



HHS Public Access

Author manuscript

Econometrica. Author manuscript; available in PMC 2024 November 01.

Published in final edited form as:

Econometrica. 2023 November ; 91(6): 2155–2185. doi:10.3982/ecta19367.

Nonrandom Exposure to Exogenous Shocks

Kirill Borusyak,

UC Berkeley and CEPR

Peter Hull*

Brown and NBER

Abstract

We develop a new approach to estimating the causal effects of treatments or instruments that combine multiple sources of variation according to a known formula. Examples include treatments capturing spillovers in social or transportation networks and simulated instruments for policy eligibility. We show how exogenous shocks to some, but not all, determinants of such variables can be leveraged while avoiding omitted variables bias. Our solution involves specifying counterfactual shocks that may as well have been realized and adjusting for a summary measure of non-randomness in shock exposure: the average treatment (or instrument) across shock counterfactuals. We use this approach to address bias when estimating employment effects of market access growth from Chinese high-speed rail construction.

1 Introduction

Many questions in economics involve the causal effects of treatments which are computed from multiple sources of variation, and sometimes observed at different “levels,” according to a known formula. Consider three examples. First, when estimating spillovers from a randomized intervention, one might count the number of an individual’s neighbors who were selected for the intervention. This spillover treatment combines variation in who was selected with variation in who neighbors whom. Second, in studies of transportation infrastructure effects, one might measure the growth of regional market access: a treatment computed from the location and timing of transportation upgrades and the spatial distribution of economic activity in a country. A third example is a treatment capturing individual eligibility for a public program, such as Medicaid, which is jointly determined by the eligibility policy in the individual’s state and her household’s demographics and income.¹

*Contact: peter_hull@brown.edu.

We are grateful to Rodrigo Adão, Gabriel Ahlfeldt, Dmitry Arkhangelsky, Nate Baum-Snow, Sophie Calder-Wang, Vasco Carvalho, Gabriel Chodorow-Reich, Dave Donaldson, Raffaella Giacomini, Paul Goldsmith-Pinkham, Richard Hornbeck, Kilian Huber, Xavier Jaravel, Tetsuya Kaji, Vishal Kamat, Michal Kolesár, Whitney Newey, Aureo de Paula, Andrés Rodríguez-Clare, Jonathan Roth, Cyrus Samii, Jesse Shapiro, Ben Sommers, Chenzi Xu, and numerous seminar participants for helpful comments. Ruixue Li, Elise Parrish, and Steven Shi provided outstanding research assistance.

¹Examples of these three settings include Miguel and Kremer (2004), Donaldson and Hornbeck (2016), and Currie and Gruber (1996), respectively. Our working paper (Borusyak and Hull, 2021) discusses other common treatments and instruments nested in our framework: linear and nonlinear shift-share variables, model-implied optimal instruments, instruments based on centralized school assignment mechanisms, “free-space” instruments for access to mass media, and variables leveraging weather shocks.

This paper develops a new approach to estimating the effects of such composite variables when some, but not all, of their determinants are generated by a true or natural experiment. We ask, for example, how one can estimate market access effects by leveraging the timing of new railroad line construction as exogenous shocks, when the other determinants of market access (such as the pre-determined location of large markets and planned lines) are non-random.

We first show that omitted variable bias (OVB) may confound conventional regression approaches in such settings. Bias arises from different observations receiving systematically different values of the treatment because of their individual non-random “exposure” to the exogenous shocks. For example, even when construction is delayed for a random set of lines, regions that are economically or geographically more central will tend to see a larger growth in market access because they are closer to a typical potential line (and thus closer to a typical *constructed* line). Regression estimation of market access effects then fails without an additional assumption on the exogeneity of economic geography: that more exposed (e.g., central) regions do not differ in their relevant unobservables, such as changes in local productivity or amenities. Intuitively, randomizing transportation upgrades does not randomize the market access growth generated by them.

Our solution to the OVB challenge is based on the specification of counterfactual exogenous shocks that might as well have been realized. This approach views the observed shocks as one realization of some data-generating process—what we call the shock *assignment process*—which can be simulated to obtain counterfactuals. In a true experiment, the shock assignment process is given by the randomization protocol. In natural experiments, shock counterfactuals make explicit the contrasts which the researcher wishes to leverage, for instance by specifying permutations of the shocks that were as likely to have occurred. For example, if line construction delays are considered as-good-as-random, one might produce counterfactual network maps by randomly exchanging the lines which were completed earlier and later.

Valid shock counterfactuals can be used to avoid OVB by a “recentering” procedure which involves measuring and appropriately adjusting for a single confounder: the *expected treatment*. To do so, a researcher draws counterfactual shocks from the assignment process and recomputes the instrument many times. Then, for each observation, the treatment is averaged across these many draws to obtain the expected treatment. Finally, the expected treatment is subtracted from the realized treatment to obtain the *recentered treatment*. We show that using this recentered treatment as an instrument for the realized treatment removes the bias from non-random shock exposure. Intuitively, observations only get high vs. low values of the recentered treatment because the observed shocks were drawn instead of the counterfactuals, which is assumed to happen by chance. For example, when the expected treatment is constructed by permuting the timing of new line construction, regressions that instrument with recentered market access growth compare regions which received higher vs. lower market access growth because proximate lines were constructed early vs. late, and not because of the economic geography. Another closely related solution to OVB is to include the expected treatment as a control in the regression of an outcome on the realized treatment;

this can be viewed as recentering the treatment while also removing some residual variation in the outcome, in a control function approach.²

This approach to causal inference with composite variables, in which some determinants are labeled as exogenous and characterized by an assignment process, can be seen as formalizing the natural experiment of interest and bringing composite variables to familiar econometric territory.³ Indeed, the conditions we impose on the exogenous shocks are similar to those which might be used if the shocks were directly used as treatments: e.g., if shocks to the timing of railroad line upgrades were used in a regression of outcomes defined at the “level” of those lines. Recentering ensures identification from the natural experiment, even when the regression is estimated at a different level (e.g., across regions instead of lines).

Our framework further allows the treatment to have endogenous or unobserved determinants. In this case one may construct candidate composite instruments based on the treatment’s exogenous and predetermined components. The same OVB problem arises in this instrumental variable (IV) case, and it can again be solved by recentering the candidate instrument by its expectation over the shock assignment process. Controlling for the expected instrument is again another solution.

We establish several attractive properties of the recentering approach, beyond our primary results on OVB. First, recentered estimators are consistent provided the exogenous shocks induce sufficient cross-sectional variation in the instrument and treatment—regardless of the correlation structure of unobservables. Second, shock counterfactuals can be used for exact finite-sample inference and specification tests via randomization inference (RI). Finally, while our consistency and RI results rely on an assumption of constant treatment effects, recentered IV estimators generally capture a convex average of heterogeneous effects under a natural first-stage monotonicity condition.

We apply this framework to estimate the employment effects of market access (MA) growth due to new high-speed railway system in China. We show how recentering can help leverage variation in the timing of transportation upgrades to purge OVB. Simple regressions of employment growth on MA growth suggest a large and statistically significant effect, which is only partially reduced by conventional geography-based controls. But this effect is eliminated when we adjust for expected MA growth, measured by permuting constructed HSR lines with similar ones that were planned but not built. The unadjusted estimates thus reflect the fact that employment grew in regions which were more exposed to planned high-speed rail construction, whether or not construction actually occurred.

²While recentering is the key step that removes OVB, removing residual variation is likely to increase the efficiency of estimation in large samples. We give practical recommendations for each adjustment in the paper’s conclusion.

³Our approach is “design-based,” in that identification is achieved by specifying the assignment process of some observed shocks (see, e.g., Lee (2008), Athey and Imbens (2022), Shaikh and Toulis (2021), and de Chaisemartin and Behaghel (2020)). This strategy for analyzing observational data builds on a long tradition in the analysis of randomized experiments, going back to Neyman (1923). It contrasts with other identification strategies that instead model the residual determinants of the outcome, such as difference-in-difference strategies (e.g. de Chaisemartin and D’Haultfoeuille (2020) and Athey et al. (2021)) or fully-specified structural models.

Econometrically, expected treatment and instrument adjustment is similar to propensity score methods for removing OVB (Rosenbaum and Rubin 1983), with two key differences. First, we propose using the structure of composite treatments and instruments to compute their expectation from more primitive assumptions on the assignment process for exogenous shocks. This approach is similar to how Borusyak et al. (2022) and Aronow and Samii (2017) address OVB when using linear shift-share instruments and network treatments, respectively. It differs from conventional methods of directly estimating propensity scores; such methods are typically infeasible in the settings we consider because the exposure to exogenous shocks is intractably high-dimensional. Second, our regression-based adjustment differs from conventional approaches of weighting by or matching on propensity scores.⁴ Regression adjustment is more popular in applied research, avoids practical issues of limited overlap (due to, e.g., propensity scores that are close to zero or one), does not require the treatments or instruments to be binary, and is natural for estimating constant structural parameters or convex averages of heterogeneous treatment effects.

The remainder of this paper is organized as follows. The next section motivates our analysis with three examples related to network spillovers, market access effects, and Medicaid eligibility effects. Section 3 develops our general framework and results. Section 4 presents our application, and Section 5 concludes. Additional results and extensions are given in an earlier working paper, Borusyak and Hull (2021, henceforth BH).

2 Motivating Examples

We develop three stylized examples, inspired respectively by the settings of Miguel and Kremer (2004), Donaldson and Hornbeck (2016), and Currie and Gruber (1996), to illustrate the main insights of this paper. In each example we consider estimating the parameter β of a causal or structural model which relates an outcome y_i to a treatment x_i ,

$$y_i = \beta x_i + \varepsilon_i, \tag{1}$$

for a set of units $i = 1, \dots, N$ with an unobserved error ε_i . The common feature of the examples is that x_i is computed from multiple sources of variation by a known formula.

Network spillovers:

Suppose y_i is student i 's educational achievement and x_i counts the number of i 's neighbors who have been dewormed in an intervention:

$$x_i = \sum_{k=1}^N \text{Neighbor}_{ik} \text{Dewormed}_k.$$

⁴A notable exception of a recentering-type regression adjustment in the traditional propensity scores setting is the E-estimator of Robins et al. (1992).

Here $Dewormed_k \in \{0, 1\}$ is an indicator for student k being selected for the deworming intervention and $Neighbor_{ik} \in \{0, 1\}$ indicates that i and k are neighbors (i.e., connected by an observed network link). The error term ε_i captures i 's educational outcome when none of her neighbors are dewormed. This example is a stylized version of the main specification in Miguel and Kremer (2004).⁵

Market access:

Suppose y_i is the growth of land values in region i between two dates $t \in \{0, 1\}$ and $x_i = \log MA_{i1} - \log MA_{i0}$ is the growth of regional market access (MA) due to improvements to the interregional railroad network. Market access is computed

$$MA_{it} = \sum_{j=1}^N \frac{Pop_j}{\tau(Network_t, Loc_i, Loc_j)},$$

following standard models of economic geography (e.g. Redding and Venables (2004)). Here Pop_j is the time-invariant population of region j , $Network_t$ is the set of railway lines and other types of transit which comprise the transportation network in operation at time t , Loc_j is the location of region j on the map, and $\tau(\cdot)$ is a function giving the travel time between regions i and j . The error term ε_i captures location i 's land value growth in the absence of market access growth, due to some regional amenity and productivity shocks. Similar market access growth specifications are considered in, for example, Donaldson and Hornbeck (2016).

Medicaid eligibility:

Suppose y_i is individual i 's health outcome and $x_i \in \{0, 1\}$ indicates her eligibility for Medicaid. Let $IncDem_i$ be a vector of individual income and demographics, $State_i \in \{1, \dots, 50\}$ index i 's state of residence, and $Policy_k$ be state k 's eligibility policy: i.e. the set of income and demographic groups eligible for Medicaid in that state. Then

$$x_i = \mathbf{1}[IncDem_i \in Policy_{State_i}].$$

The error term ε_i captures individual i 's outcome when she is ineligible for Medicaid. This example comes from Currie and Gruber (1996).

To estimate β in each example, we consider a true or natural experiment that manipulates some of the determinants of x_i . Formally, we partition the variables from which $x = (x_1, \dots, x_N)$ is computed into two groups: a set of shocks g and a set of predetermined variables w . The shocks $g = (g_1, \dots, g_k)$ are assumed to be exogenous, i.e. independent of the errors $\varepsilon = (\varepsilon_1, \dots, \varepsilon_N)$. Shock exogeneity combines two conceptually distinct assumptions—that g is as-good-as-randomly assigned, and that this assignment only affects the outcome

⁵Nothing is changed in what follows if one instead considers the number of not-dewormed neighbors as the treatment. For simplicity here we consider only the spillover treatment and not also the direct deworming treatment; see Section 3.6 on the extension to multiple treatments.

of each unit i via its treatment x_i (an exclusion restriction). The shocks can be assigned at a different “level” than the observations, with $K \neq N$. The remaining variables w have an arbitrary structure and $f_i(\cdot; w)$ governs the mapping from the exogenous shocks to each unit’s treatment, i.e. the observation’s “exposure” to the shocks.⁶ We assume that w is determined prior to the (natural) experiment and is unaffected by the shocks.

Network spillovers (cont.)

Suppose deworming is assigned in a randomized control trial (RCT) and βx_i fully captures its spillover effects. Then $g = (Dewormed_k)_{k=1}^K$ collects the exogenous shocks, for $K = N$.⁷ The remaining determinants of the spillover treatments of all units, $w = (Neighbor_{ik})_{i,k=1}^N$, are fixed in the experiment.

Market access (cont.)

Suppose the timing of new railroads is exogenous. Specifically, suppose that among K lines planned to be constructed by $t = 1$ some are randomly delayed by unexpected engineering problems (unrelated to the trends in regional land values). Suppose also the model of economic geography is correctly specified, so $\beta \log M A_i$ fully captures the effects of transportation upgrades. Then $g = (Open_k)_{k=1}^K$ collects the exogenous shocks, where $Open_k$ is an indicator for whether planned line k faces no delays. Assuming no other changes to the network at $t = 1$, we can partition the determinants of MA growth into g and $w = ((Loc, Pop)_{i=1}^N, Network_0)$ as $Network_1$ is fully determined by $Network_0$ and the set of newly opened lines.

Medicaid eligibility (cont.)

Suppose Medicaid policies across the $K = 50$ states are exogenous, i.e. are chosen irrespective of the potential health outcomes and affect individual outcomes only via Medicaid eligibility. Then $g = (Policy_k)_{k=1}^K$ collects the exogenous shocks, with the other determinants of eligibility collected in $w = (IncDem, State)_{i=1}^N$.

The first point of this paper is that ordinary least squares (OLS) estimation of β can suffer from OVB, despite the experimental variation underlying x_i .⁸ The OVB problem arises because some units receive systematically higher values of x_i than others, as a consequence of their non-random exposure to the shocks. This systematic variation may be cross-sectionally correlated with the errors ε_i , generating bias in OLS estimation of equation (1).

⁶Aronow and Samii (2017) use a similar “exposure mapping” terminology for objects like $f_i(\cdot; w)$ in the network spillover context. We depart from this literature by referring to the realized $f_i(g; w)$ as the “treatment” or “candidate instrument” and not the realized “exposure.”

⁷For simplicity here we assume away any direct effects of deworming; see Section 3.6 on the extension to multiple treatments.

⁸Such OVB may arise even if (as in the network spillovers and market access examples) variation in the treatment “results” from the experimental shocks, in the sense that $x_1 = \dots = x_N = 0$ whenever $g_1 = \dots = g_K = 0$.

Network spillovers (cont.)

Even when deworming is randomly assigned to students, those with more neighbors (e.g., because they live in dense urban areas) will tend to have more dewormed neighbors and therefore be more exposed to the deworming intervention. Urban areas may have different educational outcomes for reasons unrelated to deworming, generating OVB.

Market access (cont.)

Even when the opening status of lines is as-good-as-randomly assigned, regions in the economic and geographic center of the country will tend to see more market access growth than peripheral regions as the former are closer to a typical potential line. Central regions may face different amenity and productivity shocks, generating OVB.

Medicaid eligibility (cont.)

Even when Medicaid policies are as-good-as-randomly assigned to states, poorer individuals will tend to see higher rates of eligibility. Poor individuals may face different health shocks, generating OVB.

Our second insight is that this OVB problem has a conceptually simple solution, which follows from viewing the set of realized g as one draw from a shock *assignment process* and considering what counterfactual sets of exogenous shocks could have as likely been drawn. The specification of such counterfactuals allows one to measure and remove the systematic component of variation in the treatment which drives OVB. Specifically, the researcher recomputes the treatment x_i of each unit i across many counterfactual sets of shocks and takes their average to measure the *expected treatment*, μ_i . We show that this μ_i , which is co-determined by the exposure of x_i to the shocks g and the shock assignment process, is the sole confounder in equation (1). OVB can then be purged by “recentering” the treatment, i.e. instrumenting x_i with $\tilde{x}_i = x_i - \mu_i$ in equation (1), or by simply adding μ_i as a control in OLS estimation. The key to removing bias with this approach is thus to credibly specify and average over shock counterfactuals—a task which is trivial in true experiments and which otherwise formalizes the natural experiment of interest.

Network spillovers (cont.)

With deworming assigned in an RCT, the shock assignment process is given by the known randomization protocol. If, say, each student has a 30% chance of being dewormed, the expected number of i 's dewormed neighbors over repeated draws of deworming shocks μ_i is 0.3 times their number of neighbors $\sum_{k=1}^K Neighbor_{ik}$. OVB is thus purged by controlling for the number of neighbors, or by using the recentered number of dewormed neighbors $\tilde{x}_i = x_i - \mu_i$ to instrument for x_i . With either adjustment, the regression will only compare students who had more neighbors dewormed than expected (given the network) to those with fewer-than-expected dewormed neighbors.

Market access (cont.)

The as-good-as-random assignment of opening status can be formalized by each planned line facing an equal and independent chance of opening. Then, if $\sum_{k=1}^K Open_k = K$, railway

lines open by $t = 1$, every counterfactual network in which K_1 lines from the plan opened was as likely to have occurred. One can thus compute expected MA growth μ_i as the average MA growth of region i across these counterfactuals (or a random subset of them). Recentering by or controlling for this μ_i ensures that the regressions only compare regions which saw higher MA growth than expected—given pre-existing economic geography and the plan—to those which saw less-than-expected MA growth.

Medicaid eligibility (cont.)

The as-good-as-random assignment of Medicaid policies can be formalized by each state randomly drawing from a pool of potential policies, such that every permutation of the realized policies was equally likely to have occurred. Averaging individual i 's eligibility across these permutations yields an expected eligibility μ_i which equals the share of states in which she would be eligible. Our solution is to instrument actual eligibility with recentered eligibility $x_i - \mu_i$, or control for μ_i in an OLS regression. Either approach would, for example, effectively remove from the sample “always-eligible” or “never-eligible” individuals (with $x_i = \mu_i = 1$ or $x_i = \mu_i = 0$) whose income and demographics make them unaffected by policy variation.

The recentering solution generally dominates more conventional ones, such as instrumenting directly by the shocks or controlling for the other determinants of x_i . Instrumenting with the shocks is infeasible when the shocks are assigned at a different level than the units, and generally discards variation in treatment due to w . Controlling for an observation's non-random shock exposure flexibly is typically infeasible, because such exposure is high-dimensional. Conversely, low-dimensional controls are only guaranteed to purge OVB (absent additional non-experimental restrictions on the error term) when they linearly span μ_i , which is difficult to establish except when μ_i is known and recentering is feasible. If either the assignment process or shock exposure mapping is complex, μ_i is unlikely to be a simple function of observed characteristics.

Network spillovers (cont.)

Using student i 's own deworming status as an instrument is infeasible as it does not predict the number of dewormed neighbors; incorporating the non-random network adjacency matrix is necessary. Controlling for the entire row of the adjacency matrix (which characterizes student's exposure) is also infeasible, as it would absorb all cross-sectional variation in the treatment. Controlling for the number of i 's neighbors is enough to purge OVB under completely random assignment of deworming, since this control is proportional to μ_i . However, such simple controls would not linearly span μ_i with more complex randomization protocols, such as with two tiers (by school, then by student) or stratification (e.g., with girls dewormed with a known higher probability). Simple controls are also generally insufficient with more complex specifications of spillovers.⁹

⁹An example is given by Carvalho et al. (2021), where i is a Japanese firm and x_i is the distance in the firm-to-firm supply network from i to the nearest firm located in the area hit by an earthquake. Unlike the number of treated neighbors, this spillover treatment is a nonlinear function of the earthquake shock dummies. The earthquake assignment process is also more complex, exhibiting spatial correlation. Our recentering approach still applies naturally in cases like this.

Market access (cont.)

Railroad timing shocks vary at the level of lines, so it is infeasible to use them as instruments for regional market access without incorporating some non-random features of economic geography. Controlling perfectly for these features is also infeasible, as each region's market access depends on the entire spatial distribution of economic activity. Simple sets of controls, such as polynomials in the latitude and longitude of a region, need not linearly span μ_i given the complexity of x_i , and thus are not guaranteed to purge OVB.

Medicaid eligibility (cont.)

Currie and Gruber (1996) propose instrumenting individual eligibility with a measure of the overall policy generosity of her state—a so-called “simulated instrument.” Such instruments are simple functions of $Policy_k$ for all individuals in state k and are thus exogenous and relevant under random policy assignment. However they discard relevant within-state variation in i 's income and demographics and are thus likely to yield a less powerful first-stage prediction of x_i than recentered eligibility.¹⁰

We conclude this section by noting that the OVB problem and recentering solution both extend to the case with an arbitrary endogenous x_i and a candidate instrument z_i which is constructed from exogenous shocks and other variables by a known formula. This approach is natural when the treatment can be represented as a function of exogenous shocks g , predetermined variables w , and endogenous (and possibly unobserved) variables u : i.e. when $x_i = h_i(g, w, u)$ for a known $h_i(\cdot)$. An intuitive candidate instrument for x_i is the prediction of x_i in the scenario when the u shocks are ignored: $z_i = h_i(g, w, 0)$. Our framework shows that these candidate instruments are generally invalid, again because of the non-random exposure of z_i to g . Yet OVB can again be purged by measuring the expected instrument μ_i —the average z_i across counterfactual g —and either instrumenting x_i with the recentered IV $\tilde{z}_i = z_i - \mu_i$ or controlling for μ_i while instrumenting with z_i .

Market access (cont.)

Suppose population sizes also change between $t = 0$ and $t = 1$, and the observable changes $u = (Pop_{t1} - Pop_{t0})_{i=1}^N$ are not exogenous (e.g. they respond to house price shocks in ϵ). Then one can consider instrumenting the realized change in MA by a predicted change in MA which keeps population sizes fixed at $t = 0$ levels. Without recentering, this IV regression may suffer from the same OVB as the OLS regression discussed above. OVB is now avoided by recentering the MA prediction via counterfactual railroad networks.

Medicaid eligibilty (cont.)

Suppose one is interested in the effects of Medicaid takeup, instead of eligibility. Takeup is the product of eligibility and $1 - NeverTaker_i$, where $NeverTaker_i$ indicates that individual

¹⁰With completely random policy assignment, flexibly controlling for $IncDem_i$ may purge OVB as this is the only source of variation in μ_i . However, even in this setting the relevant demographics in $IncDem_i$ and their interactions can be high-dimensional, as discussed by Gruber (2003). This problem is exacerbated under more complex assignment processes, e.g. if policies can be viewed as random only within some groups of states, in which case group indicators and their interactions with the demographics would also have to be included. Recentering extends naturally and avoids the curse of dimensionality.

i would decline Medicaid if eligible and $u = (NeverTaker_i)_{i=1}^N$ is unobserved. Under the appropriate exclusion restriction one can consider instrumenting takeup with eligibility; our recentering strategy then again removes OVB from non-random variation in policy exposure.

3 Theory

We now develop a general econometric framework for settings with non-random exposure to exogenous shocks. We introduce the baseline setting, develop our approach to estimation based on the recentering procedure, and discuss how this recentering can be performed by specifying counterfactual shocks in Sections 3.1-3.3. We then discuss conditions for consistency of recentered IV estimators in Section 3.4, and how inference can be conducted in Section 3.5. Several extensions are summarized in Section 3.6.

3.1 Setting

We consider estimation of β in the causal or structural model

$$y_i = \beta x_i + \varepsilon_i, \tag{2}$$

from a dataset of scalar and demeaned y_i and x_i , $i = 1, \dots, N$. Below we discuss extensions to heterogeneous causal effects, nonlinear models, multiple treatments, and additional control variables. Although we use a single index of i for observations, we note our framework accommodates repeated cross-sections and panel data.

Importantly for the applicability of our framework, we do not assume that the observations of y_i and x_i are independently or identically distributed (*iid*) as when arising from random sampling. This allows for complex dependencies across the units due to their common exposure to observed and potentially unobserved shocks. It is also consistent with settings where the N units represent a population—for example, all regions of a country—and conventional random sampling assumptions are inappropriate (Abadie et al., 2020).¹¹

We suppose that to estimate β a researcher has constructed a candidate instrument

$$z_i = f_i(g; w), \tag{3}$$

where $f_1(\cdot), \dots, f_N(\cdot)$ is a list of known non-stochastic functions, g is a $K \times 1$ vector of shocks, and w is a list of other variables of unrestricted dimension. Equation (3) is very general: any z_i that can be computed from a set of observed data, according to a known formula, can be described in this way.¹² It also allows $x_i = z_i$, in which case β is the causal effect of the composite treatment.

¹¹Formally, we assume $(x_i, \varepsilon_i)_{i=1}^N$ and the g and w variables introduced below are all drawn from some joint distribution which is unrestricted at this point.

We assume that the shocks g are exogenous, which we formalize by their conditional independence from the vector of errors given the other sources of instrument variation:

Assumption 1. (Shock exogeneity): $g \perp\!\!\!\perp \varepsilon \mid w$

As noted in Section 2, this notion of shock exogeneity combines two conceptually distinct conditions. First, it imposes an exclusion restriction, reflecting an economic model of how g can affect y . Second, it requires as-good-as-random shock assignment. This latter condition is satisfied when the shocks are fully randomly assigned, as in an RCT (i.e., $g \perp\!\!\!\perp (\varepsilon, w)$), but also allows w to contain variables that govern the shock assignment process.¹³ Importantly, Assumption 1 allows $\mathbb{E}[\varepsilon_i \mid w]$ to vary arbitrarily across i ; this reflects the lack of non-experimental assumptions, such as parallel trends, constraining the error in equation (2).¹⁴ Assumption 1 is consistent with a two-step data-generating process, where w is determined prior to the realization of shocks g and errors ε which then together determine (x, y) .¹⁵

We start by considering an instrumental variable (IV) regression of y_i on x_i that instruments with z_i . As usual, this strategy requires z_i to be relevant to the treatment and orthogonal to the error term. In our non-*iid* setting, we formalize these two conditions in terms of the full-sample IV moments $\mathbb{E}\left[\frac{1}{N} \sum_i z_i x_i\right]$ and $\mathbb{E}\left[\frac{1}{N} \sum_i z_i y_i\right]$. Since (2) implies $\mathbb{E}\left[\frac{1}{N} \sum_i z_i y_i\right] = \beta \mathbb{E}\left[\frac{1}{N} \sum_i z_i x_i\right] + \mathbb{E}\left[\frac{1}{N} \sum_i z_i \varepsilon_i\right]$, β is recoverable from the ratio of these moments (what we term *identification*) under the relevance condition of $\mathbb{E}\left[\frac{1}{N} \sum_i z_i x_i\right] \neq 0$ and the orthogonality condition of $\mathbb{E}\left[\frac{1}{N} \sum_i z_i \varepsilon_i\right] = 0$.¹⁶ To start we assume the two IV moments are known, in order to focus on the potential for OVB when the orthogonality condition fails. We discuss conditions for consistent estimation in Section 3.4.

¹²In some cases (such as the network spillover and medicaid eligibility examples in the previous section) the candidate instrument can be naturally written as $z_i = f(g, w_i)$ with a common function $f(\cdot)$ and a unit-specific measure of exposure w_i . In other cases a more general notation is necessary: in the market access example, for example, the MA of each region depends on the entire country's economic geography. An alternative way to formalize general composite variables is $z_i = f(g; w, \tilde{w}_i)$ for a common w and unit-specific \tilde{w}_i . This notation is equivalent to (3); we use $f_i(g; w)$ as it is more compact. We also note also that equation (3) does not contain a residual: it formalizes an algorithm for computing an instrument rather than characterizing an economic relationship.

¹³The exclusion and as-good-as-random assignment assumptions are isolated in Appendix C.1 of BH, via a general potential outcomes model.

¹⁴Our identification results hold under the weaker conditional mean independence assumption of $\mathbb{E}[\varepsilon \mid g, w] = \mathbb{E}[\varepsilon \mid w]$. This assumption can be understood as defining a partially linear model, as in Robinson (1988): $y_i = \beta x_i + \psi_i(w) + \tilde{\varepsilon}_i$ where $\psi_i(w) = \mathbb{E}[\varepsilon_i \mid w]$ and $\mathbb{E}[\tilde{\varepsilon}_i \mid g, w] = 0$ for $\tilde{\varepsilon}_i = \varepsilon_i - \psi_i(w)$. A difference from Robinson (1988) arises because we do not assume *iid* data; for instance, we do not assume $\psi_i(w) \equiv \psi(w)$ for *iid* w_i .

¹⁵Throughout, we allow (ε, w) to be stochastic (as when some components are sampled from a superpopulation) or fixed (as in a more conventional "design-based" analysis; e.g. Athey and Imbens (2022)). In the Medicaid eligibility example it may be more natural to view the observed $(IncDem_i, state_i)$ as sampled from the national population, along with untreated potential outcomes ε_i . Conversely, in the market access example, it may be more natural to view the set of observed regions as a finite population, with fixed geography $(Loc_i)_{i=1}^N$. With fixed (ε, w) , Assumption 1 holds trivially but Assumption 2 below is still restrictive.

¹⁶It is worth emphasizing that in our non-*iid* setup these conditions combine two dimensions of variation: over the stochastic realizations of g, w, x , and ε , and across the cross-section of observations $i = 1, \dots, N$. In the *iid* case they reduce to the more familiar conditions of $\mathbb{E}[z_i x_i] \neq 0$ and $\mathbb{E}[z_i \varepsilon_i] = 0$.

3.2 OVB and Instrument Recentering

We define the *expected instrument* $\mu_i = \mathbb{E}[f_i(g; w) \mid w]$ as the average value of z_i across different realizations of the shocks, conditional on w . Our first result shows that OVB may arise when predetermined exposure to the natural experiment is endogenous, and that the potential for such bias is entirely governed by the relationship between μ_i and the error ε_i . Formally, under Assumption 1 instrument orthogonality need not hold: $\mathbb{E}\left[\frac{1}{N} \sum_i z_i \varepsilon_i\right] \neq 0$ in general. Rather,

$$\mathbb{E}\left[\frac{1}{N} \sum_i z_i \varepsilon_i\right] = \mathbb{E}\left[\frac{1}{N} \sum_i \mu_i \varepsilon_i\right]. \quad (4)$$

This result follows from the law of iterated expectations:

$\mathbb{E}[z_i \varepsilon_i] = \mathbb{E}[\mathbb{E}[f_i(g; w) \varepsilon_i \mid w]] = \mathbb{E}[\mu_i \mathbb{E}[\varepsilon_i \mid w]] = \mathbb{E}[\mu_i \varepsilon_i]$ for all i , where the second equality uses Assumption 1 and the definition of μ_i .

The central role of μ_i in governing OVB suggests the recentering solution: even though OVB results from potentially high-dimensional variation in units' exposure to shocks, adjustment for the one-dimensional confounder μ_i is sufficient for instrument orthogonality. We adjust z_i by defining the *recentered instrument* $\tilde{z}_i = z_i - \mu_i$. By equation (4), orthogonality always holds for this instrument:

$$\mathbb{E}\left[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i\right] = \mathbb{E}\left[\frac{1}{N} \sum_i z_i \varepsilon_i\right] - \mathbb{E}\left[\frac{1}{N} \sum_i \mu_i \varepsilon_i\right] = 0.$$

Thus, if \tilde{z}_i is also relevant, β is identified by the IV regression which uses the recentered instrument \tilde{z}_i instead of z_i .¹⁷

A closely related solution, also suggested by equation (4), is to include the expected instrument μ_i as a control in specification (2) while using the original z_i as the instrument in a control function approach (Wooldridge, 2015). Controlling for μ_i can be thought of as recentering z_i while also removing the residual variation in y_i which is cross-sectionally correlated with μ_i . As usual, removing this residual variation may generate precision gains in large samples; similar gains may arise from including (a fixed number of) any predetermined controls in a recentered IV regression.¹⁸

¹⁷There exist $f_i(\cdot)$ constructions that yield a relevant recentered instrument whenever the shocks induce some variation in treatment. Formally, when $\text{Var}[\mathbb{E}[x_i \mid g, w] \mid w]$ is not almost-surely zero at least for some i , the recentered instrument constructed as $\tilde{z}_i = \mathbb{E}[x_i \mid g, w] - \mathbb{E}[x_i \mid w]$ is relevant. This again follows by the law of iterated expectations:

$$\mathbb{E}\left[\frac{1}{N} \sum_i \tilde{z}_i x_i\right] = \mathbb{E}\left[\frac{1}{N} \sum_i \tilde{z}_i \mathbb{E}[x_i \mid g, w]\right] = \mathbb{E}\left[\frac{1}{N} \sum_i \text{Var}[\mathbb{E}[x_i \mid g, w] \mid w]\right].$$

¹⁸Formally, the regression with μ_i as a control yields the reduced-form and first-stage moments $\mathbb{E}\left[\frac{1}{N} \sum_i z_i y_i^\perp\right]$ and $\mathbb{E}\left[\frac{1}{N} \sum_i z_i x_i^\perp\right]$, where v_i^\perp denotes the residuals from a cross-sectional projection of v_i on μ_i . We show in Appendix B.1 of BH that these moments also

Equation (4) further shows that adjusting for μ_i is generally *necessary* for identification, absent additional restrictions on the unobserved error. Conventional controls and fixed effects are only guaranteed to purge OVB when they linearly span μ_i : a condition that is difficult to verify except when recentering is also feasible.¹⁹

Adjustments based on μ_i , as the sole confounder of z_i , are similar to more conventional propensity score methods. There are three key differences, concerning the setting, adjustment method, and computation of μ_i . First, propensity score methods have mostly been applied to binary treatments, starting from Rosenbaum and Rubin (1983). While generalizations to binary instruments (e.g. Abadie (2003)) and non-binary treatments (e.g. Imbens (2000)) have been proposed, our setting allows for arbitrary treatments or instruments. Second, the propensity score literature has mostly used non-regression adjustment methods, such as matching or binning (Abadie and Imbens, 2016; King and Nielsen, 2019). A notable exception is the E-estimator of Robins et al. (1992), which similarly leverages linearity of an outcome model like (2) to recenter by a scalar variable. Third, and most importantly, propensity scores are usually estimated from the data by relating the treatment to a vector of observation-specific covariates. This approach is generally not feasible because exposure to exogenous shocks is high-dimensional: for instance, as noted in Section 2, the expected market access of any region i depends on the entire economic geography of the country. We therefore take a different approach to computing μ_i , which we turn to next.

3.3 Computing the Expected Instrument via Shock Counterfactuals

We propose computing the expected instrument by specifying an assignment process for the shocks, drawing many sets of counterfactual shocks from this process, recomputing the candidate instrument each time, and averaging it across the counterfactuals. Here we formalize this approach, discuss general ways in which counterfactual shocks can be specified, and highlight the advantages of our approach over alternatives.

We define the shock assignment process as the conditional distribution of $g \mid w$, with cumulative distribution function $G(g \mid w)$. When $G(\cdot)$ is known, the expected instrument $\mu_i = \int f_i(\gamma; w) dG(\gamma \mid w)$ can be computed and either used to recenter z_i or added as a regression control.²⁰ To emphasize the importance of a known shock assignment process, we write it as an assumption:

identify β under Assumption 1. Appendix C.9 of BH shows that controlling for μ_i always reduces asymptotic variance of the estimator when $z_i \mid w$ is homoskedastic, while also giving a counterexample under heteroskedasticity.

¹⁹In panel data with $z_{it} = f_{it}(g_t, w_t)$, for example, unit fixed effects generally purge OVB only when the expected instrument is time-invariant, which generally requires the $f_{it}(\cdot)$ mapping, the value of w_t , and the distribution of g_t to be time-invariant. While plausible in some applications, these conditions (in particular, stationarity of the shock distribution) can be quite restrictive. For instance, when new railroad lines tend to be built more than destroyed, expected market access will tend to grow over time.

²⁰For the identification results, it is enough to approximate μ_i by an average of $f_i(g^{(s)}; w)$ for any number S of $g^{(s)}$ drawn from $G(g \mid w)$, independently of each other and of g . We have, for example, $\mathbb{E}\left[\frac{1}{N} \sum_i \left(z_i - \frac{1}{S} \sum_s f_i(g^{(s)}; w)\right) \varepsilon_i\right] = 0$ by iterated expectations, since $\mathbb{E}[z_i \mid w, \varepsilon] = \mathbb{E}[f_i(g^{(s)}; w) \mid w, \varepsilon]$. We discuss how the number of draws affects the asymptotic behavior of the recentered IV estimator below.

Assumption 2. (Known assignment process): $G(g | w)$ is known in the support of w .

This assumption is unrestrictive when the shocks are determined by a known randomization protocol, as in an RCT or with policy randomizations (such as tie-breaking lottery numbers in centralized assignment mechanisms; Abdulkadiroglu et al. (2017)). The assignment process may also be given by scientific knowledge when the shocks are randomized naturally, such as when g captures weather or seismic shocks governed by meteorological or geological processes (e.g., Carvalho et al. (2021); Madestam et al. (2013)). Policy discontinuities (as in regression discontinuity designs) can also yield a known $G(\cdot)$ when viewed as generating local randomization around known cutoffs (Lee, 2008; Cattaneo et al., 2015).

In observational data, where the distribution of shocks is unknown, Assumption 2 can be satisfied by specifying some permutations of shocks that were as likely to have occurred. For instance, if one is willing to assume the shocks g_k are *iid* across k , it follows that all permutations of the observed g are equally likely. In this case $G(g | w)$ is uniform when w is augmented by the permutation class $\Pi(g) = \{\pi(g) | \pi(\cdot) \in \Pi_K\}$, where Π_K denotes the set of permutation operators $\pi(\cdot)$ on vectors of length K (e.g. Lehmann and Romano, 2006, p. 634). The distribution of each g_k (conditionally on other components of w) then needs not be specified; the expected instrument is the average z_i across all permutations of shocks, which serve as counterfactuals:

$$\mu_i = \frac{1}{K!} \sum_{\pi(\cdot) \in \Pi_K} f_i(\pi(g); w).$$

Such μ_i are easy to compute (or approximate with a random set of permutations).

Similar expected instrument calculations follow under weaker shock exchangeability conditions, such as when the g_k are *iid* within, but not across, a set of known clusters and the class of within-cluster permutations is used to draw counterfactuals. We illustrate this approach in Section 4. In BH we discuss how our framework can also apply with $G(g | w)$ specified up to a low-dimensional vector of consistently estimable parameters (Appendix C.5); we also show how Assumption 2 can derive from an economic model (e.g. of transportation network formation) with stochastic shocks or from symmetries of the joint shock distribution (Appendices D.1 and D.2)

We note that even when $G(\cdot)$ is challenging to specify, a possibly incorrect specification can be useful as a sensitivity check. Specifically, if Assumption 1 holds and there is already no OVB because the included regression controls perfectly capture either the endogenous features of exposure or the expected instrument, then controlling for any candidate expected instrument $m_i(w)$ cannot introduce bias. In this case the researcher may safely control for one or several $m_i(w)$ based on some guesses of the assignment process.²¹ More generally, researchers may achieve additional robustness by controlling for multiple candidate $m_i(w)$ based on multiple shock assignment process guesses; only one such guess needs to be right to purge OVB.

3.4 Recentered IV Consistency

With the ratio of recentered IV moments identifying β , we now consider whether the corresponding IV estimator $\hat{\beta} = (\frac{1}{N} \sum_i \tilde{z}_i y_i) / (\frac{1}{N} \sum_i \tilde{z}_i x_i)$ is consistent, i.e. whether $\hat{\beta} \xrightarrow{P} \beta$ as the number of observed outcomes and treatments grows large ($N \rightarrow \infty$). To formalize consistency in our *non-iid* context we consider a sequence of distributions P_N for the complete data (y, x, g, w) . Only in this section, to make the asymptotic sequence explicit, we index moments by P_N : e.g., we write the recentered IV moments as $\mathbb{E}_{P_N}[\frac{1}{N} \sum_i \tilde{z}_i y_i]$ and $\mathbb{E}_{P_N}[\frac{1}{N} \sum_i \tilde{z}_i x_i]$. We allow the number of observed shocks, $K_N = \dim(g)$, and the dimensions of w to change arbitrarily with N .

We first consider mean-square convergence of $\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i$ to $\mathbb{E}_{P_N}[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i] = 0$ under Assumption 1: i.e., whether $\text{Var}_{P_N}[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i] \rightarrow 0$. Since $\hat{\beta} = \beta + (\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i) / (\frac{1}{N} \sum_i \tilde{z}_i x_i)$ by (2), such convergence implies $\hat{\beta} \xrightarrow{P} \beta$ so long as the instrument is asymptotically relevant (a condition we return to below). We establish this convergence under a regularity condition on ε_i and a substantive restriction on \tilde{z}_i , which we term weak mutual dependence:

Assumption 3. (Weak mutual dependence):

$$\mathbb{E}_{P_N} \left[\frac{1}{N^2} \sum_{i,j} |\text{Cov}_{P_N}[\tilde{z}_i, \tilde{z}_j \mid w]| \right] \rightarrow 0.$$

Proposition 1. Suppose Assumptions 1-3 hold and $\mathbb{E}_{P_N}[\varepsilon_i^2 \mid w] \leq U_\varepsilon$ uniformly across N and $i = 1, \dots, N$. Then $\text{Var}_{P_N}[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i] \rightarrow 0$.

Proof. See Appendix D.

Assumption 3 holds when the shocks induce rich cross-sectional variation in the recentered instrument, through heterogeneous exposure, such that most pairs of $(\tilde{z}_i, \tilde{z}_j)$ have a weak covariance across possible realizations of g . The proof shows this is enough for a law of large numbers to apply to $\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i$.²²

Note that in line with our approach to identification, Proposition 1 makes no substantive restrictions on the errors ε_i beyond Assumption 1 (in particular, it puts no restrictions on the dependence of ε_i across observations). In the absence of such restrictions, Proposition 2

²¹Formally, suppose either $\mathbb{E}[z_i \mid w] = 0$ or $\mathbb{E}[\check{\varepsilon}_i \mid w] = 0$ for each i , where $\check{\varepsilon}_i$ denotes the cross-sectional residualization of variable v_i on some functions of w used as controls. Then $\mathbb{E}[\frac{1}{N} \sum_i \check{\varepsilon}_i^\perp] = 0$, where here v_i^\perp denotes the residuals from a cross-sectional projection of v_i on $m(w)$. See Appendix C.6 of BH for our framework extended to predetermined controls.

²²The proof to Proposition 1 shows Assumption 3 is also sufficient for consistency when μ_i is approximated by an average of $f_i(g^{(s)}; w)$ for $g^{(s)}$ drawn as in footnote 20. Intuitively, the variance of $\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i$ is higher with fewer simulations S but converges to zero for any fixed S under Assumption 3.

in Appendix C shows the Assumption 3 is not only sufficient for $\text{Var}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i \right] \rightarrow 0$ but, under regularity conditions, also necessary. Of course, more conventional restrictions on the mutual dependence of errors (such as *iid* or clustered ε_i) may also suffice for convergence when weak mutual dependence of \tilde{z}_i fails.

Three additional results in Appendix C, which extend the results on consistency with linear shift-share instruments from Borusyak et al. (2022), unpack Assumption 3 further. First, a large number of exogenous shocks is essentially necessary for the recentered instrument to not have many strong cross-sectional dependencies. Proposition 3 formalizes this intuition by showing that, with sufficiently smooth $f_i(\cdot; w)$, Assumption 3 can only hold with $K_N \rightarrow \infty$. Moreover, the concentration of exposure to this growing number of shocks matters. Proposition 4 formalizes this idea by considering a concentration measure for average shock exposure which is similar to a Herfindahl-Hirschman Index (HHI): $\mathbb{E}_{P_N} \left[\sum_{k=1}^{K_N} \left(\frac{\partial \bar{f}(g; w)}{\partial g_k} \right)^2 \right]$, where $\bar{f}(g; w) = \frac{1}{N} \sum_i (f_i(g; w) - \mu_i)$. For binary g_k and weakly monotone $f_i(\cdot; w)$ (as in the network spillovers and market access examples) and with mutually-independent shocks, Assumption 3 is satisfied when this measure converges to zero such that the impact of any finite set of shocks on $\frac{1}{N} \sum_i \tilde{z}_i$ vanishes.²³ Proposition 5 considers a different low-level condition in a case covering the medicaid eligibility example: Assumption 3 holds when most pairs of observations of \tilde{z}_i are affected by non-overlapping sets of shocks.

Convergence of $\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i$ implies consistency of the recentered IV estimator so long as

(i) $\mathbb{E}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i x_i \right]$ remains bounded away from zero and (ii) $\frac{1}{N} \sum_i \tilde{z}_i x_i - \mathbb{E}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i x_i \right] \xrightarrow{P} 0$.²⁴

Condition (i) follows when the relationship between z_i and x_i is strong and when most observations of \tilde{z}_i have exposure concentrated in a small number of exogenous shocks, such that $\text{Var}_{P_N}[\tilde{z}_i]$ does not dissipate even as $K_N \rightarrow \infty$. Proposition 6 in Appendix C formalizes these conditions with a linear first stage model of $x_i = \pi z_i + u_i$, with $g \perp u \mid w$, and a different measure of shock exposure concentration: the HHI of the effects of different shocks on \tilde{z}_i , $\sum_{k=1}^{K_N} \left(\mathbb{E}_{P_N} \left[\frac{\partial f_i(g; w)}{\partial g_k} \mid w \right] \right)^2$, averaged across observations i . In the case of mutually-independent binary shocks, we require $\pi \neq 0$ and that the expectation of this concentration measure is bounded above zero. The HHI conditions from Propositions 4 and 6 may simultaneously hold when most observations are mostly exposed to a small number of

²³Proposition 4 also applies when shocks are normally distributed, and when $f_i(\cdot; w)$ is linear regardless of the shock distribution. It further gives a necessary condition for Assumption 3 in all of these cases, which is a slightly stronger notion of vanishing average shock exposure concentration. For binary shocks, $\partial \bar{f}(g; w) / \partial g_k$ is defined as the difference in \bar{f} when g_k switches from 0 to 1, keeping all other shocks fixed.

²⁴This follows because $\hat{\beta} - \beta = h_N(r_1, r_2) \equiv r_1 / \left(\mathbb{E}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i x_i \right] + r_2 \right)$ for $r_1 = \frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i$ and

$r_2 = \frac{1}{N} \sum_i \tilde{z}_i x_i - \mathbb{E}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i x_i \right]$. Since $h_N(\cdot)$ is Lipschitz-continuous at $(0, 0)$ with a Lipschitz constant uniformly bounded from above, $\hat{\beta} \xrightarrow{P} \beta$ if r_1 and r_2 converge to zero in probability.

shocks, differentially across a large number of shocks. Conditions similar to Assumption 3 can be derived to ensure convergence of the sample first stage, (ii).²⁵

3.5 Randomization Inference and Specification Tests

In some applications of our framework, natural assumptions on the mutual independence of \tilde{z}_i or ε_i across observations can make conventional (e.g. clustered) asymptotic inference valid. Generally, however, the common exposure of observations to observed and unobserved shocks generates complex dependencies across observations making conventional asymptotic analysis inapplicable.²⁶ In such cases, it may be attractive to construct confidence intervals for the constant effect β and tests for Assumptions 1 and 2 based on the specification of the shock assignment process, following a long tradition of randomization inference (RI; Fisher, 1935). The RI approach guarantees correct coverage in finite samples of both observations and shocks.²⁷ We focus on a particular type of RI test which is tightly linked to the recentered IV estimator $\hat{\beta}$.

RI tests and confidence intervals for β are based on a scalar test statistic $T = \mathcal{T}(g, y - bx, w)$, where b is a candidate parameter value. Under the null hypothesis of $\beta = b$ and Assumption 1, the distribution of $T = \mathcal{T}(g, \varepsilon, w)$ conditional on ε and w is implied by the shock assignment process $G(g \mid w)$. One may simulate this distribution, by redrawing the g shocks and recomputing T . If the original value of T is far in the tails of the simulated distribution, one has grounds to reject the null. Inversion of such tests yields confidence interval for β by collecting all b that are not rejected. These intervals have correct size, both conditionally on (ε, w) and unconditionally (see Appendix C.3 of BH for details).

We propose addressing the practical issue of choosing a randomization test statistic by picking a T that is tightly linked to the recentered IV estimator, building on the theory of Hodges and Lehmann (1963) and Rosenbaum (2002). Specifically, we consider the sample covariance of the recentered instrument and implied residual:

$T = \frac{1}{N} \sum_i (f_i(g; w) - \mu) \cdot (y_i - bx_i)$. Lemma 2 of BH shows that $\hat{\beta}$ is a Hodges-Lehmann estimator corresponding to this T , meaning that $\hat{\beta}$ equates T with its expectation across counterfactual shocks (specifically, zero).²⁸ This connection makes RI tests and confidence intervals based on T inherit the consistency of $\hat{\beta}$: the test power asymptotically increases to

²⁵For example, with a linear first stage, $\text{Var}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i x_i \right] = \text{Var}_{P_N} \left[\pi \frac{1}{N} \sum_i \tilde{z}_i^2 + \frac{1}{N} \sum_i \tilde{z}_i (\pi \mu_i + u_i) \right]$. Here, with mutually-independent binary shocks, Lemma 4 in Appendix C ensures $\text{Var}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i^2 \right] \rightarrow 0$ when the expected sum of squared effects of individual shocks on $\frac{1}{N} \sum_i \tilde{z}_i^2$ converges to zero, and Proposition 1 implies $\text{Var}_{P_N} \left[\frac{1}{N} \sum_i \tilde{z}_i (\pi \mu_i + u_i) \right] \rightarrow 0$ when Assumption 3 holds and $\mathbb{E}_{P_N} \left[(\pi \mu_i + u_i)^2 \mid w \right]$ is uniformly bounded.

²⁶An exception is Adão et al. (2019) who derive non-standard asymptotic inference in one such setting: when z_i is a linear shift-share variable, i.e. with $f_i(g; w) = \sum_k w_{ik} g_k$.

²⁷Specifically, RI guarantees the validity of tests for the model parameter β , which can be interpreted as a constant treatment effect. Valid inference with heterogeneous effects in the kind of interdependent data we study is a difficult challenge, even with an asymptotic approach (Adão et al., 2019).

²⁸With additional predetermined controls included in the regression (e.g. μ_i), the same property is satisfied by the residualized statistic $\frac{1}{N} \sum_i \tilde{z}_i (y_i^\perp - bx_i^\perp)$, where here v_i^\perp denotes the residuals from a cross-sectional projection of v_i on the included controls.

one for any fixed alternative $b \neq \beta$ under additional regularity conditions (Proposition S2 of BH).²⁹

Randomization inference can also be used to perform falsification tests on our key Assumptions 1 and 2. Recentering implies a testable prediction that \tilde{z}_i is orthogonal to any variable $r = (r_i)_{i=1}^N$ satisfying $g \perp r \mid w$, such as any function of w or other observables (either predetermined or contemporaneous) thought to be conditionally independent of g . To test this restriction, one may check that the sample covariance $T = \frac{1}{N} \sum_i \tilde{z}_i r_i$ is sufficiently close to zero by drawing counterfactual shocks and checking that T is not in the tails of its conditional-on- (w, r) distribution. Multiple falsification tests, based on a vector of predetermined variables R_i , can be combined by an appropriate RI procedure, e.g. by taking T to be the sample sum of squared fitted values from regressing \tilde{z}_i on R_i .³⁰

Falsification tests can be useful in two ways. First, when r_i is a lagged outcome or another variable thought to proxy for ϵ_i , they provide an RI implementation of conventional placebo and covariate balance tests of Assumption 1. While the use of RI for inference on causal effects may be complicated by treatment effect heterogeneity, the sharp hypothesis of zero placebo effects is a natural null. Second, RI tests will generally have power to reject false specifications of the shock assignment process, i.e. violations of Assumption 2, even when r_i does not proxy for ϵ_i . For $r_i = 1$, for example (which is trivially conditionally independent of g), the test verifies that the sample mean of z_i is typical for the realizations of the specified assignment process. Setting $r_i = \mu_i$ instead checks that the recentered instrument is not correlated with the expected instrument that it is supposed to remove.

3.6 Extensions

While we analyze the constant-effect model (2), identification by μ_i -adjusted regressions extends to settings with heterogeneous treatment effects. Namely, Appendix C.1 of BH shows that the recentered IV estimator generally identifies a convex-weighted average of heterogeneous effects under an appropriate monotonicity condition, extending Imbens and Angrist (1994). The weights are proportional to the conditional variance of $\tilde{z}_i \mid w$ across counterfactual shocks, σ_i^2 . These σ_i^2 , like μ_i , are given by the shock assignment process (Assumption 2) and therefore can be computed by the researcher. Moreover, they can be used to identify more conventional weighted average effects. For example, in reduced-form models of the form $y_i = \beta_i z_i + \epsilon_i$ a recentered and rescaled IV $(z_i - \mu_i) / \sigma_i^2$ identifies the average effect $\mathbb{E}\left[\frac{1}{N} \sum_{i=1}^N \beta_i\right]$. Similarly, in IV settings with binary x_i and z_i , this rescaled instrument identifies the local average treatment effect of Imbens and Angrist (1994).³¹

²⁹RI confidence intervals based on this statistic are still obtained by test inversion, and not from the distribution of the recentered estimator itself across counterfactual shocks g^* . The latter idea fails in IV since the re-randomized instrument $f_i(g^*; w) - \mu_i$ has a true first-stage of zero. The distribution of reduced-form coefficients across counterfactual shocks is also not useful, except for testing $\beta = 0$, as that distribution is centered around zero rather than β .

³⁰This $T = \tilde{z}' R(R', R)^{-1} R' \tilde{z}$ extends our single-dimensional test: it is a quadratic form of the vector-valued statistic $\frac{1}{N} \sum_i \tilde{z}_i R_i$, weighted by $(R', R)^{-1}$, where R is the matrix collecting R_i and \tilde{z} is the vector collecting \tilde{z}_i .

Further extensions are given in the appendix of BH. Appendix C.6 shows how predetermined observables can be included as regression controls to reduce residual variation and potentially increase power. Appendix C.7 discusses identification and inference with multiple treatments or instruments. Finally, Appendix C.8 extends the framework to nonlinear outcome models.

4 Application: Effects of Transportation Infrastructure

We now present an empirical application showing how our theoretic framework can be used to avoid OVB in practice. Specifically, we estimate the effect of market access growth on Chinese regional employment growth over 2007–2016, leveraging the recent construction of high-speed rail (HSR). We show how counterfactual HSR shocks can be specified, and how correcting for expected market access growth can help purge OVB.

The recent construction of Chinese HSR has produced a network longer than in all other countries combined (Lawrence et al., 2019). The network mostly consists of dedicated passenger lines and has developed rapidly since 2007.³² Construction objectives included freeing up capacity on the low-speed rail network and supporting economic development by improving regional connectivity (Lawrence et al., 2019; Ma, 2011). While affordable fares make HSR popular for multiple purposes, business travel is an important component of rail traffic, ranging between 28% and 62%, depending on the line (Ollivier et al., 2014; Lawrence et al., 2019). The role of HSR may also extend beyond directly connected regions, as passengers frequently transfer between HSR and traditional lines (and between intersecting HSR lines). An early analysis by Zheng and Kahn (2013) finds positive effects of HSR on housing prices, while Lin (2017) similarly finds positive effects on regional employment.

We analyze HSR-induced market access effects for 340 sub-province-level administrative divisions in mainland China, referred to as prefectures.³³ We measure market access growth between 2007 and 2016 by combining data on the development of the HSR network and each prefecture's location and population (as measured in the 2000 census). A total of 83 HSR lines opened between these years, with the first in 2008; a further 66 lines were completed or under construction as of April 2019.³⁴ We compute a simple market access measure in each prefecture i and year t based on the formula in Zheng and Kahn (2013): $MA_{it} = \sum_j \exp(-0.02\tau_{ijt}) \cdot Pop_{j,2000}$, where $Pop_{j,2000}$ denotes the year-2000 population of prefecture j and τ_{ijt} denotes predicted travel time between regions i and j in year t (in minutes). Travel time predictions are based on the operational speed of each HSR line as well as geographic distance, which proxies for the travel time by car or a low-speed train. We relate MA growth, $x_i = \log MA_{i,2016} - \log MA_{i,2007}$, to the corresponding growth in

³¹We note that this heterogeneous effects extension applies to identification but not randomization-based confidence intervals which, as noted above, require a sharp null hypothesis $\beta_i = b$ for all i .

³²Construction was started by the Medium- and Long-Term Railway Plan in 2004; this plan was later expanded in 2008 and again in 2016.

³³Most prefectures are officially called “prefecture-level cities,” but typically include multiple urban areas.

³⁴We define a line by a contiguous set of inter-prefecture HSR links that were proposed together and opened simultaneously. One pilot HSR line between Qinhuangdao and Shenyang opened in 2003. We include it in our market access measure but focus on the bulk of HSR growth over 2007–2016.

prefecture's urban employment y_i from Chinese City Statistical Yearbooks. This yields a set of 275 prefectures with non-missing outcome data; see Appendix A for details on the sample construction and MA measure. Panel A of Figure 1 shows the Chinese HSR network as of the end of 2016, along with the implied MA growth of relative to 2007.

Column 1 of Table I, Panel A, reports the coefficient from a regression of employment growth on MA growth.³⁵ The estimated elasticity of 0.23 is large. With an average MA growth of 0.54 log points, it implies a 12.4% employment growth attributable to HSR for an average prefecture—almost half of the 26.6% average employment growth. The estimate is also highly statistically significant using Conley (1999) spatially-clustered standard errors.

Panel A of Figure 1, however, gives reason for caution against causally interpreting the OLS coefficient. Prefectures with high MA growth, which serve as the effective treatment group, tend to be clustered in the main economic areas in the southeast of the country where HSR lines and large markets are concentrated. A comparison between these prefectures and the economic periphery may be confounded by the effects of unobserved policies, both contemporaneous and historical, that differentially affected the economic center.

We quantify the systematic nature of spatial variation in MA growth in Column 1 of Table II, by regressing it on a prefecture's distance to Beijing, latitude, and longitude. These predictors capture over 80% of the variation in MA growth (as measured by the regression's R^2), reinforcing the OVB concern: for a causal interpretation of the Table I regression, one would need to assume that all unobserved determinants of employment growth (e.g. local productivity shocks) are uncorrelated with these geographic features. While one could of course control for the specific geographic variables from Table II (as we explore below), controlling perfectly for geography is impossible without removing all variation in x_i .

Our solution is to view certain features of the HSR network as realizations of a natural experiment. By specifying a set of counterfactual HSR networks we can compute the appropriate function of geography μ_i which removes the systematic variation in MA growth.

Our specification of counterfactuals exploits the heterogeneous timing of HSR construction. Specifically, we permute the 2016 completion status of the built and unbuilt (but planned) lines, assuming that the timing of line completion is conditionally as-good-as-random. Panel B of Figure 1 compares the built and unbuilt lines which form our counterfactuals. Unbuilt lines tend to be concentrated in the same areas of China as built lines, reinforcing the fact that construction is not uniformly distributed in space. Moreover, built lines tend to connect more regions: the average number of cross-prefecture "links" is 3.19 and 2.44 for built and unbuilt lines, respectively, with a statistically significant difference ($p = 0.048$). To account for this difference we construct counterfactual upgrades by permuting the 2016 completion status only among lines with the same number of links. For example the main Beijing to Shanghai HSR line, which has the greatest number of links, is always included in the

³⁵This regression can be viewed as a reduced form of a hypothetical IV regression, in which the treatment is a measure of market access that accounts for changes in population. We focus on the reduced form because of data constraints: we only observe the population of all 340 prefectures in the 2000 Census.

counterfactuals. This procedure generates 1,999 counterfactual HSR maps that are visually similar to the actual 2016 network; Appendix Figure A1 gives an illustrative example.

Columns 2–4 of Table II validate this specification of the HSR assignment process by the test described in Section 3.5. Column 2 shows that this recentering successfully removes the systematic geographic variation in market access. Specifically, we regress recentered MA growth on a constant and the same geographic controls as in Column 1. The regression coefficients and R^2 fall dramatically relative to Column 1, while a permutation-based p-value for their joint significance (based on the regression's sum-of-squares, as suggested in footnote 30) is 0.44. Columns 3 and 4 further show that recentered MA growth is uncorrelated with expected MA growth.³⁶

Figure 2 plots expected and recentered MA growth given by the permutations of built and unbuilt lines. The effect of recentering is apparent by contrasting the solid and striped regions in Panel B of Figure 2 (indicating high and low recentered MA growth) with the dark- and light-shaded regions in Panel A of Figure 1 (indicating high and low MA growth). The recentered treatment no longer places western prefectures in the effective control group, as their MA growth is as low as expected. Similarly, some prefectures in the east (such as Tianjin) are no longer in the effective treatment group, as they saw an expectedly large increase in MA. At the same time, recentering provides a justification for retaining other regional contrasts. Hohhot, for example, expected higher MA growth than Harbin due to the planned connection to Beijing. This line was still under construction in 2016, however, resulting in lower MA growth in Hohhot than Harbin.

Column 2 of Table I, Panel A, shows that instrumenting MA growth with recentered MA growth reduces the estimated employment elasticity substantially, from 0.23 to 0.08. Controlling for expected MA growth yields a similar estimate of 0.07 in Column 3. Neither of the two adjusted estimates is statistically distinguishable from zero according to either Conley (1999) spatial-clustered standard errors or permutation-based inference (which yields a wider confidence interval in this setting). The difference between the unadjusted and adjusted estimates is explained by the fact that employment growth is strongly predicted by expected MA growth. In Column 3 we find a large coefficient on μ_t , of 0.32, meaning that employment grew faster in prefectures that were more highly exposed to potential HSR construction, whether or not the nearby lines were built yet.

Panel B of Table I shows that the geographic controls from Table II do not isolate the same variation as expected MA growth adjustment. Including these controls in the unadjusted regression of Column 1 yields a smaller but still economically and statistically significant coefficient of 0.13. In contrast, Columns 2 and 3 show that the finding of no significant MA effect after adjusting for μ_t is robust to including geographic controls. The μ_t adjustment alone appears sufficient to remove the geographic dependence of MA, as Table II also showed.³⁷

³⁶These results are consistent with correct specification of counterfactuals (i.e. we cannot reject Assumption 2), though we note they do not provide direct support for the exogeneity of HSR construction to the unobserved determinants of employment (Assumption 1).

While our primary interest is to illustrate the recentering approach, we note that there are several possible explanations for the substantive finding of a small employment effect of MA. Unlike other transportation networks used for trading goods, the Chinese HSR network primarily operates passenger trains. Its scope for directly affecting production is therefore smaller, although it could still facilitate cross-regional business relationships. In addition, the employment effects of growing market access could be positive for some regions but negative for others, as easier commuting between regions relocates employers. We leave analyses of such mechanisms and heterogeneity for future study.

In BH we discuss how market access recentering relates to other approaches in the long literature estimating transportation infrastructure upgrade effects (Redding and Turner, 2015). We first contrast the well-known challenge of strategically chosen transportation upgrades with the less discussed problem that regional exposure to exogenous upgrades may be unequal. We then explain how common strategies to address the former issue (e.g. by leveraging historical routes or inconsequential places) can be incorporated in our framework, at least in principle. At the same time, we highlight that recentering may still be needed to address the latter issue. We further discuss how some of the existing approaches naturally yield specifications of counterfactual networks (e.g. the placebos in Donaldson (2018) and Ahlfeldt and Feddersen (2018)) and summarize the conceptual and practical advantages of our approach relative to employing more conventional controls. We emphasize that even when it is challenging to obtain a convincing specification of counterfactuals, *any* specification can yield a robustness check on these alternative strategies (see footnote 21).

5 Conclusion

Many studies in economics use treatments or instruments which combine multiple sources of variation, which are sometimes observed at different “levels,” according to a known formula. We develop a general approach to causal inference when some, but not all, of this variation is exogenous. Non-random exposure to the exogenous shocks can bias conventional regression estimators, but this problem can be solved by specifying a shock assignment process: namely, a set of counterfactual shocks that might as well have been realized. Averaging the treatment or instrument over these counterfactuals yields a single μ_i which can be adjusted for to achieve identification and consistency. The specification of counterfactuals also yields a natural form of valid finite-sample inference.

In practice, researchers face a choice in how to use μ_i in a regression analysis: recentering by it or controlling for it. When the assignment process is given by a true randomization protocol, as in a RCT, we recommend researchers recenter first to purge OVB. Then *any* predetermined controls (i.e. functions of exposure) can be included to remove variation in the error term and likely increase estimation efficiency. While μ_i is one possible control, which automatically recenters the treatment or instrument, it need not be the best choice in terms of predicting the residual variation. Our recommendation is different in natural

³⁷In BH we provide additional robustness checks, adjusting the definitions of MA and outcome variables, using a binary measure of connectivity to the HSR network, including province fixed effects, dropping influential prefectures, and examining the role of treatment effect heterogeneity.

experiments where assumptions must be placed on the assignment process. Then controlling for candidate μ_i instead of recentering can have a valuable “double-robustness” property. Researchers can compute and control for several candidate μ_i based on different assignment processes, such that OVB is purged if at least one of the processes is specified correctly (or if there is no OVB to begin with).

We conclude by noting that our framework bears practical lessons for a range of common treatments and instruments, well beyond the market access measure in our empirical application. In our working paper (Borusyak and Hull, 2021), we discuss and illustrate some of these implications for policy eligibility treatments, network spillover treatments, linear and nonlinear shift-share instruments, model-implied instruments, instruments from centralized school assignment mechanisms, “free-space” instruments for mass media access, and weather instruments. We expect other settings may also benefit from explicit specification of shock counterfactuals and appropriate adjustment for non-random shock exposure.

A: Data Appendix

Our analysis of market access effects uses data on 340 prefectures of mainland China. This excludes the islands of Hainan and Taiwan and the special administrative regions of Hong Kong and Macau, but includes six sub-prefecture-level cities (e.g. Shihezi) that do not belong to any prefecture. We use United Nations shapefiles to geocode each prefecture by the location of its main city (or, in a few cases, by the prefecture centroid).³⁸

We use a variety of sources to assemble a comprehensive database of the HSR network in 2016 as well as the lines planned (and in many cases under construction) as of April 2019 but not opened yet by the end of 2016. Our starting points are Map 1.2 of Lawrence et al. (2019), China Railway Yearbooks (China Railway Yearbook Editorial Board, 2001-2013), and the replication files of Lin (2017). We cross-check network links across these sources and use Internet resources such as Wikipedia and Baidu Baike to confirm and fill in missing information. Our database includes various types of HSR lines, including the National HSR Grid (4+4 and 8+8) and high-speed intercity railways. However, we only consider newly built HSR lines, excluding traditional lines upgraded to higher speeds. We do not put further restrictions on the class of trains (e.g. to G- and D-classes only) or specify an explicit minimum speed. The operating speed therefore ranges between 160 and 380kph, although the majority of lines are at 250kph. For each line we collect the date of its official opening (if it has opened), the actual or planned operating speed, and the list of prefecture stops. When different sections of the same line opened in a staggered way, we classify each section as a separate line for the purposes of constructing our 1,999 counterfactuals, following the definition of a line in footnote 34. We include only one contiguous stop per prefecture and drop lines that do not cross prefecture borders.

We compute travel time τ_{ij} between all pairs of prefectures i and j as of the ends of 2007 and 2016 for both the actual and counterfactual networks. Travel time combines

³⁸The shapefiles are obtained from OCHA Regional Office for Asia and the Pacific (2018, 2020), accessed on April 4, 2020.

traditional modes of transportation (car or low-speed train) with HSR, where available. We allow for unlimited changes between different HSR lines and between HSR and traditional modes without a layover penalty, as HSR trains tend to operate frequently and traditional modes also involve downtime. Following the existing literature, we proxy for travel time by traditional modes by the straight-line distance, and specify the speed of $100 = 120/1.2$ kph, where 120kph is their typical speed and the 1.2 adjustment for actual routes that are longer than a straight line. For two prefectures connected by an HSR line, we compute the distance along the line as the sum of straight-line distances between adjacent prefectures on the line. We use the operating speed of each line divided by an adjustment factor of 1.3 to capture the fact that the average speed is lower than the nominal speed we record. Computing MA further requires the population of each of the 340 prefectures from the 2000 population Census, which we obtain from Brinkhoff (2018).³⁹

We measure prefecture employment in the 2008–2017 China City Yearbooks (China Statistics Press, 2000–2017).⁴⁰ Each yearbook covers the previous year (so our data cover 2007–2016). While the yearbooks provide several employment variables, we use “The Average Number of Staff and Workers” (from the “People’s Living Conditions and Social Security” chapter), as measured in the entire prefecture and not just the main urban core. This employment series has by far the lowest number of strong year-to-year deviations which may indicate data quality issues.

We finally apply a data cleaning procedure to the outcome variable. We first mark a prefecture-year observation as exhibiting a “structural break” if (i) the outcome changes by more than twice in either direction relative to the previous non-missing value for the prefecture, (ii) it is not followed by a change in the opposite direction that is between 3/4 and 4/3 as large in terms of log-changes (which we view as a one-off jump and ignore), and (iii) the previous change does not satisfy (i). We view the outcome change between 2007 and 2016 as valid only if there are no structural breaks in any year in between. This reduces the sample from 283 to the final set of 275 prefectures.

³⁹Accessed on November 20, 2018.

⁴⁰Data for 2008–2015, excluding 2009 and 2011, are from <http://oversea.cnki.net.proxy.uchicago.edu/kns55/default.aspx>; data from 2009, 2011, 2016, and 2017, are from <http://tongji.oversea.cnki.net/chn/navi/HomePage.aspx?id=N2018050234&name=YZGCA> (all accessed on January 23, 2019 via a University of Chicago portal).

B: Additional Exhibits

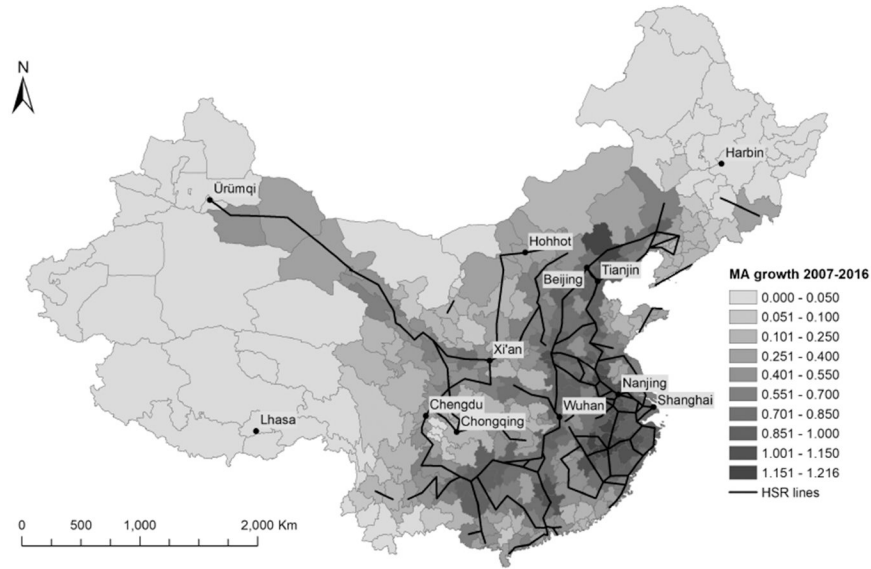


Figure A1: Simulated HSR Lines and Market Access Growth

Notes: This figure shows an example map of simulated Chinese HSR lines and market access growth over 2007–2016, obtained by permuting the opening status of built and unbuilt lines with the same number of cross-prefecture links.

C: Additional Results

Throughout the results and later proofs we omit the phrase “almost surely with respect to w ” for brevity. We also abbreviate weak mutual dependence (Assumption 3) as WMD.

Proposition 2 (Convergence for all errors implies WMD). Suppose $\text{Var}_{P_N}[\tilde{z}_i | w] \leq U_z$ uniformly. If, for some $U_\epsilon > 0$, $\text{Var}_{P_N}\left[\frac{1}{N} \sum_i \tilde{z}_i \epsilon_i\right] \rightarrow 0$ for every sequence of distributions of ϵ such that $\text{Var}_{P_N}[\epsilon_i] \leq U_\epsilon$ for all N and $i = 1, \dots, N$ then WMD holds.

Proposition 3 (WMD implies growing number of shocks). Suppose the support of g_k is bounded uniformly across across N and $k = 1, \dots, K_N$. Suppose further that, uniformly across N , $\tilde{f}_i(g; w) \equiv f_i(g; w) - \mathbb{E}_{P_N}[f_i(g; w) | w]$ is Lipschitz-continuous in g with the Lipschitz constant below U_{Lip} and that $\text{Var}_{P_N}[\tilde{f}_i(g; w) | w] \geq L_z > 0$ for at least $N \cdot L_V$ units i , with $L_V > 0$. Then WMD implies $K_N \rightarrow \infty$.

Proposition 4 (WMD and dispersed shock exposure). Suppose $f_i(g; w)$ is weakly monotone in g for all i , the components of g are jointly independent conditionally on w , and $\text{Var}_{P_N}[g_k | w] \in [L_\sigma, U_\sigma]$ for $0 < L_\sigma < U_\sigma < \infty$. Consider three cases: conditionally on w , (i) all components of g are normally distributed and $\mathbb{E}_{P_N}\left[\left|\frac{\partial \tilde{f}(g)}{\partial g_k}\right| \mid w\right] < \infty$, (ii) all components of g have the Bernoulli distribution, or (iii) \tilde{f} is linear in g . In each case:

- a. If $\sum_k \mathbb{E}_{P_N} \left[\left(\frac{\partial \tilde{f}(g; w)}{\partial g_k} \right)^2 \right] \rightarrow 0$, WMD holds;
- b. If WMD holds, $\mathbb{E}_{P_N} \left[\sum_k \left(\mathbb{E}_{P_N} \left[\frac{\partial \tilde{f}(g; w)}{\partial g_k} \mid w \right] \right)^2 \right] \rightarrow 0$,

where we define $\frac{\partial \tilde{f}(g)}{\partial g_k} \equiv \bar{f}(g_1, \dots, g_{k-1}, 1, g_{k+1}, \dots, g_{K_N}) - \bar{f}(g_1, \dots, g_{k-1}, 0, g_{k+1}, \dots, g_{K_N})$ in the Bernoulli shock case.

Proposition 5 (WMD and non-overlapping exposure sets). For each N , let $G_1(\cdot), \dots, G_N(\cdot)$ be a fixed set of functions of w to subsets of $\{1, \dots, K_N\}$ such that $f_i(\cdot; w)$ does not depend on g_k for any $k \notin G_i(w)$. Suppose the components of g are jointly independent conditionally on w , and $\text{Var}_{P_N}[\tilde{z}_i \mid w] \leq U_z$, uniformly across N and $i = 1, \dots, N$. Then WMD holds if

$$\mathbb{E}_{P_N} \left[\frac{1}{N^2} \sum_{i,j=1}^N \mathbf{1}[G_i(w) \cap G_j(w) \neq \emptyset] \right] \rightarrow 0.$$

Proposition 6 (First stage and concentrated individual exposure). Suppose $x_i = \pi z_i + u_i$ with $u_i \perp\!\!\!\perp g \mid w$, for all i and $\pi \neq 0$. Suppose further, conditionally on w , the components of g are mutually independent with $\text{Var}_{P_N}[g_k \mid w] \geq L_\sigma > 0$ uniformly across N and $k = 1, \dots, K_N$. Moreover, one of three conditions hold: (i) all components of g are mutually independent and $\mathbb{E}_{P_N} \left[\left| \frac{\partial \tilde{f}_i(g; w)}{\partial g_k} \mid w \right| \right] < \infty$, (ii) all components of g have the Bernoulli distribution, or (iii) all $\tilde{f}_i(\cdot; w)$ are linear in g . Then $\mathbb{E}_{P_N} \left[\frac{1}{N} \sum_i \sum_k \left(\mathbb{E}_{P_N} \left[\frac{\partial \tilde{f}_i(g; w)}{\partial g_k} \mid w \right] \right)^2 \right] \geq L_{HH1} > 0$ uniformly across N , $\mathbb{E}_{P_N} \left[\frac{1}{N} \sum_i x_i \tilde{z}_i \right]$ is also uniformly bounded away from zero (by $|\pi| \cdot L_{HH1} L_\sigma$).

D: Proofs

We drop the P_N subscripting of moments for all proofs to simplify notation. We again abbreviate weak mutual dependence (Assumption 3) as WMD.

Proof of Proposition 1. By Assumption 1 and the Cauchy-Schwarz inequality

$$\begin{aligned} \text{Var} \left[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i \right] &= \mathbb{E} \left[\left(\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i \right)^2 \right] = \frac{1}{N^2} \sum_{i,j} \mathbb{E}[\tilde{z}_i \tilde{z}_j \varepsilon_i \varepsilon_j] \\ &= \frac{1}{N^2} \sum_{i,j} \mathbb{E}[\mathbb{E}[\tilde{z}_i \tilde{z}_j \mid w] \mathbb{E}[\varepsilon_i \varepsilon_j \mid w]] \\ &\leq \frac{1}{N^2} \sum_{i,j} \mathbb{E} \left[|\mathbb{E}[\tilde{z}_i \tilde{z}_j \mid w]| \sqrt{\mathbb{E}[\varepsilon_i^2 \mid w] \mathbb{E}[\varepsilon_j^2 \mid w]} \right] \\ &\leq U_\varepsilon \mathbb{E} \left[\frac{1}{N^2} \sum_{i,j} |\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w]| \right] \rightarrow 0. \end{aligned} \tag{5}$$

When μ_i is approximated by $\frac{1}{S} \sum_{s=1}^S z_i^{(s)}$ for $z_i^{(s)} = f_i(g^{(s)}; w)$ and for a finite number S of random draws $g^{(s)}$ from $G(g | w)$, the same argument holds with a variance upper bound that is at most twice as large. Indeed, since

$$\begin{aligned} & \mathbb{E} \left[\left(z_i - \frac{1}{S} \sum_{s=1}^S z_i^{(s)} \right) \left(z_j - \frac{1}{S} \sum_{s=1}^S z_j^{(s)} \right) \mid w \right] \\ &= \text{Cov} \left[z_i - \frac{1}{S} \sum_{s=1}^S z_i^{(s)}, z_j - \frac{1}{S} \sum_{s=1}^S z_j^{(s)} \mid w \right] + \mathbb{E} \left[z_i - \frac{1}{S} \sum_{s=1}^S z_i^{(s)} \mid w \right] \mathbb{E} \left[z_j - \frac{1}{S} \sum_{s=1}^S z_j^{(s)} \mid w \right] \\ &= \frac{S+1}{S} \text{Cov} [z_i, z_j \mid w] = \frac{S+1}{S} \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w], \end{aligned}$$

we have, repeating the steps in (5),

$$\text{Var} \left[\frac{1}{N} \sum_i \varepsilon_i \left(z_i - \frac{1}{S} \sum_{s=1}^S z_i^{(s)} \right) \right] \leq \frac{S+1}{S} U_\varepsilon \mathbb{E} \left[\frac{1}{N^2} \sum_{i,j} | \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w] | \right] \rightarrow 0.$$

Proof of Proposition 2. For each N consider $\varepsilon = \sqrt{U_\varepsilon / U_z} \cdot \tilde{\varepsilon}$ where $(\tilde{\varepsilon}, w)$ is distributed as (\tilde{z}, w) , and $\tilde{\varepsilon} \perp\!\!\!\perp g \mid w$. Then $\text{Var} [\varepsilon_i] = U_\varepsilon \text{Var} [\tilde{z}_i] / U_z \leq U_\varepsilon$. Moreover,

$$\begin{aligned} \text{Var} \left[\frac{1}{N} \sum_i \tilde{z}_i \varepsilon_i \right] &= \frac{1}{N^2} \sum_{i,j=1}^N \mathbb{E} [\mathbb{E} [\tilde{z}_i \tilde{z}_j \mid w] \mathbb{E} [\varepsilon_i \varepsilon_j \mid w]] \\ &= \frac{U_\varepsilon}{U_z} \cdot \frac{1}{N^2} \sum_{i,j=1}^N \mathbb{E} [\mathbb{E} [\tilde{z}_i \tilde{z}_j \mid w]^2] \\ &= \frac{U_\varepsilon}{U_z} \cdot \frac{1}{N^2} \sum_{i,j=1}^N \mathbb{E} [\text{Cov} [\tilde{z}_i \tilde{z}_j \mid w]^2] \rightarrow 0. \end{aligned}$$

By the Cauchy-Schwarz inequality,

$$\frac{1}{N^2} \sum_{i,j=1}^N | \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w] | \leq \left(\frac{1}{N^2} \sum_{i,j=1}^N \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w]^2 \right)^{0.5}.$$

And by Jensen's inequality,

$$\begin{aligned} & \mathbb{E} \left[\frac{1}{N^2} \sum_{i,j=1}^N | \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w] | \right] \leq \mathbb{E} \left[\left(\frac{1}{N^2} \sum_{i,j=1}^N \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w]^2 \right)^{0.5} \right] \\ & \leq \mathbb{E} \left[\frac{1}{N^2} \sum_{i,j=1}^N \text{Cov} [\tilde{z}_i, \tilde{z}_j \mid w]^2 \right]^{0.5} \rightarrow 0 \end{aligned}$$

Proof of Proposition 3.⁴¹ We prove this result by contradiction. Without loss of generality suppose $K_N = K$ is constant along the asymptotic sequence; whenever $K_N \not\rightarrow \infty$, there is a subsequence of K_N bounded by some K , and the proof follows for that subsequence without change. Also without loss, we condition on w and suppress the w notation. We denote the upper bound on the support of $|g_k|$ by $U_g / 2$ and extend the domain of each \tilde{f}_i to $[-U_g / 2, U_g / 2]^K$ preserving its Lipschitz constant, by the Kirszbraun theorem.

⁴¹We thank Mikhail Dekhtarev for help with this proof.

Let $\mathbb{R}(x) = \min \{ \lceil x / \delta \rceil \cdot \delta, U_g / 2 \}$ denote the upward-rounding function for some $\delta > 0$. Consider $\check{f}_i(g) = \mathbb{R}(\tilde{f}_i(\mathbb{R}(g_1), \dots, \mathbb{R}(g_k)))$ which rounds both the shocks and the values of \tilde{f}_i . Note that

$$| \check{f}_i(g) - \tilde{f}_i(g) | \leq | \check{f}_i(g) - \tilde{f}_i(\mathbb{R}(g_1), \dots, \mathbb{R}(g_k)) | + | \tilde{f}_i(\mathbb{R}(g_1), \dots, \mathbb{R}(g_k)) - \tilde{f}_i(g) | \leq \delta + \delta U_{Lip} \sqrt{K},$$

where the second inequality uses the Lipschitz condition and $\|g - (\mathbb{R}(g_1), \dots, \mathbb{R}(g_k))\|_2 \leq \delta \sqrt{K}$ since $|g_k - \mathbb{R}(g_k)| \leq \delta$ for each k .

By the Lipschitz condition $| \tilde{f}_i(g_A) - \tilde{f}_i(g_B) | \leq U_{Lip} U_g \sqrt{K}$ for any $g_A, g_B \in [-U_g / 2, U_g / 2]^K$. Since $\mathbb{E} | \tilde{f}_i(g) | = 0$ and $g \in [-U_g / 2, U_g / 2]^K$, this implies $| \tilde{f}_i(g) | \leq U_{Lip} U_g \sqrt{K}$ and consequently $| \check{f}_i(g) | \leq U_{Lip} U_g \sqrt{K} + \delta + \delta U_{Lip} \sqrt{K}$. Thus, for all N and $i = 1, \dots, N$ there is only a finite number U_R of possible ‘‘rounded’’ $\check{f}_i(\cdot)$ functions. Therefore, at least $N L_V / U_R$ of observations i with $\text{Var}[\check{f}_i(g)] \geq L_z$ have the same rounded function, and thus there are at least $(N L_V / U_R)^2$ such pairs of observations (i, j) . For any such pair, since $\mathbb{E}[\check{f}_i(g) - \check{f}_j(g)] = 0$,

$$\begin{aligned} 2\text{Cov}[\check{f}_i(g), \check{f}_j(g)] &= \text{Var}[\check{f}_i(g)] + \text{Var}[\check{f}_j(g)] - \mathbb{E}[(\check{f}_i(g) - \check{f}_j(g))^2] \\ &\geq 2L_z - \mathbb{E}\left[\left(|\check{f}_i(g) - \check{f}_j(g)| + |\check{f}_i(g) - \check{f}_j(g)| + |\check{f}_i(g) - \check{f}_j(g)|\right)^2\right] \\ &\geq 2L_z - \mathbb{E}\left[\left(\delta(1 + U_{Lip}\sqrt{K}) + 0 + \delta(1 + U_{Lip}\sqrt{K})\right)^2\right]. \end{aligned}$$

Setting $\delta = \frac{\sqrt{L_z}}{2(1 + U_{Lip}\sqrt{K})}$, we have $\text{Cov}[\check{f}_i(g), \check{f}_j(g)] \geq \frac{L_z}{2} > 0$, and therefore

$$\frac{1}{N^2} \sum_{i, j=1}^N |\text{Cov}[\check{f}_i(g), \check{f}_j(g)]| \geq \frac{1}{N^2} \cdot \left(\frac{N L_V}{U_R}\right)^2 \frac{L_z}{2} = \frac{L_V^2 L_z}{2 U_R^2} \rightarrow 0.$$

Thus, weak mutual dependence does not hold, establishing the contradiction.

To establish Proposition 4, we first state and prove four lemmas. We assume all moments relevant for those lemmas exist.

Lemma 1. $h : \mathbb{R}^K \rightarrow \mathbb{R}$ is weakly increasing and random variables g_1, \dots, g_k are independent, then for any $k \in \{1, \dots, K - 1\}$ the conditional expectation $\mathbb{E}[h(g_1, \dots, g_k) \mid g_1, \dots, g_k]$ is weakly increasing.

Proof. Fix $\gamma_1, \dots, \gamma_k$ and $\gamma'_1, \dots, \gamma'_k$ such that $\gamma'_r \geq \gamma_r$ for $r = 1, \dots, k$, and define the $K \times 1$ vectors $g = (\gamma_1, \dots, \gamma_k, g_{k+1}, \dots, g_K)$ and $g' = (\gamma'_1, \dots, \gamma'_k, g_{k+1}, \dots, g_K)$. Note $h(g') \geq h(g)$. For $r = k + 1, \dots, K$, denote the cumulative distribution function of g_r by $G_r(\cdot)$. Then

$$\begin{aligned} \mathbb{E}[h(g')] &= \int \cdots \int h(g') dG_{k+1}(g_{k+1}) \cdots dG_K(g_K) \\ &\geq \int \cdots \int h(g) dG_{k+1}(g_{k+1}) \cdots dG_K(g_K) = \mathbb{E}[h(g)]. \end{aligned}$$

Lemma 2. For any weakly increasing $h_1, h_2: \mathbb{R}^K \rightarrow \mathbb{R}$, $\text{Cov}[h_1(g), h_2(g)] \geq 0$ for $g = (g_1, \dots, g_K)$ with independent components.

Proof. For $K = 1$ this is well known. The proof for $K > 1$ follows by induction. Suppose it is true for $K - 1$. Then by the law of total covariance

$$\text{Cov}[h_1(g), h_2(g)] = \mathbb{E}[\text{Cov}[h_1(g), h_2(g) \mid g_1]] + \text{Cov}[\mathbb{E}[h_1(g) \mid g_1], \mathbb{E}[h_2(g) \mid g_1]].$$

The first term is the expectation of a covariance of two monotone (by Lemma 1) functions of $K - 1$ variables. The second term, again by Lemma 1, is a covariance of two monotone functions of random scalars. Thus both terms are non-negative.

Lemma 3. If $f_i(g; w)$ is weakly monotone in g for all i and components of g are jointly independent conditionally on w , then $\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w] \geq 0$ for all i and j . Furthermore WMD simplifies to $\text{Var}[\bar{z}] \rightarrow 0$ for $\bar{z} = \frac{1}{N} \sum_i \tilde{z}_i$.

Proof. Applying Lemma 2 to $\tilde{z}_i = f_i(g; w) - \mathbb{E}[f_i(g; w) \mid w]$ and $\tilde{z}_j = f_j(g; w) - \mathbb{E}[f_j(g; w) \mid w]$ (or their negations, if $f_i(g; w)$ is weakly decreasing) and conditioning on w everywhere, we obtain $\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w] \geq 0$. Thus, WMD simplifies to

$$\begin{aligned} \mathbb{E}\left[\frac{1}{N^2} \sum_{i,j} \text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w]\right] &= \mathbb{E}\left[\frac{1}{N^2} \sum_{i,j} \text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w]\right] \\ &= \mathbb{E}\left[\text{Var}\left[\frac{1}{N} \sum_i \tilde{z}_i \mid w\right]\right] = \text{Var}\left[\frac{1}{N} \sum_i \tilde{z}_i\right] \rightarrow 0, \end{aligned}$$

where the second line rearranges terms and the third line follows by $\mathbb{E}[\frac{1}{N} \sum_i \tilde{z}_i \mid w] = 0$.

Lemma 4. Suppose $g = (g_1, \dots, g_K)$ is jointly independent with $\sigma_k^2 \equiv \text{Var}[g_k]$ and consider a scalar function h on the support of g . Then if (i) all components of g are normally distributed and $\mathbb{E}\left[\left|\frac{\partial h(g)}{\partial g_k}\right|\right] < \infty$ or (ii) all components of g have the Bernoulli distribution,

$$\sum_k \sigma_k^2 \left(\mathbb{E}\left[\frac{\partial h(g)}{\partial g_k}\right]\right)^2 \leq \text{Var}[h(g)] \leq \sum_k \sigma_k^2 \mathbb{E}\left[\left(\frac{\partial h(g)}{\partial g_k}\right)^2\right], \tag{6}$$

with $\frac{\partial h}{\partial g_k}$ defined in the Bernoulli case as in Proposition 4. Further, (iii) if h is linear, (6) holds with equalities, regardless of the distributions of the components of g .

Proof. For part (i), the lower bound is established by Cacoullos (1982, Proposition 3.7), and the upper bound on $\text{Var}[h(g)]$ is established by Chen (1982, Corollary 3.2). For part (ii), the lower bound follows from restricting the results for binomial distributions in Cacoullos and Papathanasiou (1989, p. 355), and the upper bound is similarly a special case of the result in Cacoullos and Papathanasiou (1985, p. 183). Part (iii) follows trivially from the fact that $\partial h / \partial g_k$ is non-stochastic.

Proof of Proposition 4. By Lemma 3, WMD is equivalent to

$\text{Var}[\bar{f}(g; w)] = \mathbb{E}[\text{Var}[\bar{f}(g; w) \mid w]] \rightarrow 0$. Applying Lemma 4 conditionally on w and using the bounds on $\text{Var}[g_k \mid w]$,

$$\mathbb{E}\left[\sum_k L_\sigma \left(\mathbb{E}\left[\frac{\partial \bar{f}(g; w)}{\partial g_k} \mid w\right]\right)^2\right] \leq \text{Var}[\bar{f}(g; w)] \leq \mathbb{E}\left[\sum_k U_\sigma \mathbb{E}\left[\left(\frac{\partial \bar{f}(g; w)}{\partial g_k}\right)^2 \mid w\right]\right].$$

The upper bound, the law of iterated expectations, and $U_\sigma > 0$ imply that if

$$\mathbb{E}\left[\sum_k \mathbb{E}\left[\left(\frac{\partial \bar{f}(g; w)}{\partial g_k}\right)^2\right]\right] \rightarrow 0, \text{Var}[\bar{f}(g; w)] \rightarrow 0, \text{ and thus WMD holds. The lower bound and } L_\sigma > 0$$

imply that if WMD holds and thus $\text{Var}[\bar{f}(g; w)] \rightarrow 0$, we have $\mathbb{E}\left[\sum_k \left(\mathbb{E}\left[\frac{\partial \bar{f}(g; w)}{\partial g_k} \mid w\right]\right)^2\right] \rightarrow 0$.

Proof of Proposition 5. For any N and fixed \bar{w} in the support of w and for i and j such that $G_i(\bar{w}) \cap G_j(\bar{w}) = \emptyset$ we have $\tilde{z}_i \perp \tilde{z}_j \mid w = \bar{w}$ because f_i and f_j are functions of two non-overlapping subvectors of g , the components of which are conditionally independent. Thus $\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w = \bar{w}] = 0$ for such (i, j) pairs, and we obtain

$$\begin{aligned} \frac{1}{N^2} \sum_{i,j} |\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w]| &= \frac{1}{N^2} \sum_{i,j} \mathbf{1}[G_i(w) \cap G_j(w) \neq \emptyset] \mathbb{E}[|\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w]|] \\ &\leq \frac{1}{N^2} \sum_{i,j} \mathbf{1}[G_i(w) \cap G_j(w) \neq \emptyset] \mathbb{E}[\sqrt{\text{Var}[\tilde{z}_i \mid w] \text{Var}[\tilde{z}_j \mid w]}] \\ &\leq U_z \cdot \frac{1}{N^2} \sum_{i,j} \mathbf{1}[G_i(w) \cap G_j(w) \neq \emptyset] \end{aligned}$$

and therefore $\mathbb{E}\left[\frac{1}{N^2} \sum_{i,j} |\text{Cov}[\tilde{z}_i, \tilde{z}_j \mid w]| \right] \rightarrow 0$.

Proof of Proposition 6. By the law of iterated expectations,

$$\mathbb{E}\left[\frac{1}{N} \sum_i x_i \tilde{z}_i\right] = \pi \mathbb{E}\left[\frac{1}{N} \sum_i \tilde{z}_i^2\right] = \pi \mathbb{E}\left[\frac{1}{N} \sum_i \text{Var}[\tilde{z}_i \mid w]\right]. \text{ The result then follows directly from}$$

Lemma 4 applied to each $\tilde{f}_i(g; w)$:

$$\frac{1}{N} \sum_i \text{Var}[\tilde{z}_i \mid w] \geq \frac{1}{N} \sum_i \sum_k \text{Var}[g_k \mid w] \left(\mathbb{E}\left[\frac{\partial \tilde{f}_i(g; w)}{\partial g_k} \mid w\right]\right)^2 \geq L_\sigma \cdot \frac{1}{N} \sum_i \sum_k \left(\mathbb{E}\left[\frac{\partial \tilde{f}_i(g; w)}{\partial g_k} \mid w\right]\right)^2,$$

and thus $\left|\mathbb{E}\left[\frac{1}{N} \sum_i x_i \tilde{z}_i\right]\right| \geq |\pi| \cdot L_\sigma L_{HHI}$.

References

- Abadie A. (2003): “Semiparametric instrumental variable estimation of treatment response models,” *Journal of Econometrics*, 113, 231–263.
- Abadie A, Athey S, Imbens GW, and Wooldridge JM (2020): “Sampling-based vs. Design-based Uncertainty in Regression Analysis,” *Econometrica*, 88, 265–296.
- Abadie A, and Imbens GW (2016): “Matching on the Estimated Propensity Score,” *Econometrica*, 84, 781–807.
- Abdulkadiroglu A, Angrist JD, Narita Y, and Pathak PA (2017): “Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation,” *Econometrica*, 85, 1373–1432.
- Adão R, Kolesár M, and Morales E (2019): “Shift-Share Designs: Theory and Inference,” *Quarterly Journal of Economics*, 134, 1949–2010.
- Ahlfeldt GM, and Feddersen A (2018): “From periphery to core: Measuring agglomeration effects using high-speed rail,” *Journal of Economic Geography*, 18, 355–390.
- Aronow PM, and Samii C (2017): “Estimating average causal effects under general interference, with application to a social network experiment,” *Annals of Applied Statistics*, 11, 1912–1947.
- Athey S, Bayati M, Doudchenko N, Imbens GW, and Khosravi K (2021): “Matrix Completion Methods for Causal Panel Data Models,” *Journal of the American Statistical Association*, 116, 1716–1730.
- Athey S, and Imbens GW (2022): “Design-based Analysis in Difference-In-Differences Settings with Staggered Adoption,” *Journal of Econometrics*, 226, 62–79.
- Borusyak K, and Hull P (2021): “Non-Random Exposure to Exogenous Shocks: Theory and Applications,” Mimeo.
- Borusyak K, Hull P, and Jaravel X (2022): “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 89, 181–213. [PubMed: 37363108]
- Brinkhoff T. (2018): “City Population,” <http://www.citypopulation.de>.
- Cacoullos T. (1982): “On Upper and Lower Bounds for the Variance of a Function of a Random Variable,” *The Annals of Probability*, 10, 799–809.
- Cacoullos T, and Papathanasiou V (1985): “On upper bounds for the variance of functions of random variables,” *Statistics and Probability Letters*, 3, 175–184.
- (1989): “Characterizations of Distributions by Variance Bounds,” *Statistics and Probability Letters*, 7, 351–356.
- Carvalho VM, Nirei M, Saito YU, and Tahbaz-Salehi A (2021): “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake,” *Quarterly Journal of Economics*, 136, 1255–1321.
- Cattaneo MD, Frandsen BR, and Titiunik R (2015): “Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate,” *Journal of Causal Inference*, 3, 1–24.
- de Chaisemartin C, and Behaghel L (2020): “Estimating the Effect of Treatments Allocated by Randomized Waiting Lists,” *Econometrica*, 88, 1453–1477.
- de Chaisemartin C, and D’Haultfœuille X (2020): “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 110, 2964–2996.
- Chen LHY (1982): “An Inequality for the Multivariate Normal Distribution,” *Journal of Multivariate Analysis*, 12, 306–315.
- China Railway Yearbook Editorial Board (2001-2013): *China Railway Yearbook*, <https://oversea.cnki.net/KNavi/YearbookDetail?pcode=CYFD&pykm=YZGTD&>.
- China Statistics Press (2000-2017): *China City Statistical Yearbook*, <https://cnki.net/KNavi/YearbookDetail?pcode=CYFD&pykm=YZGCA>.
- Conley TG (1999): “GMM estimation with cross sectional dependence,” *Journal of Econometrics*, 92, 1–45.
- Currie J, and Gruber J (1996): “Health Insurance Eligibility, Utilization of Medical Care, and Child Health,” *The Quarterly Journal of Economics*, 111, 431–466.
- Donaldson D. (2018): “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, 108, 899–934.

- Donaldson D, and Hornbeck R (2016): "Railroads and American Economic Growth: A "Market Access" Approach," *Quarterly Journal of Economics*, 131, 799–858.
- Fisher RA (1935): *The design of experiments*: Oliver & Boyd.
- Gruber J. (2003): "Medicaid," in *Means-tested transfer programs in the United States*: University of Chicago Press, 15–78.
- Hodges JJJ, and Lehmann EL (1963): "Estimates of Location Based on Rank Tests," *The Annals of Mathematical Statistics*, 34, 598–611.
- Imbens GW (2000): "The role of the propensity score in estimating dose-response functions," *Biometrika*, 87, 706–710.
- Imbens GW, and Angrist JD (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–475.
- King G, and Nielsen R (2019): "Why Propensity Scores Should Not Be Used for Matching," *Political Analysis*, 27, 435–454.
- Lawrence M, Bullock R, and Liu Z (2019): *China's High-Speed Rail Development*, Washington, D.C.: World Bank.
- Lee DS (2008): "Randomized experiments from non-random selection in U.S. House elections," *Journal of Econometrics*, 142, 675–697.
- Lehmann EL, and Romano JP (2006): *Testing statistical hypotheses*: Springer Science & Business Media.
- Lin Y. (2017): "Travel costs and urban specialization patterns: Evidence from China's high speed railway system," *Journal of Urban Economics*, 98, 98–123.
- Ma D. (2011): "China's Long, Bumpy Road to High-Speed Rail," *The Atlantic*.
- Madestam A, Shoag D, Veuger S, and Yanagizawa-Drott D (2013): "Do Political Protests Matter? Evidence from the Tea Party Movement," *Quarterly Journal of Economics*, 128, 1633–1685.
- Miguel E, and Kremer M (2004): "Worms: Identifying impacts on education and health in the presence of treatment externalities," *Econometrica*, 72, 159–217.
- Neyman J. (1923): "On the application of probability theory to agricultural experiments. Essay on principles," *Ann. Agricultural Sciences*, 1–51.
- OCHA Regional Office for Asia and the Pacific (2018): "Province and Prefecture Capitals of China," <https://data.humdata.org/dataset/province-and-prefecture-capitals-of-china>.
- (2020): "China - Subnational Administrative Boundaries," <https://data.humdata.org/dataset/cod-ab-chn>.
- Ollivier G, Bullock R, Jin Y, and Zhou N (2014): "High-Speed Railways in China: A Look at Traffic," *China Transport Topics*, 1–12.
- Redding SJ, and Turner MA (2015): "Transportation Costs and the Spatial Organization of Economic Activity," in *Handbook of regional and urban economics*: Elsevier, 1339–1398.
- Redding SJ, and Venables AJ (2004): "Economic geography and international inequality," *Journal of International Economics*, 62, 53–82.
- Robins JM, Mark SD, and Newey WK (1992): "Estimating Exposure Effects by Modelling the Expectation of Exposure Conditional on Confounders," *Biometrics*, 48, 479–495. [PubMed: 1637973]
- Robinson P. (1988): "Root-N-Consistent Semiparametric Regression," *Econometrica*, 56, 934–954.
- Rosenbaum PR (2002): "Covariance adjustment in randomized experiments and observational studies," *Statistical Science*, 17, 286–327.
- Rosenbaum PR, and Rubin DB (1983): "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, 70, 41–55.
- Shaikh A, and Toulis P (2021): "Randomization Tests in Observational Studies with Staggered Adoption of Treatment," *Journal of the American Statistical Association*, 116, 1835–1848.
- Wooldridge JM (2015): "Control function methods in applied econometrics," *Journal of Human Resources*, 50, 420–445.
- Zheng S, and Kahn ME (2013): "China's bullet trains facilitate market integration and mitigate the cost of megacity growth," *Proceedings of the National Academy of Sciences of the United States of America*, 110, 1248–1253.

A. Completed Lines and MA Growth

B. All Planned Lines

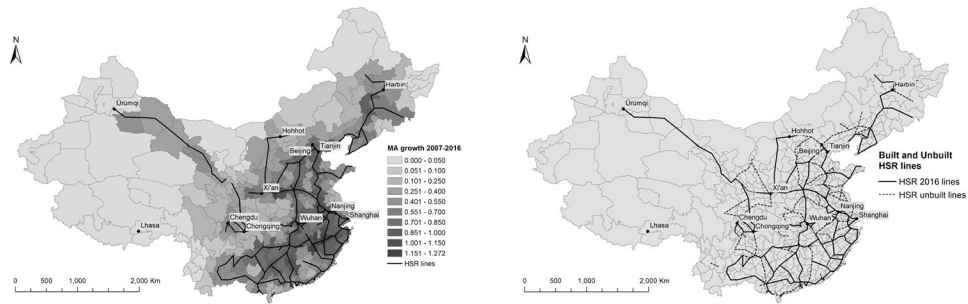


Figure 1: Chinese High Speed Rail and Market Access Growth, 2007-2016 A. Completed Lines and MA Growth B. All Planned Lines Notes: Panel A shows the completed China high-speed rail network by the end of 2016, with shading indicating MA growth (i.e. log-change in MA) relative to 2007. Panel B shows the network of all HSR lines, including those planned but not yet completed as of 2016.

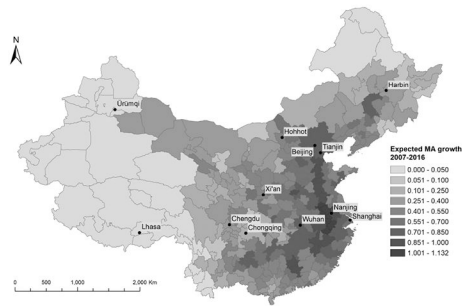
Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

A. Expected Market Access Growth



B. Recentered Market Access Growth

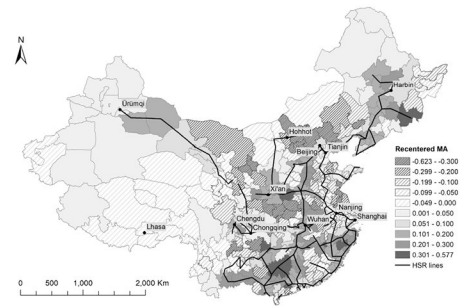


Figure 2: Expected and Recentered Market Access Growth from Chinese HSR

Notes: Panel A shows the variation in expected 2007–16 MA growth across Chinese prefectures, computed from 1,999 HSR counterfactuals that permute the opening status of built and unbuilt lines with the same number of cross-prefecture links. Panel B plots the variation in corresponding recentered MA growth: the difference between the MA growth shown in Panel A of Figure 1 and expected MA growth. The HSR network as of 2016 is also shown in this panel.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table I:

Employment Effects of Market Access: Unadjusted and Recentered Estimates

	Unadjusted OLS (1)	Recentered IV (2)	Controlled OLS (3)
Panel A: No Controls			
Market Access Growth	0.232 (0.075)	0.084 (0.097) [-0.245, 0.337]	0.072 (0.093) [-0.169, 0.337]
Expected Market Access Growth			0.317 (0.096)
Panel B: With Geography Controls			
Market Access Growth	0.133 (0.064)	0.056 (0.089) [-0.135, 0.280]	0.047 (0.092) [-0.146, 0.280]
Expected Market Access Growth			0.214 (0.073)
Recentered	No	Yes	Yes
Prefectures	275	275	275

Notes: This table reports coefficients from regressions of employment growth on MA growth in Chinese prefectures from 2007–2016. MA growth is unadjusted in Column 1. In Column 2 this treatment is instrumented by MA growth recentered by permuting the opening status of built and unbuilt HSR lines with the same number of cross-prefecture links. Column 3 instead estimates an OLS regression with recentered MA growth as treatment and controlling for expected MA growth given by the same HSR counterfactuals. The regressions in Panel B control for distance to Beijing, latitude, and longitude. Standard errors which allow for linearly decaying spatial correlation (up to a bandwidth of 500km) are reported in parentheses. 95% RI confidence intervals based on the HSR counterfactuals are reported in brackets.

Table II:

Regressions of Market Access Growth on Measures of Economic Geography

	Unadjusted		Recentered	
	(1)	(2)	(3)	(4)
Distance to Beijing	-0.291 (0.062)	0.069 (0.039)		0.088 (0.045)
Latitude/100	-3.324 (0.646)	-0.342 (0.276)		-0.182 (0.319)
Longitude/100	1.321 (0.458)	0.485 (0.237)		0.440 (0.240)
Expected Market Access Growth			0.026 (0.056)	0.054 (0.069)
Constant	0.536 (0.029)	0.018 (0.018)	0.018 (0.021)	0.018 (0.018)
Joint RI p-value		0.443	0.711	0.492
R^2	0.824	0.083	0.010	0.086
Prefectures	275	275	275	275

Notes: This table reports coefficients from regressing the unadjusted and recentered MA growth of Chinese prefectures (2007–2016) on geographic controls. Recentering is done by permuting the opening status of built and unbuilt lines with the same number of cross-prefecture links. All regressors are measured for the prefecture's main city and demeaned such that the constant in each regression captures the average outcome. Distance to Beijing is measured in 1,000km. Standard errors which allow for linearly decaying spatial correlation (up to a bandwidth of 500km) are reported in parentheses. Joint RI p-values are based on the 1,999 HSR counterfactuals and the sum-of-square fitted values statistic, as described in footnote 30.