mathbf{q}_k .

4.3. Randomization inference

The asymptotic results of AKM and BHJ are specific to linear shift-share variables with a large number of shocks. We now discuss an alternative RI approach, which BH propose for other formula instruments or settings with only a small number of shocks. This approach may be natural when Assumption 2.2 is made, because specified shock counterfactuals immediately yield confidence intervals that are exact in finite samples of observations and shocks—albeit only under the constant-effect model (2.1).¹⁷

Formally, BH propose to construct confidence intervals for β using the sample covariance between \tilde{z}_i and the residual as a randomization test statistic:

$$T = \frac{1}{N} \sum_i (f_i(\mathbf{s}, \mathbf{g}) - \mu_i)(y_i - bx_i), \quad (4.7)$$

where b is a candidate parameter value. Under the null hypothesis of $\beta = b$ and Assumption 2.1, $y_i - bx_i = \varepsilon_i$, and the distribution of T conditional on $\boldsymbol{\varepsilon}$ and \mathbf{w} is implied by the shock assignment process $G(\mathbf{g} | \mathbf{w})$. One may simulate this distribution under Assumption 2.2, by redrawing the shocks and recomputing T . The null is rejected if the original value of T is far in the tails of the simulated distribution.¹⁸ The confidence interval for β is then obtained by inversion of such tests, i.e., by collecting all b that are not rejected. These intervals have correct size, both conditionally on $(\boldsymbol{\varepsilon}, \mathbf{w})$ and unconditionally.

While any statistic $T(\mathbf{g}, \mathbf{y} - b\mathbf{x}, \mathbf{w})$ can in principle be used in a similar procedure, the choice of (4.7) is motivated by its close connection to the IV estimator $\hat{\beta}$. Specifically, Borusyak and Hull (2021b) show that $\hat{\beta}$ is the Hodges–Lehman estimator corresponding to T (Hodges and Lehmann, 1963; Rosenbaum, 2002) and provide guarantees for the consistency of the randomization tests when $\hat{\beta}$ is consistent.

REMARK 4.5. RI-based confidence intervals can be useful even when few shocks are observed in the data, making the asymptotic approach inapplicable even with linear formulas. Consider

¹⁷ RI methods are typically used to test the ‘sharp null’ of zero treatment effects for all observations; valid inference under weaker nulls of (say) no average effect is a more challenging problem in general. See, e.g., Chung and Romano (2013) for an approach motivated by this challenge.

¹⁸ Note that, while we recentre T in equation (4.7), this is not necessary since the term $\frac{1}{N} \sum_i \mu_i (y_i - bx_i)$ shifts the observed and re-randomized T equally.

for example the single earthquake studied by Carvalho et al. (2021). Even without spillovers, it is not possible to consistently estimate the effect of earthquakes if only one region is shocked. Yet RI remains informative in this case, as it can produce confidence intervals of finite length. For example, if the true effect is zero, it is unlikely that unobserved shocks hit exactly the same region where the earthquake randomly happened and the randomization test could reject $\beta = 0$.

REMARK 4.6. As with asymptotic shift-share inference, RI-based confidence intervals can be used for falsification tests of the identifying assumptions. A falsification test of Assumption 2.1 is obtained by checking that the sample covariance between \tilde{z}_i and a proxy for ε_i is close to zero (or, more precisely, is not in the tail of the distribution of this covariance across counterfactual shocks). Similarly, to test the correct specification of the shock assignment process (Assumption 2.2) and the expected instrument implied by it, one can check that \tilde{z}_i is not correlated with any pre-specified functions of \mathbf{w} . While the assumption of constant effects required for the validity of RI-based confidence intervals can be restrictive in practice, it is not a problem for falsification tests where the true effect should be zero for all i .¹⁹

5. EXTENSIONS

Identification with heterogeneous effects Borusyak and Hull (2021b, appendix C.1) and BHH (appendix A.1) show that classic results on linear IV identification with heterogeneous treatment effects (e.g., Imbens and Angrist, 1994) extend to formula instruments provided μ_i is recentred by or controlled for and a version of the usual first-stage monotonicity condition holds. Specifically, they show that when an exclusion restriction holds, but the outcome model features nonlinear and heterogeneous effects of the treatment (as opposed to the linear specification in (2.1)), the recentred IV estimand is a convex weighted average of the marginal effects of x_i on y_i . For example, in the reduced-form case with linear effect heterogeneity, i.e., $y_i = \beta_i z_i + \varepsilon_i$, recentred IV identifies $\mathbb{E} \left[\frac{1}{N} \sum_i \omega_i \beta_i \right]$ where the weights ω_i are proportional to $\text{Var} [z_i | \mathbf{w}]$. This conditional variance can be computed by a simulation analogously to μ_i , allowing researchers to study the implied weights of the estimand. When $\text{Var} [z_i | \mathbf{w}]$ is bounded away from zero, a recentred IV estimator that inversely weights by it identifies the unweighted average causal effect $\mathbb{E} \left[\frac{1}{N} \sum_i \beta_i \right]$.

In the IV case, Borusyak and Hull (2021b) generalize Imbens and Angrist (1994) by allowing the regression of x_i on \tilde{z}_i to be noncausal, in the sense that the first-stage coefficient on \tilde{z}_i does not capture the causal effect of shocks \mathbf{g} on x_i . This extension is useful for formula instruments, since \tilde{z}_i is a constructed variable that does not usually correspond to a real economic object. Thus, it is more appropriate to specify the potential values of x_i in terms of the primitive shocks \mathbf{g} . BH show that $\hat{\beta}$ retains its interpretation as a convexly-weighted average of causal effects under an appropriate monotonicity assumption, even if \tilde{z}_i does not capture all of the ways for how the shocks affect the treatment. In the context of linear shift-share instruments, BHH consider the case where use exposure shares s_{ik} are misspecified, such that they do not correlate imperfectly with the true shares through which shocks g_k affect x_i . Nevertheless, a causally interpretable IV estimand is guaranteed as long as x_i is correctly specified, the s_{ik} are nonnegative, the shocks are weakly positively correlated, and the true effects of shocks on each treatment are monotone.

These results can be contrasted to those of Blandhol et al. (2022), who show the importance of including controls flexibly for IV regressions to be ‘weakly causal’ (i.e., to identify a convex

¹⁹ Testing the first-stage relationship between x_i and \tilde{z}_i is similarly straightforward, though, as noted by Imbens and Rosenbaum (2005), there is no weak instrument problem for RI-based confidence intervals.

average of heterogeneous effects). Here design knowledge yields the single control μ_i needed for the IV estimand to be weakly causal; other low-dimensional $\mathbf{r}_i(\mathbf{w})$ can again increase precision while not affecting this estimand.

Multiple treatments and instruments The identification results of Section 3 extend immediately to cases with multiple formula instruments $f_{\ell i}(\mathbf{s}, \mathbf{g})$, $\ell = 1, \dots, L$, for either a single treatment x_i or multiple treatments x_{pi} , $p = 1, \dots, P \leq L$. As long as each expected instrument $\mathbb{E}[f_{\ell i}(\mathbf{s}, \mathbf{g}) | \mathbf{w}]$ is adjusted for via recentring or controlling, the corresponding orthogonality conditions hold—such that $(\beta_1, \dots, \beta_P)$ is identified under a standard rank condition.

The case of multiple treatments is particularly useful in network settings where it allows, including a node's own shock along with a treatment capturing network spillovers; e.g., student i 's own deworming status and the number of her neighbours who have been dewormed, in the Miguel and Kremer (2004) setting. Notably, the same set of exogenous shocks allows for identification of the effects of both treatments, via different exposure mappings.²⁰

The case of multiple instruments is also relevant when the researcher has access to multiple sets of exogenous shocks. BHJ illustrate this scenario in the ADH China shock setting. While ADH measure industry supply shocks as average Chinese import growth across eight non-US countries, in principle the shocks from each individual country may be considered as-good-as-randomly assigned leading to multiple shift-share instruments for the regional exposure to import competition from China and thus overidentification of the parameter of interest. BHJ show how to extend the shock-level representation of the shift-share IV estimator when the exposure shares used to construct the instruments are the same, but the shocks differ. This yields results for both identification and asymptotic exposure-robust inference.

Other necessary controls The framework so far may seem unable to accommodate certain controls that may be deemed necessary for identification. Specifically, the identification results in Section 3.1 only allow for controls \mathbf{r}_i which are functions of \mathbf{w} . The purpose of these controls for identification is to absorb confounding variation in μ_i ; beyond that, they only play a role in estimation efficiency. In the linear shift-share case of $z_i = \sum_k s_{ik} g_k$, for example, this means all necessary controls must have a shift-share structure of the form $\sum_k s_{ik} q_k$.

One could imagine, however, other necessary controls that remove confounding variation in ε_i . Such controls can be accommodated by augmenting the causal model (2.1) to have $\varepsilon_i = \mathbf{c}'_i \lambda + \eta_i$ for an observed \mathbf{c}_i (which is not assumed conditionally independent of \mathbf{g}) and an error η_i defined to be orthogonal to \mathbf{c}_i (without loss of generality, since λ is not interpreted causally). Then

$$y_i = \beta x_i + \mathbf{c}'_i \lambda + \eta_i, \quad (5.1)$$

and the identification question shifts to the orthogonality of z_i and η_i . Modifying Assumption 2.1 to the conditional independence of $\boldsymbol{\eta} = (\eta_i)_{i=1}^N$ and \mathbf{g} given \mathbf{w} (and thus allowing violations of $\mathbf{g} \perp\!\!\!\perp \boldsymbol{\varepsilon} | \mathbf{w}$ via \mathbf{c}) and maintaining Assumption 2.2 and appropriate instrument relevance, we obtain identification in a recentred IV regression of y_i on x_i which instruments by \tilde{z}_i and controls for \mathbf{c}_i , because the orthogonality conditions $\mathbb{E}[\frac{1}{N} \sum_i \tilde{z}_i \eta_i] = \mathbb{E}[\frac{1}{N} \sum_i \mathbf{c}_i \eta_i] = 0$ hold. The same result holds for linear instruments with Assumption 2.3 modified to specify that the conditional mean of g_k given $(\boldsymbol{\eta}, \mathbf{w})$ is linear in \mathbf{q}_k . Alternatively, as before, μ_i can be controlled for along with \mathbf{c}_i with the original z_i used as the instrument.²¹

To make these ideas concrete, consider again the setting of Autor et al. (2013). Imagine that regions more exposed to supply shocks in China also have higher (lagged) local unemployment

²⁰ Constant effects play a bigger role with multiple treatments, however. See Bhuller and Sigstad (2023) and Goldsmith-Pinkham et al. (2022) for discussions in the context of IV and reduced-form regressions.

²¹ Consistency and inference with such controls may be more complicated (see Remark 4.2).

rates, due to some location-specific factors (e.g., local regulations). The local unemployment rate does not have a shift-share structure, as there is residual variation in unemployment rates across locations with the same industry composition. However, one can directly control for the (lagged) local unemployment rate for identification.

Panel data While we have indexed observations by i with cross-sectional data in mind, the results of Sections 2–4 extend to panel data where observations $i = (\ell, t)$ correspond to units ℓ in periods t and shocks can also vary over time. We make two remarks specific to panel data.

First, including unit fixed effects does not generally alleviate the need to adjust the formula instrument, unless the corresponding expected instrument is time-invariant. Consider the instrument $z_{\ell t} = f_{\ell t}(s_t, \mathbf{g}_t)$ with $\mathbf{g}_t \perp \boldsymbol{\varepsilon}_t \mid \mathbf{w}_t$ (here objects like \mathbf{v}_t , with just a t subscript, refer to the vector of $v_{\ell t}$). Then $\mu_{\ell t} = \mathbb{E}[z_{\ell t} \mid \mathbf{w}_t]$ is absorbed by the unit fixed effects under the very restrictive conditions that the $f_{\ell t}(\cdot)$ mapping, the value of \mathbf{w}_t , and the shock assignment process $G(\mathbf{g}_t \mid \mathbf{w}_t)$ are all time-invariant. In linear and nonlinear shift-share settings, for example, even if exposure shares are measured in a fixed pre-period for the entire panel, the shocks often vary systematically over time. Footnote 11 above illustrates this issue in the ADH setting where industry-level China shocks are higher in the 2000s than in the 1990s, making $\mu_{\ell t}$ time-varying.

Second, an advantage of panel data is that the number of cross-sectional shocks (and observations) need not be large; instead, a long time series of shocks g_t with weak serial dependence may also suffice for consistency of $\hat{\beta}$. BHJ formalize this insight for linear shift-share settings. Nunn and Qian (2014) is an empirical example of constructing shift-share instruments from purely time-series exogenous shocks (specifically, in total US foreign aid arising from wheat production in the US) with no cross-sectional variation.

Estimated shock assignment processes Assumption 2.2 requires complete knowledge of the shock assignment process, conditionally on \mathbf{w} . Appendix C.5 of Borusyak and Hull (2021b) shows that this assumption can be relaxed with the assignment process specified up to a vector of parameters, $G(\mathbf{g} \mid \mathbf{w}; \theta)$. If these parameters can be consistently estimated by $\hat{\theta}$ from shock-level data (e.g., by maximum likelihood) and some regularity conditions hold, adjusting the instrument by the estimated $\hat{\mu}_i = \int f_i(s, \mathbf{g}) dG(\mathbf{g} \mid \mathbf{w}; \hat{\theta})$ leads to consistent estimation of β . This generalizes Assumption 2.3 (which allowed the conditional shock mean to be parameterized) to nonlinear formula instruments. This approach is especially attractive when the shocks are binary (and mutually independent conditionally on \mathbf{w}) as a model the conditional mean, e.g., via logit or probit, yields the entire conditional distribution.

Inference on β will typically be affected by the first-step estimation of θ . Appendix C.5 of Borusyak and Hull (2021b) shows how the baseline BH RI procedure can be extended to build conservative finite-sample confidence intervals, drawing on the approach of Berger and Boos (1994). It further notes that there are cases in which exact confidence intervals can be constructed by using sufficient statistics that obviate the need to estimate θ . In some settings valid asymptotic inference may also be possible.

Most efficient recentered instruments The results of Sections 2–4 take the instrument construction $f_i(\cdot)$ as given, but different constructions are likely to yield estimates with different levels of precision in large samples. Borusyak and Hull (2021a) study the question of optimal instrument construction (under constant effects), building on Chamberlain (1987, 1992). In a class of regular recentered IV estimators, the asymptotically most efficient one can be described in a three step process: the best prediction of \mathbf{x} from (\mathbf{g}, \mathbf{w}) is constructed, recentered using the shock assignment process, and then adjusted for ‘heteroskedasticity’ (formally, weighted by the inverse of $\mathbb{E}[\boldsymbol{\varepsilon}\boldsymbol{\varepsilon}' \mid \mathbf{w}]$ which also depends on $\mathbb{E}[\boldsymbol{\varepsilon} \mid \mathbf{w}]$). The prediction and recentering steps can be seen to justify the empirical practice, noted in Remark 2.2, of forming instruments from the

structure of x_i by removing or replacing some endogenous components. The heteroskedasticity adjustment step is infeasible in non-*iid* data, but it is often ignored in practice even in settings where it would be feasible (see, e.g., Coussens and Spiess, 2021).

6. COMPARISONS WITH ALTERNATIVE APPROACHES

The design-based approach leverages knowledge of the formula construction $f_i(\cdot)$ to correct for potential confounding. A natural question is whether identification is possible without such knowledge, potentially under alternative assumptions. We answer this question by comparing the design-based approach to two alternative strategies for identification in linear models: those based on conventional assumptions of conditional treatment unconfoundedness, and those based on alternative restrictions on outcome unobservables (such as parallel trend assumptions).

For both comparisons we use the fact that, under Assumption 2.1, equation (2.1) can be written as a partially linear model in the spirit of Robinson (1988). Specifically, with $h_i(\mathbf{w}) = \mathbb{E}[\varepsilon_i | \mathbf{w}]$,

$$y_i = \beta x_i + h_i(\mathbf{w}) + \eta_i, \quad \mathbb{E}[\eta_i | \mathbf{g}, \mathbf{w}] = 0. \quad (6.1)$$

One approach to identification in partially linear models, advanced by Robins et al. (1992) and in line with a long literature on propensity score methods, adjusts x_i for its association with potential confounders in \mathbf{w} before estimating β . A different approach, considered by Goldsmith-Pinkham et al. (2020) for linear shift-share instruments, restricts the form of $h_i(\mathbf{w})$ in order to directly estimate (6.1). This section connects the design-based approach to formula instruments to the first strategy—showing that explicit use of the formula is essential in this case—before contrasting it with the alternative restrictions implied by the second strategy.

6.1. Conventional unconfoundedness

Partially linear models are typically considered in *iid* data where $x_i = g_i$ is assumed to be unconfounded given observation-specific \mathbf{w}_i . That is, with $h(\mathbf{w}_i) = \mathbb{E}[\varepsilon_i | \mathbf{w}_i]$,

$$y_i = \beta x_i + h(\mathbf{w}_i) + \eta_i, \quad \mathbb{E}[\eta_i | x_i, \mathbf{w}_i] = 0. \quad (6.2)$$

Based on a suggestion by Newey (1990), Robins et al. (1992) consider estimation of β when a researcher has ‘sharper’ information on the relationship between the observed x_i and \mathbf{w}_i than on the form of unobservables as captured by $h(\mathbf{w}_i)$. This motivates their ‘E-estimator’:

$$\hat{\beta}_E = \frac{\sum_i y_i (x_i - \mathbb{E}[x_i | \mathbf{w}_i])}{\sum_i x_i (x_i - \mathbb{E}[x_i | \mathbf{w}_i])}. \quad (6.3)$$

The recentred IV estimator of BH can be seen to generalize the E-estimator to the setting where $x_i = f_i(\mathbf{s}, \mathbf{g})$ is constructed by a formula from common shocks. The $\mathbb{E}[x_i | \mathbf{w}_i]$ term becomes μ_i in this case. A further generalization is obtained by letting $\tilde{z}_i = z_i - \mu_i$ instrument for a different treatment x_i , now with $z_i = f_i(\mathbf{s}, \mathbf{g})$ and $\mu_i = \mathbb{E}[z_i | \mathbf{w}_i]$.²²

A central difference in the BHJ/BH settings, where the treatment or instrument is constructed from common shocks, is that estimation of $\mu_i = \mathbb{E}[f_i(\mathbf{s}, \mathbf{g}) | \mathbf{w}]$ via some first-stage procedure

²² When z_i is binary, μ_i can be seen to generalize the propensity score of Rosenbaum and Rubin (1983) to the formula instrument setting. Propensity scores are typically defined in settings with randomly sampled data and a conditionally unconfounded treatment, and used in weighting or matching estimators rather than with regression. Appendix C.1 of Borusyak and Hull (2021b) shows how such estimators can be adapted to formula instruments.

is not possible without using the formula. In the conventional setting Robins et al. (1992) study, $\mathbb{E}[x_i | \mathbf{w}_i]$ may be learned nonparametrically from the conditional distribution of observation-specific shocks g_i given the confounders \mathbf{w}_i . But when all observations i are exposed to the *same* vector of shocks \mathbf{g} , and only one realization of this vector is observed, such nonparametric estimation is impossible. Indeed, in conventional settings there are many observations with similar \mathbf{w}_i , but different values of x_i , such that $\mathbb{E}[x_i | \mathbf{w}_i]$ can be learned. In contrast, for the linear shift-share instruments there is no cross-sectional variation in $\sum_k s_{ik} g_k$ conditionally on s_i , while variation across possible realizations of \mathbf{g} vectors is not observed. This problem is distinct from challenges related to high-dimensional confounding, which can be solved by modern machine learning methods (e.g., Chernozhukov et al., 2018).²³

Leveraging the knowledge contained in the formula, however, allows the researcher to circumvent the common shocks issue. This is achieved in two steps: by first specifying or estimating the assignment process for shocks, similar to the design-based approach of Robins et al. (1992), and then ‘translating’ it to the level at which observations and treatment are observed. The shock-level equivalence results of BHJ make this translation especially clear in the linear shift-share case.

6.2. Outcome model restrictions

An alternative approach to estimating β in the partially linear model (6.1) restricts the unobserved outcome error ε_i , without restricting the assignment process of conditionally unconfounded shocks. Specifically, one posits a model for $h_i(\mathbf{w})$ and jointly estimates it along with β . Specification of the shock design plays no role in this strategy. In fact, the strategy is coherent when the shocks are considered nonrandom, in which case Assumptions 2.1 and 2.2 hold trivially. More generally, consider a linear outcome model replacing Assumptions 2.1 and 2.2 (or 2.3):

ASSUMPTION 6.1. (*Outcome model*) *There exists unknown γ such that $\mathbb{E}[\varepsilon_i | \mathbf{g}, \mathbf{w}] = \mathbf{q}'_i \gamma$ for observed \mathbf{q}_i included in \mathbf{w} and for all i .*

Assumption 6.1 requires a weaker version of Assumption 2.1, $\mathbb{E}[\varepsilon_i | \mathbf{g}, \mathbf{w}] = \mathbb{E}[\varepsilon_i | \mathbf{w}]$, which is trivially satisfied when the shocks are nonrandom.

A concrete example of Assumption 6.1 in the linear shift-share setting comes from Goldsmith-Pinkham et al. (2020). Conditioning on the shocks \mathbf{g} , they assume $\boldsymbol{\varepsilon}$ is mean-independent of the shares s after partialling out a vector of controls. Then $z_i = \sum_k s_{ik} g_k$ satisfies orthogonality. Justification for the key mean-independence assumption is similar to that underlying conventional ‘parallel trends’ restrictions in difference-in-differences estimation and related strategies for causal inference in panel data. Indeed, shift-share regressions are often specified in first-differences, in which case $\mathbb{E}[\varepsilon_i | s] = 0$ amounts to an assumption that observations with different exposure shares would have been on similar outcome trends if not for a change in x_i . Assumption 6.1 relaxes this condition by allowing for time-varying controls \mathbf{q}_i and stochastic shocks.

Assumption 6.1 is very powerful; this can be seen from the fact that it justifies the orthogonality of *any* formula instrument $f_i(s, \mathbf{g})$ when including the same set of controls \mathbf{q}_i . Indeed, in the Goldsmith-Pinkham et al. (2020) case, any function of s_i (including individual s_{ik}) satisfies instrument orthogonality. In this sense, identification strategies based on such models for ε_i can be said to not leverage the specific formula construction. Restrictions on ε_i may be especially

²³ Consider, for instance, a shift-share instrument z_i where $\mathbb{E}[g_k | s] = 0$ for all but a small set of $k \in \mathcal{K}$. The expected instrument $\mu(s_i)$ is then sparse in s_i , suggesting one might estimate it via least absolute shrinkage and selection operator (LASSO) methods. However, this intuition is misleading: the observed z_i are based on a single draw of each g_k , which may be far from zero even for $k \in \mathcal{K}$, giving no hope to identify \mathcal{K} or $\mu(\cdot)$ even as $N \rightarrow \infty$.

valuable when design-based estimation is imprecise (e.g., because the shocks are too few or too mutually correlated), if credible assumptions on their assignment process are lacking, or if the shocks are unobserved.

The cost of this alternative model-based approach is that Assumption 6.1 can be very restrictive. In particular, BBJ show that in the shift-share setting the Goldsmith-Pinkham et al. (2020) approach generally fails when there are *any* unobserved shocks ν_k that affect the outcomes through the same (or correlated) exposure shares; e.g., when

$$\varepsilon_i = \sum_k s_{ik} \nu_k + \tilde{\varepsilon}_i,$$

for some idiosyncratic $\tilde{\varepsilon}_i$.²⁴ Ruling out such ν_k is implausible in applications where g_k are ‘specific’ shocks that are of interest, while other unobserved shocks varying at the same level are likely.²⁵ For instance, in the ADH China shock setting, Assumption 6.1 is generally violated if there are any unobserved industry shocks affecting the outcome of interest—such as those arising from automation or changing tastes. Indeed, Goldsmith-Pinkham et al. (2020) find evidence against their ‘exogenous shares’ assumption in the ADH setting.

Models for ε_i can more generally be unsatisfying for the same reason that motivates Robins et al. (1992): one may have limited information about the ‘right’ form of $h_i(\mathbf{w})$, particularly since ε_i is unobserved. It may be unclear, for example, whether linearity of the unobservables more plausibly holds when the outcome is specified in levels vs. in logs (while both generally cannot hold; see Roth and Sant’Anna, 2023). Moreover, it may be challenging to pick the appropriate observation-specific features \mathbf{q}_i in the kinds of non-*iid* settings where formula instruments are deployed, such as with spatial or network data. The challenge is amplified in the presence of heterogeneous effects where a correct model of untreated potential outcomes may not suffice for the regression to estimate a convex average of treatment effects.²⁶ The design-based approach helps meet these challenges, as the formula for z_i and specification of the distribution of observed shocks guide the choice of controls via μ_i and, as noted in Section 5, the design-based approach guarantees a causally interpretable estimand even with heterogeneous effects under a first-stage monotonicity condition.

7. CONCLUSION

The design-based approach to formula instruments can be seen to bring new insights to longstanding identification strategies in economics (shift-share IV, in particular) while also pointing a way forward for novel econometric strategies that leverage more complex instrument constructions. An ever-increasing richness of data, sophistication of economic models, and creativity in discovering plausibly exogenous shocks are likely to yield many new settings where Assumptions 2.1,

²⁴ If $\mathbb{E}[\nu_k | s] = 0$ for all k , Assumption 6.1 technically holds. But the presence of the ν_k shocks may still affect the consistency of the nonrecentered shift-share IV estimator; see appendix A.2 of BBJ.

²⁵ Conversely, BBJ argue this problem may be less concerning when the shares are ‘tailored’ to the treatment of interest such that they are unlikely to mediate other unobserved shocks. One candidate for this is Card (2009), who uses lagged immigrant enclave shares and builds on an earlier difference-in-difference strategy in Card (1990).

²⁶ This is now well known in the context of staggered difference-in-differences designs—e.g., de Chaisemartin and D’Haultfœuille (2020), Borusyak et al. (2023). Restrictions on treatment effects can only be avoided in a restricted class of settings where ‘imputation’ estimators are feasible—e.g., Wooldridge (2021). Borusyak et al. (2023), Borusyak and Hull (2024) link this literature on nonconvex ‘ex post’ weighting schemes to the convex ‘ex ante’ weight results of BBJ and BH, and to design-based IV and OLS estimation more broadly.

2.2, or 2.3 credibly hold. A focus on shock design in such settings can allow researchers to avoid extraneous or undesirably strong assumptions on how model unobservables relate to predetermined observables, provided the formula construction is leveraged appropriately for identification and inference.

Several open paths remain in this agenda. First, asymptotic inference results are only available in some special cases. It would be valuable to develop asymptotically valid inference techniques outside the linear shift-share case, and for shift-share instruments with unrestricted control vectors, as well as to extend results on RI with heterogeneous treatment effects (e.g., Chung and Romano, 2013) to formula instruments. Second, it would be interesting to extend recent double/debiased machine learning results (e.g., Chernozhukov et al., 2018) to shift-share instruments—allowing the vector \mathbf{q}_k of shock-level confounders to be high-dimensional. Third, it would be useful to characterize properties of the recentring approach when the shock assignment process is estimated in a flexible way; e.g., by nonparametric estimation of the distribution of $g_k | \mathbf{q}_k$ when (g_k, \mathbf{q}_k) are *iid*. Fourth, the estimand of the recentred IV procedure is not known when the exclusion restriction is violated. For instance, is recentring guaranteed to reduce bias in some situations even if the network spillovers embedded in x_i are misspecified? Finally, open questions remain on how recentring can relax the assumptions of more elaborate econometric methods, such as structural models for differentiated product demand. On these and other issues we expect many interesting developments in the growing literature on formula instruments.

ACKNOWLEDGEMENT

Borusyak acknowledges financial support from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement number 949995).

REFERENCES

- Abadie, A., S. Athey, G. W. Imbens and J. M. Wooldridge (2020). Sampling-based vs. design-based uncertainty in regression analysis. *Econometrica* 88(1), 265–96.
- Adão, R., M. Kolesár and E. Morales (2019). Shift-share designs: theory and inference. *Quarterly Journal of Economics* 134(4), 1949–2010.
- Aghion, P., C. Antonin, S. Bunel and X. Jaravel (2022). Modern manufacturing capital, labor demand, and product market dynamics: Evidence from France. Discussion paper, Centre for Economic Performance.
- Aronow, P. M. and C. Samii (2017). Estimating average causal effects under general interference, with application to a social network experiment. *Annals of Applied Statistics* 11(4), 1912–47.
- Autor, D. H., D. Dorn and G. H. Hanson (2013). The China syndrome: local labor market impacts of import competition in the United States. *American Economic Review* 103(6), 2121–68.
- Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W. E. Upjohn Institute for Employment Research.
- Basso, G. and G. Peri (2015). The association between immigration and labor market outcomes in the United States. IZA Discussion Papers 9436, Institute of Labor Economics (IZA).
- Berger, R. L. and D. D. Boos (1994). P values maximized over a confidence set for the nuisance parameter. *Journal of the American Statistical Association* 89(427), 1012–6.
- Berman, N., A. Berthou and J. Héricourt (2015). Export dynamics and sales at home. *Journal of International Economics* 96(2), 298–310.

- Berthou, A., J. J.-H. Chung, K. Manova and C. Sandoz Dit Bragard (2019). Productivity, (mis)allocation and trade. Working Paper 749, Banque de France.
- Bhuller, M. and H. Sigstad (2023). 2SLS with Multiple Treatments. Working paper, arXiv:2205.07836.
- Blanchard, O. J. and L. F. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity*, 1, 1–75.
- Blandhol, C., J. Bonney, M. Mogstad and A. Torgovitsky (2022). When is TSLS Actually LATE?. NBER Working Paper 29709, National Bureau of Economic Research.
- Borusyak, K. and P. Hull (2021a). Efficient estimation with nonrandom exposure to exogenous shocks. Working paper.
- Borusyak, K. and P. Hull (2021b). Nonrandom exposure to exogenous shocks: theory and applications. Working paper.
- Borusyak, K. and P. Hull (2023). Nonrandom exposure to exogenous shocks. *Econometrica* 91(6), 2155–85.
- Borusyak, K. and P. Hull (2024). Negative weights are no concern in design-based specifications. *AEA Papers and Proceedings*. Forthcoming.
- Borusyak, K., P. Hull and X. Jaravel (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies* 89(1), 181–213.
- Borusyak, K., X. Jaravel and J. Spiess (2023). Revisiting event study designs: robust and efficient estimation. *Review of Economic Studies*. Forthcoming.
- Boustan, L., F. Ferreira, H. Winkler and E. M. Zolt (2013). The effect of rising income inequality on taxation and public expenditures: evidence from US municipalities and school districts, 1970–2000. *Review of Economics and Statistics* 95(4), 1291–302.
- Card, D. (1990). The impact of the marinel boatlift on the Miami labor market. *ILR Review* 43(2), 245–57.
- Card, D. (2009). Immigration and inequality. *American Economic Review: Papers & Proceedings* 99(2), 1–21.
- Carvalho, V. M., M. Nirei, Y. U. Saito and A. Tahbaz-Salehi (2021). Supply chain disruptions: evidence from the great east Japan earthquake. *Quarterly Journal of Economics* 136(2), 1255–321.
- Cattaneo, M. D., B. R. Frandsen and R. Titiunik (2015). Randomization inference in the regression discontinuity design: an application to party advantages in the US Senate. *Journal of Causal Inference* 3(1), 1–24.
- Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics* 34(3), 305–34.
- Chamberlain, G. (1992). Efficiency bounds for semiparametric regression. *Econometrica* 60(3), 567–96.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal* 21(1), C1–C68.
- Chung, E. and J. P. Romano (2013). Exact and asymptotically robust permutation tests. *Annals of Statistics* 41(2), 484–507.
- Costa, R., S. Dhingra and S. Machin (2019). Trade and worker deskilling. NBER Working Paper 25919.
- Coussens, S. and J. Spiess (2021). Improving inference from simple instruments through compliance estimation. Working paper, arXiv:2108.03726.
- Currie, J. and J. Gruber (1996). Health insurance eligibility, utilization of medical care, and child health. *Quarterly Journal of Economics* 111(2), 431–66.
- de Chaisemartin, C. and X. D’Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110(9), 2964–96.
- Diamond, R. (2016). The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000. *American Economic Review* 106(3), 479–524.

- Donaldson, D. and R. Hornbeck (2016). Railroads and American economic growth: a ‘market access’ approach. *Quarterly Journal of Economics* 131(2), 799–858.
- Freedman, D. A. (2008). On regression adjustments to experimental data. *Advances in Applied Mathematics* 40(2), 180–93.
- Goldsmith-Pinkham, P. and G. Imbens (2013). Social networks and the identification of peer effects. *Journal of Business and Economic Statistics* 31(3), 253–64.
- Goldsmith-Pinkham, P., P. Hull and M. Kolesár (2022). Contamination bias in linear regressions. NBER Working Paper 30108, National Bureau of Economic Research.
- Goldsmith-Pinkham, P., I. Sorkin and H. Swift (2020). Bartik instruments: what, when, why, and how. *American Economic Review* 110(8), 2586–624.
- Greenstone, M., A. Mas and H.-L. Nguyen (2020). Do credit market shocks affect the real economy? Quasi-experimental evidence from the great recession and ‘normal’ economic times. *American Economic Journal: Economic Policy*, 12(1), 200–25.
- Hodges, J. J. and E. L. Lehmann (1963). Estimates of location based on rank tests. *Annals of Mathematical Statistics* 34(2), 598–611.
- Hummels, D., R. Jorgensen, J. Munch and C. Xiang (2014). The wage effects of offshoring: evidence from Danish matched worker–firm data. *American Economic Review* 104(6), 1597–629.
- Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of local average treatment effects. *Econometrica* 62(2), 467.
- Imbens, G. W. and P. R. Rosenbaum (2005). Robust, accurate confidence intervals with a weak instrument: quarter of birth and education. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 168(1), 109–26.
- Jaravel, X. (2019). The unequal gains from product innovations: evidence from the US retail sector. *Quarterly Journal of Economics* 134(2), 715–83.
- Lee, D. S. (2008). Randomized experiments from nonrandom selection in US House elections. *Journal of Econometrics* 142(2), 675–97.
- Miguel, E. and M. Kremer (2004). Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72(1), 159–217.
- Nakamura, E. and J. Steinsson (2014). Fiscal stimulus in a monetary union: evidence from US regions. *American Economic Review* 104(3), 753–92.
- Newey, W. K. (1990). Efficient instrumental variables estimation of nonlinear models. *Econometrica* 58(4), 809–37.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments: essay on principles. *Annals of Agricultural Sciences*, 1–51.
- Nunn, N. and N. Qian (2014). US food aid and civil conflict. *American Economic Review* 104(6), 1630–66.
- Oberfield, E. and D. Raval (2021). Micro data and macro technology. *Econometrica* 89(2), 703–32.
- Peri, G., K. Shih and C. Sparber (2016). STEM workers, H-1B visas, and productivity in US cities. *Journal of Labor Economics* 49(3), 277–307.
- Robins, J. M., S. D. Mark and W. K. Newey (1992). Estimating exposure effects by modelling the expectation of exposure conditional on confounders. *Biometrics* 48(2), 479–95.
- Robinson, P. (1988). Root-N-consistent semiparametric regression. *Econometrica* 56(4), 931–54.
- Rosenbaum, P. R. (2002). Covariance adjustment in randomized experiments and observational studies. *Statistical Science* 17(3), 286–327.
- Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), 41–55.
- Roth, J. and P. H. Sant’Anna (2023). When is parallel trends sensitive to functional form? *Econometrica* 91(2), 737–47.

- Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics* 125(3), 1253–96.
- Wooldridge, J. M. (2021). Two-way fixed effects, the two-way Mundlak regression, and event study estimators. Unpublished Working paper.
- Xu, C. (2022). Reshaping global trade: the immediate and long-run effects of bank failures. *Quarterly Journal of Economics* 137(4), 2107–61.

Managing editor Victor Chernozhukov handled this manuscript.