Online Appendix for:
A Practical Guide to Shift-Share Instruments

Kirill Borusyak — UC Berkeley
Peter Hull — Brown
Xavier Jaravel — LSE

December 2024

Contents

A Other Practical Concerns with Shift-Share IVs

In this appendix we discuss some questions which come up frequently in shift-share instrumental variable (IV) designs. These twelve questions are: What if shares are observed in a panel? What kind of local average treatment effect does the shift-share IV estimate? Do the "shares" have to really be shares, between zero and one? Should the shares be normalized to add up to one? Can shift-share instruments be constructed by apportioning national changes to units? Can the shifts be unit-specific? Can one take a shift-share average of shift-share IVs? What if a log, or another transformation, of a shift-share variable is used? Can one use multiple shift-share instruments? What about interaction terms in shift-share regressions? Should the instruments in Card (2009) and Bartik (1991), which measure the shifts as the national growth of some equilibrium outcome (industry employment or total migration by origin), be viewed through the lens of exogenous shifts or exogenous shares? And what is the role of leave-one-out construction of shifts?

A.1 What if shocks are observed over multiple periods in a panel?

While our main discussion considered a single cross-section (typically, with first-differenced outcomes), in many applications researchers have access to shifts g_{kt} happening in multiple periods t. In a panel of units i over periods $t = 1, \dots, T$, one may consider an IV specification

$$y_{it} = \beta x_{it} + \gamma' w_{it} + \varepsilon_{it}, \tag{1}$$

with a shift-share IV $z_{it} = \sum_k s_{ikt} g_{kt}$ and some controls w_{it}, potentially including unit and period fixed effects. Here we indexed the shares s_{ikt} by the period when they are used, not when they are measured: for instance, s_{ikt} can be time-invariant, fixed in an early "base" period.

The panel setting offers new possibilities. In the exogenous shifts approach, if a natural experiment generates exogenous shifts in several periods, "stacking" them provides more estimation power. Moreover, panels with *many* periods make it possible to apply the many exogenous shifts approach even when K is small, thanks to the time-series variation in the shifts. In the simplest case, there may be just one shift in each period and heterogeneous unit exposure to this shift, such that $z_{it} = s_{it} g_t$ (where s_{it} is often time-invariant). When many periods are observed, exogeneity of the time series variation in g_t is sufficient for consistent estimation. Exposure-robust inference also follows from the time-series properties of the shifts. The shift-level equivalent IV regression in this case is just a time-series regression regardless of the number of units in the panel, and standard errors should correspondingly be clustered in a time-series way (e.g., by period or allowing for serial correlation in the shifts). Below we provide a detailed illustration of these points in the setting of Nunn and Qian (2014).

In the exogenous shares approach, the number of available instruments is also larger in a panel, as primitive instruments are individual shares interacted with period dummies, all assumed to be valid. Correspondingly, Rotemberg weights are computed for each pair of k and t, although Goldsmith-Pinkham et al. (2020) recommend reporting their sums over time for clarity.

Panel data also pose new challenges. We focus on the exogenous shifts approach in this discussion. First, the shifts may have different means in different periods. In conventional panel models, time-varying means are addressed by including period fixed effects (FEs), γ_t. Correspondingly, in shift-share designs, time-varying shift means are addressed by a share-weighted aggregate of period FEs: $\sum_k s_{ikt}\gamma_t$. With complete shares, i.e. when $\sum_k s_{ikt} = 1$, this control coincides with the period FEs. But in the incomplete shares case, the sum of shares control needs to be interacted with period FEs. The setting of Autor et al. (2013) discussed in the main text illustrates this point.

Second, shifts can be serially correlated, in which case each period cannot be viewed as a separate natural experiment. Then, as mentioned in Step 4 of the exogenous shifts checklist, the static specification in (1) suffers from an omitted variables bias problem when there are dynamic causal effects, i.e. if lagged shifts affect current outcomes (Jaeger et al., 2017). Intuitively, the estimated coefficient for the treatment in specification (1) is biased because it also includes the dynamic causal effect of past treatments. Moreover, if the shares can respond to past shifts which are correlated with contemporaneous shifts, the shares cannot be viewed as measured before the natural experiment in shifts began (see footnote 5).

There are two solutions to the problems of serial correlation in shifts. One involves estimating richer specifications which include the relevant lagged treatments, as well as lagging the shares underlying the shift-share IV further. The shares need to be measured at a date before the sequence of serially correlated shifts began if such a date exists. Jaeger et al. (2017), for instance, show that migration rates by country of origin are very serially correlated since 1970s, but not correlated with those from earlier decades. Thus, year 1970 can be viewed as the beginning of the natural experiment in their setting.

An alternative solution is based on isolating the unpredictable component of the contemporaneous shifts before constructing the shift-share IV. For instance, if the shifts follow a first-order autoregressive process, one can control for the lagged shifts (by controlling for a share-aggregated version of them at the unit level). If the shift-share IV leverages the idiosyncratic component of shifts, the issues stemming from serial correlation disappear.[1] This approach only yields the contemporaneous effect but does not require a correct specification of the dynamic effects.

Example of shift-share IV with time-series variation Nunn and Qian (2014) study the impact of US food aid on conflicts in a long panel of recipient countries. Simple ordinary least square (OLS) estimates, or even those with country fixed effects, are subject to several potential biases: the presence of conflict may increase the demand for food aid; there might be many omitted variables—such as political and economic crises—affecting both conflict and food aid; or donors may decide to reduce food aid to countries engaged in conflict.

To resolve these issues, the authors leverage exogenous time variation in US wheat production

[1] There are different ways of extracting the idiosyncratic component of shifts. Instead of controlling for lagged shifts, another natural approach could be to control for the time-invariant component of shifts. Implementing this strategy is easy when time-invariant shares are used: then including *unit* fixed effects in the control vector w_{it} implicitly removes any shift-level confounders α_k, since the corresponding share-aggregated control $\sum_k s_{ikt}\alpha_k$ is time-invariant.

over time. Due to price stabilization policies requiring the US government to buy wheat from US farmers at a set price, the US government accumulates excess reserves in high production years, which is shipped to developing countries as food aid. The shift-share design leverages these time series shifts, using as exposure weights a country's likelihood of being a US food aid recipient. Specifically, the quantity of wheat aid shipped from the US to recipient i in year t is instrumented by $z_{it} = s_i g_t$, where g_t is the amount of US wheat production in the previous year and s_i is the fraction of years that recipient country i receives a positive amount of US food aid during the sample period, 1971–2006.

How can one follow our exogenous shifts checklist in this context? For steps 1–2, the researcher would clarify whether all time series variation in wheat production is considered as-good-as-random. This requires an exclusion restriction, that US wheat production affects conflict in other countries only through US aid. Moreover, if US wheat production is correlated with key economic indicators such as oil prices, which can have a direct effect on conflict, the researcher would need to control for these variables interacted with s_i. Indeed, the interaction of the oil price with s_i is one of the controls Nunn and Qian (2014) include. They also include other controls, such as dummies for six geographic regions of the world interacted with year dummies. For step 3, the incomplete share control here is simply s_i, the time-invariant exposure to US aid, since each observation is exposed to only one shift; in Nunn and Qian's regression it is absorbed by country fixed effects. For step 4, one would measure s_i before, rather than during, the sample period. For step 5, it would be useful to plot the time series of wheat prices, which serves as identifying variation. Christian and Barrett (2024, Fig. 3) finds strong serial correlation and an inverse U-shaped trend in wheat prices. In this case, it may be appropriate to analyze dynamic causal effects or extract an unpredictable component of the time series of US wheat production. For step 6, one can check whether the time series of wheat production is correlated with potential confounders, such as the aforementioned oil prices. At the country-by-year level, an IV regression with lagged conflict as the outcome would constitute a standard pre-trend test. Finally, for step 7, one can cluster standard errors at the level of identifying variation, i.e. by year (rather than by country, which is more conventional in panel regressions). Given the shifts (and, likely, errors) are serially correlated, heteroskedasticity and autocorrelation-consistent standard errors may be more appropriate. Indeed, Christian and Barrett (2024) show that conventional standard errors can lead to spuriously significant relationships in this setting. These standard errors are easier to obtain from the time-series regression equivalent to the original panel regression.

A.2 Does shift-share IV estimate a LATE when the effects are heterogeneous?

With many exogenous shifts, yes, and different units and shifts receive a different weight in shift-share IV regressions. Otherwise, a local average treatment effect (LATE) interpretation is more challenging.

Provided the shifts are as-good-as-randomly assigned and mutually uncorrelated, as if arising from a lottery, Adão et al. (2019) and Borusyak et al. (2022) prove that shift-share regressions (both IV and OLS) identify convex averages of heterogeneous treatment effects under a monotonicity condition similar to the one imposed by Imbens and Angrist (1994) to establish identification of LATEs in standard IV regressions. Note that as-good-as-random assignment here is formally stronger than the necessary condition on shift exogeneity described in the main text. For example, it requires the shifts to be independent of treatment effect heterogeneity as well.

What makes the shift-share setting unique is that effect heterogeneity can arise in two dimensions: across units i and shifts k. Thus, the shift-share IV estimate can be interpreted as averaging across both dimensions, with certain weights. We derive and interpret these weights in a heterogeneous-effects causal model inspired by the decomposition formula (3).[2] For concreteness, we consider our labor supply example. The model is as follows:

$$\begin{aligned} x_{ik} &= \pi_{ik} g_k + u_{ik}, \\ x_i &= \sum_k s_{ik} x_{ik}, \\ y_i &= \sum_k \beta_{ik} s_{ik} x_{ik} + \varepsilon_i. \end{aligned}$$

Here x_{ik} are changes in employment by region and industry—the local shifts. They are affected by the industry subsidies g_k with a coefficient of π_{ik}. Regional employment growth x_i is an aggregate of industry-by-region growth rates x_{ik} weighted by regional employment shares s_{ik}, as in equation (3). But the effects of x_{ik} on the wage change y_i are not necessarily proportional to s_{ik}, as captured by the heterogeneous effects β_{ik}. Variation in β_{ik} across i captures the idea that the local labor supply elasticity may depend on the region. In turn, variation in β_{ik} across k reflects the scenario in which employment changes coming different industries (say, tradable and nontradable ones) would have different wage impacts through labor supply.

Following the logic of Adão et al. (2019) and Borusyak et al. (2022), it is easy to show that, when the shifts g_k have mean zero, variance σ_k^2, and no mutual correlation conditional on all other sources of unobserved heterogeneity $(u_{ik}, \varepsilon_i, \pi_{ik}, \beta_{ik})$, the shift-share IV estimand equals

$$\beta = \frac{\sum_i \sum_k \pi_{ik} s_{ik}^2 \sigma_k^2 \cdot \beta_{ik}}{\sum_i \sum_k \pi_{ik} s_{ik}^2 \sigma_k^2}.$$

That is, heterogeneous effects β_{ik} are averaged with weights proportional to: (1) the strength of the first-stage effect π_{ik}, (2) the local share s_{ik} squared, and (3) the shift variance σ_k^2. Evidently, the weights are non-negative under a monotonicity condition $\pi_{ik} \geq 0$ (which holds trivially in shift-share OLS regressions, which correspond to $\pi_{ik} = 1$, $u_{ik} = 0$).

To gain more intuition for the weights, suppose $\pi_{ik} \equiv \pi$ and $\sigma_k^2 \equiv \sigma^2$. Then, if all heterogeneity

[2] This formulation generalizes Proposition 2 of Adão et al. (2019) to IV, rather than reduced-form regressions. Footnote 16 in Adão et al. (2019) considers IV regression but does not allow the effects β_{ik} to vary by k. The analysis in Appendix A.1 of Borusyak et al. (2022) is very general, allowing further for nonlinear effects, but they do not discuss the intuition for the resulting LATE. One limitation of our formulation is that employment growth in industry k is not allowed to be affected by subsidies to other industries; the model in Appendix A.7 of Borusyak et al. (2022) relaxes that assumption but does not study heterogeneous effects.

is by region ($\beta_{ik} = \beta_i$), the weight on β_i is equal to the Herfindahl–Hirschman index of local industry concentration, $\sum_k s_{ik}^2$. A region exposed to many different industry shifts will not be useful for the regression because the law of large numbers eliminates most variation in the shift-share instrument. Conversely, when all heterogeneity is by industry ($\beta_{ik} = \beta_k$), the weight on β_k is equal to $\sum_i s_{ik}^2$.[3] Naturally, this weight is higher for larger industries. More interestingly, it is also higher when the local shares of this industry are very unequal across regions. For instance, tradable industries will play a larger role than nontradable industries of a similar national size, as typical tradable industries concentrate in a small number of regions while nontradable ones are present in every region with relatively homogeneous shares.

The heterogeneous-effect interpretation of shift-share IVs without as-if random shifts is less established. de Chaisemartin and Lei (2023) raise concerns of non-convex weighting of unit-specific causal effects when shift-share IVs are justified by parallel trend assumptions, with respect to either shares (as in the exogenous shares approach) or shifts (i.e., with a weaker restriction on the shift exogeneity), adding to a large literature noting similar issues for popular two-way fixed effect specifications (e.g. de Chaisemartin and D'Haultfoeuille (2020) and Borusyak et al. (2023)). Part of this issue is apparent in the Goldsmith-Pinkham et al. (2020) Rotemberg weight decomposition since, as Goldsmith-Pinkham et al. (2020) note, some weights may be negative. To the best of our knowledge, the case of heterogeneous causal effects across shifts has not—to our knowledge—been studied without as-if random shifts.

A.3 Do the "shares" have to really be shares?

No, they can be any exposure weights.

In most applications s_{ik} is non-negative and typically they are some initial shares; notably this is the case when the shift-share IV follows from the decomposition (3). But econometric results go through when s_{ik} are any weights that measure the exposure of observation i's treatment to the shift g_k.

As an example, consider the Miguel and Kremer (2004) study of spillover effects of deworming. In their OLS specification, the key explanatory variable z_i is the number of student i's neighbors who have received a randomized deworming treatment. Upon inspection, one may notice that this is a shift-share variable: $z_i = \sum_{k=1}^{N} s_{ik} g_k$ where k indexes all students, s_{ik} is a dummy that equals one if students i and k are neighbors, and g_k is a dummy that student k has been selected for deworming. Here the exposure weights are not shares of anything: they take values of zero and one and their total is the number of neighbors student i has. There is no problem with this, as long as the sum of shares (i.e., the number of student's neighbors) is controlled for.

A.4 Should I normalize the shares to add up to one?

You could, but controlling for the sum of shares is probably a better solution.

[3] Note that this is not a Herfindahl–Hirschman index because the shares add up to one across industries, not regions.

From our earlier discussion of how the incomplete shares case requires extra care (specifically, picking appropriate controls in the many exogenous shifts approach), one might conclude that this case is something to be avoided. This can be done by constructing the instrument using shares normalized to add up to one. For instance, while Autor et al. (2013) define s_{ik} as employment shares of manufacturing industry k relative to total employment in labor market i, one could consider redefining the shares to have local manufacturing employment in the denominator.

Such a conclusion would be misguided, however. First consider IV regressions, where the treatment x_i is given by the economic question. Then the researcher needs to choose the best shift-share IV z_i, and in particular the shares, to maximize instrument strength. Whether identification leverages exogenous shifts or exogenous shares, power is maximized when the shares reflect the relationship between the treatment and the shifts, e.g. following the treatment decomposition (3). For example, in the Autor et al. (2013) setting, using the local manufacturing employment in the denominator would reduce power because the shift-share instrument would exhibit large variation even in areas where manufacturing is a low share of total employment and the treatment (import competition) is close to zero. Including appropriate controls is a better way to avoid OVB while retaining statistical power, compared to modifying the shares.

Second, consider OLS shift-share analyses, such as spillover regressions, where the researcher is deciding on the right-hand side variable $x_i = z_i$. This choice is about specifying the most plausible functional form for how the shifts affect the outcome, such that the coefficient is economically meaningful. Again, this is achieved by setting the shares to reflect the exposure of observations to the exogenous shifts. For example, the fraction of treated friends, as in Cai et al. (2015), is a shift-share variable with the shares adding up to one, while the number of treated friends, as in Miguel and Kremer (2004), is an incomplete shares example. Still, if the researcher believes that the outcome is determined by the *number* of treated friends, they should use that specification, and include appropriate controls to avoid bias.

A.5 Can shift-share instruments be constructed by apportioning national changes to units?

Yes, and in fact Bartik (1991), Card (2009), and Autor et al. (2013) all derived their instruments this way. However, to apply the tools from this paper correctly, the resulting instruments must be rewritten with different shares and shifts, as in equation (2).

We illustrate the apportioning logic with the labor supply example. Recall that the percent change in regional employment is an aggregate across industries: $x_i = \left(\sum_k \Delta X_{ik}\right)/X_{i0}$. The researcher can then replace the local industry employment change (in levels), ΔX_{ik}, with a prediction that allocates the national change in the industry employment, ΔX_k, to regions proportionally to the initial regional composition of the industry, $\frac{X_{ik0}}{X_{k0}}$. Region i therefore "gets" $\frac{X_{ik0}}{X_{k0}} \cdot \Delta X_k$ workers in industry k. Adding up such predictions and rescaling them by the initial regional employment yields the instrument:

$$z_i = \frac{\sum_k \frac{X_{ik0}}{X_{k0}} \cdot \Delta X_k}{X_{i0}}. \tag{2}$$

While this expression *looks* different from the shift-share instrument $\sum_k \frac{X_{ik0}}{X_{i0}} \cdot \frac{\Delta X_k}{X_{k0}}$ that follows from the decomposition of x_i in equation (3), a simple rearrangement of terms shows that they are actually the same:

$$\frac{\sum_k \frac{X_{ik0}}{X_{k0}} \Delta X_k}{X_{i0}} = \frac{\sum_k X_{ik0} \cdot \frac{\Delta X_k}{X_{k0}}}{X_{i0}} = \sum_k \frac{X_{ik0}}{X_{i0}} \cdot \frac{\Delta X_k}{X_{k0}}. \tag{3}$$

This rearranging step is crucial for applying the theoretical results and taking the practical steps in both exogenous shifts and exogenous shares approaches. The left-hand side of (3) is based on employment shares relative to the industry total, whereas the shares on the right-hand side are relative to the regional total. The left-hand side suggests that the national shifts are industry employment changes in levels, ΔX_k, although this leaves the denominator unaccounted for by either shares or shifts; on the right-hand side of (3), the shifts are relative changes in national industry employment.

Both conceptual and practical issues arise if the apportioning formula (2) is used without rewriting it as in (3). In the exogenous shifts approach, assuming ΔX_k is as-good-as-randomly assigned is untenable, as larger industries of course get larger employment changes on average (provided national employment is growing).[4] This assumption is also not sufficient because the denominator X_{i0} in (2) is ignored, while it affects the identification conditions. Measuring shifts in relative terms instead makes their as-if random assignment a more plausible assumption. In the exogenous shares approach, using the shares relative to the industry total, X_{ik}/X_k, as instruments is the same as using initial employment levels X_{ik}, since the share denominators in (2) do not vary across observations. Thus, variation in the local industry size is used instead of the local composition of industries that is usually intended in shift-share IV designs. Moreover, since the remaining terms in the summation, $\Delta X_k/X_{i0}$, vary across i, z_i cannot be viewed as pooling variation in the shares (relative to the industry total).

More practically, applying the checklists above to the wrong shares and shifts would lead to incorrect controls (e.g., incomplete share controls) and diagnostic tests (e.g., based on wrong Rotemberg weights). In (2), it looks like there is an incomplete share problem, while (3) makes it clear there is not (since $\sum_k \frac{X_{ik0}}{X_{k0}} \neq 1$ while $\sum_k \frac{X_{ik0}}{X_{i0}} = 1$).

A.6 Can the shifts be unit-specific?

Yes.

While we introduced shift-share variables as combining heterogeneous shares with a common set of shifts, the econometric framework also nests settings where each unit is exposed to a distinct

[4] In Appendix A.11 we argue that the exogenous shifts lens may not be appealing for the Bartik (1991) instrument. However, the issues we discuss here are not specific to that application, and they arise similarly with the Autor et al. (2013) instrument.

set of shifts. One can define k to index the shifts to all observations and redefine the shares such that the exposure of a unit to another unit's shift is zero.

A set of examples is considered by Borusyak and Kolerman-Shemer (2024) who study "regression discontinuity aggregation" designs in which a shift-share treatment aggregates policy discontinuities defined at smaller geographic units. For instance, Clots-Figueras (2011) estimates the effect of the fraction of women in state legislatures in India, using the fraction of women who won against a man in a close election as the IV. Although each state has a distinct set of constituencies, this instrument is a shift-share where each state has non-zero exposure only to its own constituencies' shifts.

A.7 Can I take a shift-share average of shift-share IVs?

Yes, and the result is also a shift-share IV, with the same shifts but more complicated shares.

This situation commonly arises when studying spillovers from treatments (or instruments) that already have a shift-share structure. Adão et al. (2023), for instance, study spatial spillovers from regional import competition with China. Let $z_i = \sum_k s_{ik} g_k$ be the Autor et al. (2013) instrument, capturing direct exposure of commuting zone i to Chinese imports based on industry shifts g_k and local employment shares s_{ik}. Slightly simplifying, Adão et al. (2023) define the indirect exposure of commuting zone j as the inverse-distance weighted average of direct exposures of all other commuting zones: $z_j^* = \sum_i s_{ji}^* z_i$, where the shares s_{ji}^* decay with the distance between j and i (and $s_{jj}^* = 0$). One can see that this variable can be rewritten as $z_j^* = \sum_k s_{jk}^{**} g_k$ with compound shares $s_{jk}^{**} = \sum_i s_{ji}^* s_{ik}$ and original shifts g_k.

Representing the shift-share instrument with the resulting shares and shifts, in one step, makes it clear that exogeneity of g_k is still sufficient for identification. It also yields appropriate incomplete share and other share-aggregated controls, and correct standard errors.

A.8 What if I take logs of a shift-share?

A log — or any other nonlinear transformation — of a shift-share variable is not a shift-share variable. This may or may not complicate IV exogeneity.

In the exogenous shares approach, which views the shift-share IV as a particular function of the shares (where the shifts serve as weights), a nonlinear function of a shift-share IV is just another function of the same shares. If all individual shares are exogenous instruments, i.e. $\mathbb{E}\left[\varepsilon_i \mid s_{i1}, \ldots, s_{iK}\right] = 0$, then any function of them is exogenous, too.

On the contrary, shift exogeneity does not imply exogeneity of nonlinear transformations of the shift-share IV, such as taking the log; such transformations can lead to a new type of bias. To see this, imagine the shares add up to one and the exogenous shifts are assigned in a lottery with positive values. Then, regardless of how the shares are correlated with the error term, the share-weighted average of the lottery shifts $z_i = \sum_k s_{ik} g_k$ is not correlated with the error. That logic fails for $\log z_i$: because of Jensen's inequality, units with dispersed shares will on average have a higher $\log z_i$ than units with concentrated shares, potentially leading to bias. Similar issues arise with other transformations of shift-share IVs, e.g. using a dummy that a shift-share variable is in the lowest quartile of its distribution, as in Greenstone et al. (2020).

There are two ways to avoid this bias. First, Borusyak and Hull (2023) propose a "recentering" adjustment to the nonlinear instrument, such as $\log z_i$, based on rerandomizing the shifts, e.g. by permuting them. Second, putting the log inside the sum, i.e. replacing $\log \sum_k s_{ik} g_k$ with $\sum_k s_{ik} \log g_k$, yields an actual shift-share IV with shares s_{ik} and shifts $\log g_k$.

For a concrete example, Berman et al. (2015) estimate the effects of log firm exports on the log of its domestic sales to measure returns to scale. While our discussion so far has focused on outcomes and treatments measured as changes, consistent with the decomposition (3), Berman et al. (2015) perform the analysis in logs of levels, using a panel of firms and controlling for firm fixed effects. They instrument log exports with $z_{it} = \log \sum_k s_{ik} G_{kt}$, where k denotes product-by-country pairs, s_{ik} is the share of this pair in firm's exports (on average across periods), and G_{kt} is the total world exports of this product to this country. Leveraging exogeneity of G_{kt}, or the log-changes in G_{kt} over time, would require the corrections discussed above.

Borusyak and Hull (2021, footnote 82) show an additional problem with this IV: it implicitly uses shares that are not the s_{ik} and may not capture the intended economic intuition. For instance, one may think that for firm i that has 50% of initial exports in a certain product-country cell k ($s_{ik} = 0.5$), a 10% increase of world exports in that cell raises z_{it} by approximately 0.05. This is not the case. To see the issue, suppose changes in G_{kt} over time are sufficiently small and consider how $z_{it} = \log \sum_k s_{ik} G_{kt}$ changes in response, relative to some base period 0 (recalling that, with firm fixed effects, changes over time play the key role). It is easy to show that

$$z_{it} - z_{i0} \approx \sum_k \frac{s_{ik} G_{k0}}{\sum_{k'} s_{ik'} G_{k'0}} \left(\log G_{kt} - \log G_{k0}\right) \neq \sum_k s_{ik} \left(\log G_{kt} - \log G_{k0}\right). \tag{4}$$

Thus, the response of z_{it} to a 10% shift to G_{kt} is determined not only by the share of k in firm i's initial exports but also by the world supply of k in the initial period — which was presumably not intended when constructing the instrument. To avoid this issue, one can replace $\log \sum_k s_{ik} G_{it}$ with $\sum_k s_{ik} \log G_{kt}$.

A.9 What if I have multiple shift-share instruments?

This is fine, both when multiple shift-share variables instrument for a single treatment and when multiple IVs are necessitated by multiple treatments. One should just perform the relevant steps for each of the shift-share IVs: e.g., include incomplete share controls in the exogenous shifts approach and check sensitivity to how shares are combined in the exogenous shares approach.

Getting exposure-robust standard errors may be more challenging in this case. When the shares are the same but there are several sets of exogenous shifts, Borusyak et al. (2022) show how the shift-level equivalent IV regression extends in this case, yielding correct standard errors. Appendix B.1 below extends this result by allowing for several shift-share IVs that use different shares and different shifts, as long as all shifts are defined at the same "level" k. We derive an equivalent shift-level representation of the estimator in terms of a set of moment conditions (but no longer as a simple IV). This equivalence result yields exposure-robust standard errors. A Stata example is available in our GitHub repository, https://github.com/borusyak/shift_share_jep.

We give two examples. First, Dauth et al. (2014) consider the impacts of two import competition shifts in Germany, originating from the growth of China and from the accession of Eastern European countries into the European Union. Both are shift-share variables that combine the local employment shares of different industries with two national industry import competition shifts.

Second, including both direct and spillover effects of a certain treatment in the same specification can be viewed as using two shift-share variables with the same shifts but different shares. For instance, the right-hand side variables in Miguel and Kremer (2004) are the student i's own deworming dummy and the number of her dewormed friends. We explained above how their spillover treatment is a shift-share IV that uses deworming dummies as the shifts g_k and the patterns of friendship as exposure weights. Mechanically, one's own deworming status is also a shift-share with the exposure weight being one for $i = k$ and zero otherwise.

A.10 What if I have interaction terms in a shift-share regression?

This is similar to having multiple shift-share variables.

There can be two types of interaction terms in shift-share regressions. A more conventional one interacts z_i with some unit-level variable a_i. For instance, in the Autor et al. (2013) context, one may be interested in understanding whether labor market responses to import competition vary by the share of college graduates in the region. This interaction can be written as a shift-share IV with the same shifts and different exposure weights: $a_i z_i = \sum_k (a_i s_{ik}) g_k$.

The second type — albeit not exactly an interaction — aims to identify the heterogeneous responses to different groups of *shifts*. For instance, Bombardini and Li (2020) consider the health effects of two treatments: regional exposure to the national industry growth of exports for all industries and for pollution-intensive industries in particular. The former is a standard shift-share variable $z_i = \sum_k s_{ik} g_k$ while the latter can be written as $z_i' = \sum_k s_{ik}(b_k g_k)$ where b_k is industry's pollution intensity.[5] This z_i' is a shift-share IV with shares s_{ik} and shifts $b_k g_k$.[6]We refer the reader to Appendices A.9 and B.1 for a discussion of incomplete share and other appropriate controls, as well as exposure-robust standard errors with multiple shift-share instruments.

[5]Note that z_i' is not the same as the interaction of z_i with the regional share of pollution-intensive industries, which would be an interaction term of the first type.

[6]It can also be viewed as a shift-share with shares $s_{ik} b_k$ and shifts g_k. Both interpretations lead to the same practical conclusions, in different ways. For instance, with as-good-as-random g_k, one needs to control for $\sum_k s_{ik} b_k$. In the former interpretation this follows because the shifts $g_k b_k$ can be considered as-good-as-random only controlling for b_k (while the shares add up to one). In the latter interpretation this follows because the shares $s_{ik} b_k$ add up to $\sum_k s_{ik} b_k$ (while the shifts are already as-good-as-random).

A.11 Can the instruments in Bartik (1991) and Card (2009) be valid without exogenous shares?

This depends on the underlying model, but probably not. We consider shift-share instruments with shifts constructed as national averages of endogenous local shifts correlated with the error terms, as in Bartik (1991) and Card (2009). The researcher might argue that these shifts proxy for some latent exogenous shifts. Here we show that the proxy error in the shifts is innocuous when the local *shares* are exogenous (and if there are many more observations than shifts), while otherwise the proxy error typically makes the instrument invalid. Thus, there is little value in focusing on the shifts for justifying the validity of Bartik (1991) and Card (2009) type instruments, except in a special case discussed below.

For concreteness, we illustrate the general insight in the setting of Bartik (1991); we discuss Card (2009) at the end. Bartik (1991) estimated the (inverse) elasticity of local labor supply by using a shift-share instrument that combined local employment shares of different industries with the national growth rate of employment in each industry (see Table 1). For the Bartik (1991) instrument to be valid, it has to capture labor demand conditions. Interestingly, Blanchard and Katz (1992) focus on the exogeneity of the shifts, rather than local employment shares, when introducing the Bartik (1991) instrument: *"This series will be valid for our purposes [of isolating a labor demand shift] as long as the national growth rates are not correlated with labor supply shifts in the state"* (p. 25). Is the exogenous shifts approach appropriate in this setting? In particular, is it a problem that the shifts are equilibrium outcomes which may also be affected by labor supply factors?

We give intuition before the formal analysis. Suppose high net migration—internal or foreign—into a region makes employment in all local industries grow. Then, industries that are concentrated in regions with growing net migration will systematically have higher employment growth in most regions, and therefore nationally. That, however, is precisely the situation when the industry growth rate shifts are econometrically endogenous. The main (although not the only) case when this does not happen is if no industry is concentrated in regions with growing or falling net migration. But that corresponds to the case where the local shares of all industries are exogenous with respect to the local net migration rate. It is further required that industries are not too concentrated in a small number of regions, such that random local migration shocks do not have a big impact on national industry growth rates.

We now formally characterize how labor supply shocks affect national industry growth rates. We model employment growth by region and industry as

$$x_{ik} = g_k^* + u_{ik},$$

where g_k^* is the latent national industry labor demand shock and u_{ik} captures labor supply factors.[7] Since labor demand conditions are unobserved, Bartik (1991) proxy for them by the national industry employment growth rate g_k when constructing the shift-share instrument. Denoting by E_{ik},

[7] The results extend directly to the case where labor demand shifts vary across regions.

E_i, and E_k the initial employment levels by region-industry, region, and industry respectively, we have:

$$g_k = \sum_i \frac{E_{ik}}{E_k} x_{ik} = g_k^* + \tilde{g}_k,$$

where the proxy error is given by

$$\tilde{g}_k = \frac{\sum_i E_{ik} u_{ik}}{E_k}. \tag{5}$$

We consider a favorable case where the labor demand conditions g_k^* are exogenous, and thus the only concern is whether $\tilde{g}_k$ affects the exogeneity of the instrument.

As discussed in the main text, the necessary and sufficient condition for the shift-share instrument to be valid is that the measured shifts g_k have no covariance with a particular industry confounder. Specifically, provided the regional analysis is performed with initial employment E_i as importance weights, as is commonly done, this confounder is the average of regional error terms ε_i weighted by initial employment in industry k:[8]

$$\bar{\varepsilon}_k = \frac{\sum_i E_{ik} \varepsilon_i}{E_k}. \tag{6}$$

The expressions for the proxy noise (5) and the confounder (6) exhibit a striking similarity: if some labor supply conditions in ε_i affect employment local in all industries (u_{ik}), we may expect employment-weighted averages of those shocks to be correlated, too. Indeed, in the simple model of local labor markets in Appendix A.7 of Borusyak et al. (2022), regional labor supply shocks affect industry employment growth rates equally, such that $u_{ik} = \gamma \varepsilon_i$ for some $\gamma > 0$. In that model, $\tilde{g}_k$ and $\bar{\varepsilon}_k$ would be perfectly correlated.

There are, however, some special cases in which the problem does not arise, both linked to the properties of the local employment shares. First, if the shares of all industries are exogenous with respect to the regional error term ε_i and the number of industries is small, $\bar{\varepsilon}_k \xrightarrow{p} 0$ for each k. In this case, the exogeneity of the shifts is not required so any proxy noise is fine (Goldsmith-Pinkham et al., 2020).[9]

Second, if for each industry the share s_{ik} is exogenous with respect to the local employment change in that industry due to labor supply, u_{ik}, the proxy noise will average out: $\tilde{g}_k \xrightarrow{p} 0$. This is the case, in particular, for labor supply shocks that induce reallocation of workers across domestic regions in the sample without changing industry. Naturally, such reallocation does not affect the national industry employment growth. However, it seems unlikely that this scenario constitutes the only source of local shift endogeneity, particularly since industry switching (or, similarly, international migration or mobility out of unemployment or non-employment) is necessary to generate nontrivial national industry growth rates to begin with. In that case, the Bartik (1991) instrument

[8] The weights in this averaging combine the shares underlying the instrument, $s_{ik} = E_{ik}/E_i$, and importance weights E_i.

[9] One can see is that $\bar{\varepsilon}_k$ is the weighted covariance between s_{ik} and ε_i weighted by E_i and rescaled by E_k; see Appendix A.2 in Borusyak et al. (2022) for further details on how $\bar{\varepsilon}_k \xrightarrow{p} 0$ constitutes the relevant notion of share exogeneity.

cannot be valid without shares being exogenous (with respect to ε_i) too.

Analogous issues are likely present in Card (2009), who constructs the shifts as the national growth rate of migration from origin country k which are aggregates of local migration rates from that country. Here problems arise if the local migration rates by origin are endogenous, i.e. correlated with the error term—which for Card (2009) reflects the relative demand for migrant vs. native labor. As long as the local migration rates from all origins respond to the same local relative demand conditions, the resulting national shifts cannot be viewed as exogenous without the local shares of migration from different origins being exogenous, too.

We close by noting that our discussion here concerned problems with shifts constructed as national averages of endogenous local shifts, but in practice researchers often use leave-one-out averages.We discuss the role of leave-one-out adjustments in the next section.

A.12 What is the role of leave-one-out construction of shifts?

This is a useful way of mitigating bias when pooling variation from many exogenous share instruments. It is an open question whether this practice can help to extract exogenous latent shifts when the shares are endogenous.

In settings like Bartik (1991) where, as explained above, the shifts can be mechanically confounded by the errors, it is common since Autor and Duggan (2003) to use "leave-out" constructions of shift-share instruments: $z_i = \sum_k s_{ik} g_{k,-i}$, where $g_{k,-i}$ is, say, the industry growth rate in all regions except i (or perhaps except nearby regions, too).[10] With many exogenous *shares*, Borusyak et al. (2022, Appendix A.6) show that using leave-one-out means to construct the national growth rates is useful to address the finite sample bias that can mechanically arise when using own-observation information. This approach is similar to how jackknife instrument variable estimators avoid bias of 2SLS in presence of many instruments (Angrist et al., 1999).

In practice, Autor and Duggan (2003) observed that including own region in shift construction made the IV substantially stronger, raising concerns about the mechanical relationship. Other authors (e.g., Goldsmith-Pinkham et al. (2020)) found that the leave-out correction is empirically minor when the measured shifts average over sufficiently many observations.

It is an open question whether leave-one-out constructions can help address the problem of proxy bias in the shifts discussed in Appendix A.11 when the shares are endogenous. On the one hand, the leave-one-out construction can be viewed as similar to jackknife instrumental variable estimation which Kolesar et al. (2015) show can be consistent under a particular orthogonality condition even when there are many invalid instruments. On the other hand, the error terms of observations with shares similar to i can be correlated with ε_i, in which case leaving out i may not suffice. In a Monte Carlo simulation available by request, we confirm that leave-one-out need not fully eliminate the bias when the shares are endogenous.

We finally note that the leave-out constructions of shift-shares are distinct from a practice of

[10] Strictly speaking, such z_i is not a shift-share as defined by equation (2), since $g_{k,-i}$ has some variation across units.

measuring the shifts from entirely different data, e.g. in different countries. Autor et al. (2013), for instance, instrument regional exposure to Chinese imports in the US using industry shifts measured in other developed countries; Hummels et al. (2014) and Aghion et al. (2022) use similar approaches when instrumenting firm-level imports. Unlike leave-out constructions, here the shifts g_k are the same for all units in the sample. Moreover, the mechanical correlation between the error term and the shifts does not arise, such that the exogenous shifts approach can be applied under the appropriate assumptions (e.g., that import demand shifts in the US and other developed economies are uncorrelated in the Autor et al. (2013) context).[11]

B Theoretical Results

In this appendix, we present three new theoretical results.

B.1 Exposure-Robust Standard Errors for Shift-Share IV Regressions with Multiple Treatments

We derive exposure-robust standard errors for a IV (or OLS) regression with multiple shift-share instruments, by recasting the estimator as a method of moments estimator at the level of shifts. This result builds on Borusyak et al. (2022), Proposition 5.

Consider a just-identified shift-share IV regression:

$$y_i = \beta_1 x_{1i} + \cdots + \beta_R x_{Ri} + \gamma' w_i + \varepsilon_i \tag{7}$$

where $x_{1i}, \ldots, x_{Ri}$ are instrumented with a set of shift-shares $z_{1i}, \ldots, z_{Ri}$ for $z_{ri} = \sum_{k=1}^{K} s_{rik} g_{rk}$. Both the shares and the shifts can differ across r but we require the shifts to vary at the same level for all r (and thus with the same number of shifts K). Assuming the shifts g_{rk} are as-good-as-randomly assigned after controlling for some vector of shift-level controls q_{rk} (which can vary across r), we require the vector of controls w_i to include $\sum_k s_{rik} q_{rk}$ for each r. The vector w_i further includes the intercept and possibly other controls.

The IV estimator $\hat{\beta}$ for $\beta = (\beta_1, \ldots, \beta_R)'$ in (7) satisfies a system of R equations:

$$\frac{1}{N} \sum_i \left(y_i^{\perp} - \sum_{j=1}^{R} \hat{\beta}_j x_{ji}^{\perp} \right) z_{ri} = 0, \qquad r = 1, \ldots, R,$$

where for any variable v_i we let $v_i^{\perp}$ denote the in-sample projection of v_i on w_i. Expanding the expression for z_{ri}, exchanging the order of summation, denoting $\tilde{v}_k^{(r)} = \frac{1}{N} \sum_i s_{rik} v_i^{\perp}$, and combining

[11] In settings like Hummels et al. (2014), the researcher may therefore entertain two options: to measure the shifts in a different country and follow the exogenous shifts approach, or to measure the shifts in the country of interest in a leave-out way and follow the exogenous shares approach.

terms yields a set of R shift-level moment conditions satisfied by $\hat{\beta}$:

$$\sum_k \left(\tilde{y}_k^{(r)} - \sum_{j=1}^{R} \hat{\beta}_j \tilde{x}_{jk}^{(r)}\right) g_{rk} = 0, \qquad r = 1, \ldots, R.$$

Letting $\tilde{g}_{rk}$ be the projection of g_{rk} on q_{rk} weighted by $s_{rk} = \frac{1}{N}\sum_i s_{rik}$ and noting that $\sum_k \tilde{v}_k^{(r)} q_{rk} = 0$ since w_i includes $\sum_k s_{rik} q_{rk}$, we further have a set of R equations on the residualized shifts:

$$\sum_k \tilde{\psi}_k^{(r)} = 0, \quad \text{for } \tilde{\psi}_k^{(r)} = \left(\tilde{y}_k^{(r)} - \sum_{j=1}^{R} \hat{\beta}_j \tilde{x}_{jk}^{(r)}\right) \tilde{g}_{rk}.$$

In matrix form, this can be rearranged as

$$\Omega \hat{\beta} = M,$$

where $\Omega_{rj} = \sum_k \tilde{x}_{jk}^{(r)} \tilde{g}_{rk}$ and $M_r = \sum_k y_k^{(r)} \tilde{g}_{rk}$. Thus, $\hat{\beta} = \Omega^{-1} M$. Moreover, since (7) implies $\tilde{y}_k^{(r)} - \sum_{j=1}^{R} \beta_j \tilde{x}_{jk}^{(r)} = \tilde{\varepsilon}_k^{(r)}$ for true β, we also have

$$\hat{\beta} - \beta = \Omega^{-1} E \qquad \text{for } E_r = \sum\nolimits_k \varepsilon_k^{(r)} \tilde{g}_{rk}.$$

We assume that the appropriate relevance condition holds and suppose that vectors of shift residuals $\tilde{g}_k = (\tilde{g}_{1k}, \ldots, \tilde{g}_{Rk})'$ are asymptotically independent across some shift clusters c. Letting $\tilde{\psi}_k = \left(\tilde{\psi}_k^{(1)}, \ldots \tilde{\psi}_k^{(R)}\right)'$, we then have an asymptotic approximation of the exposure-robust variance-covariance matrix of $\hat{\beta}$:

$$\begin{aligned} \mathrm{Var}\left[\hat{\beta} - \beta\right] &\approx \Omega^{-1} \mathrm{Var}\left[E\right] (\Omega^{-1})' \\ &\approx \Omega^{-1} \left(\sum_c \left(\sum_{k \in c} \tilde{\psi}_k\right) \left(\sum_{k \in c} \tilde{\psi}_k\right)'\right) (\Omega^{-1})'. \end{aligned} \tag{8}$$

We note that the derivation here simplifies if the shares are the same for all shift-share instruments (and q_{rk} are also the same) and only the shifts vary across r. In this case, the coefficients and exposure-robust standard errors can be obtained by an IV estimator at the shift-level, as shown by Borusyak et al. (2022). This includes instruments constructed as $\sum_k s_{ik} b_k g_k$ for different "interaction" variables b_k, as discussed in Appendix A.10.

B.2 Visual IV weights with Share Exogeneity

In this appendix, we show the shift-share IV estimate equals the slope of the regression line through the points on the visual IV graph for the exogenous shares approach, without intercept and with appropriate weights.

Let $\hat{\beta}_k$ be the share-IV estimate for industry k and let ω_k denote the Rotemberg weights, which

Figure 1: Visual IV for Exogenous Shares, Applied to Card (2009)

A: High-school graduates B: College graduates

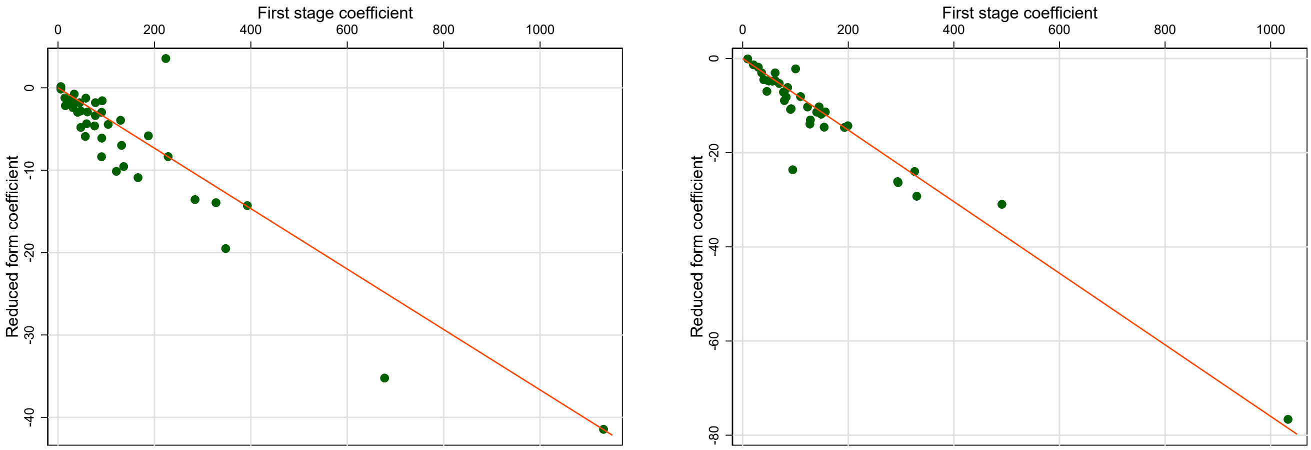

Notes: The visual IV graph in the setting of Card (2009) using the replication data from Goldsmith-Pinkham et al. (2020), estimating the relationship between the log wage gap between immigrant and native workers (as the outcome) and the ratio of immigrant to native hours worked (as the treatment). Card (2009) instruments the local ratio of immigrant to native hours with a shift-share instrument, leveraging immigration patterns from 38 countries. We plot the reduced-form coefficient against the first-stage coefficient for each share IV, using the immigration shares from each of the 38 countries one at a time as instruments. Panel A focuses on high-school graduates while Panel B considers college graduates. The shift-share IV estimate is visualized as the slope of the ray through the origin.

sum to one and are such that the shift-share IV coefficient is $\hat{\beta} = \sum_k \omega_k \hat{\beta}_k$.[12] Write $\hat{\beta}_k = \hat{\rho}_k / \hat{\pi}_k$, where $\hat{\rho}_k$ and $\hat{\pi}_k$ are reduced-form and first-stage estimates for the kth share-IV. Then we have:

$$\hat{\beta} = \sum_k \omega_k \frac{\hat{\rho}_k}{\hat{\pi}_k} = \frac{\sum_k (\omega_k / \hat{\pi}_k^2) \hat{\rho}_k \hat{\pi}_k}{\sum_k (\omega_k / \hat{\pi}_k^2) \hat{\pi}_k^2},$$

which is the slope from a regression of $\hat{\rho}_k$ on $\hat{\pi}_k$, with no intercept and with weights $\omega_k / \hat{\pi}_k^2$ (which are not necessarily convex since Rotemberg weights can take negative values).

References

Adão, Rodrigo, Costas Arkolakis, and Federico Esposito, "General Equilibrium Effects in Space: Theory and Measurement," *Working Paper*, 2023.

__ **, Michal Kolesar, and Eduardo Morales**, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel, "Modern Manufacturing Capital, Labor Demand, and Product Market Dynamics: Evidence from France," *Working Paper*, 2022.

[12] See Proposition 3 in Goldsmith-Pinkham et al. (2020) for the definition of Rotemberg weights and Section IV.B for the adjustments needed when the shares add up to one.

Angrist, J D, Guido W. Imbens, and AB Krueger, "Jackknife instrumental variables estimation," *Journal of Applied Econometrics*, 1999, *67* (95), 57–67.

Autor, David H. and Mark G. Duggan, "The rise in the disability rolls and the decline in unemployment," *Quarterly Journal of Economics*, 2003, *118* (1), 157–205.

_ , David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Impacts of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Bartik, Timothy J., *Who Benefits from State and Local Economic Development Policies?*, W. E. Upjohn Institute for Employment Research, 1991.

Berman, Nicolas, Antoine Berthou, and Jerome Hericourt, "Export Dynamics and Sales at Home," *Journal of International Economics*, 2015, *96* (2), 298–310. Publisher: Elsevier B.V.

Blanchard, Olivier Jean and Lawrence F. Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, (1), 1–75.

Bombardini, Matilde and Bingjing Li, "Trade , Pollution and Mortality in China," *Journal of International Economics*, 2020.

Borusyak, Kirill and Matan Kolerman-Shemer, "Regression Discontinuity Aggregation, with an Application to the Union Effects on Inequality," *Mimeo*, 2024.

_ and Peter Hull, "Non-Random Exposure to Exogenous Shocks: Theory and Applications," *Mimeo*, 2021.

_ and _ , "Non-Random Exposure to Exogenous Shocks," *Econometrica*, 2023, *91* (6), 2155–2185.

_ , _ , and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

_ , Xavier Jaravel, and Jann Spiess, "Revisiting Event Study Designs: Robust and Efficienct Estimation," *Working Paper*, 2023.

Cai, Jing, Alain De Janvry, and Elisabeth Sadoulet, "Social Networks and the Decision to Insure," *American Economic Journal: Applied Economics*, April 2015, *7* (2), 81–108.

Card, David, "Immigration and Inequality," *American Economic Review: Papers & Proceedings*, 2009, *99* (2), 1–21.

Christian, Paul and Christopher B Barrett, "Spurious Regressions and Panel IV Estimation: Revisiting the Causes of Conflict," *The Economic Journal*, 2024, *134* (659), 1069–1099.

Clots-Figueras, Irma, "Women in politics," *Journal of Public Economics*, August 2011, *95* (7-8), 664–690.

Dauth, Wolfgang, Sebastian Findeisen, and Jens Suedekum, "The rise of the east and the far east: German labor markets and trade integration," *Journal of the European Economic Association*, 2014, *12* (6), 1643–1675. ISBN: 9783863041267.

de Chaisemartin, Clement and Xavier D'Haultfoeuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996. arXiv: 1803.08807v1.

_ **and Ziteng Lei**, "More Robust Estimators for Instrumental-Variable Panel Designs, With An Application to the Effect of Imports from China on US Employment," *Working Paper*, 2023.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen, "Do Credit Market Shocks Affect the Real Economy? Quasi-Experimental Evidence from the Great Recession and "Normal" Economic Times," *American Economic Journal: Economic Policy*, 2020, *12* (1), 200–225.

Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence From Danish Matched Worker-Firm Data," *American Economic Review*, 2014, *104* (6), 1597–1629. arXiv: 1011.1669v3 ISBN: 9788578110796.

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, *62* (2), 467–475. ISBN: 00129682.

Jaeger, David A, Joakim Ruist, and Jan Stuhler, "Shift-Share Instruments and the Impact of Immigration," *Working Paper*, 2017.

Kolesar, Michal, Raj Chetty, John Friedman, Edward Glaeser, and Guido W. Imbens, "Identification and Inference With Many Invalid Instruments," *Journal of Business and Economic Statistics*, 2015, *33* (4), 474–484. ISBN: 0735-0015.

Miguel, Edward and Michael Kremer, "Worms: Identifying impacts on education and health in the presence of treatment externalities," *Econometrica*, 2004, *72* (1), 159–217.

Nunn, Nathan and Nancy Qian, "US Food Aid and Civil Conflict," *American Economic Review*, 2014, *104* (6), 1630–1666.