Synthetic instruments in DiD designs with unmeasured confounding

Ahmet Gulek (MIT) and Jaume Vives-i-Bastida (MIT)

[Preliminary Draft]

Abstract

Unmeasured confounding and selection into treatment are key threats to reliable causal inference in Difference-in-Differences (DiD) designs. In practice, researchers often use instrumental variables to address endogeneity concerns, for example through shift-share instruments. However, in many settings instruments may be correlated with unobserved confounders, exhibiting pre-trends. In this paper we explore the use of synthetic controls to address unmeasured confounding in IV-DiD settings. We propose a synthetic IV estimator that partials out the unmeasured confounding and derive conditions under which it is consistent, when the standard two-stage least squares is not, and asymptotically normal. Motivated by the finite sample properties of our estimator we then propose an ensemble estimator that might address different sources of bias simultaneously. Finally, we show the relevance and pitfalls of our estimator in a simulation exercise and in an empirical application to the Syrian refugee crisis.

1. Introduction

In many applied settings researchers are interested in evaluating the impact of a treatment or an intervention that affects a set of units of interest over time. In these settings, researchers may be worried about endogeneity concerns such as the treatment affecting units selectively or differential trends amongst units that received different doses of the treatment. To address these concerns researchers may turn to differences in differences (DiD) designs (Card and Krueger, 2000) or synthetic control (Abadie and Gardeazabal, 2003) designs (SC) in which control units are used to evaluate the counterfactual in absence of the intervention. While these approaches may address part of the endogeneity problems, often valid control units may not exist, as all units may be treated, or control units and treated units may not follow similar paths (violating the parallel trends assumption). Faced with this reality, researchers may consider an instrumental variable (IV) approach in combination with the Diff-in-Diff design (for example using a shift-share instrument, e.g. Jaeger et al. (2018a)). Unfortunately, in practice the endogeneity problems may persist as the instrument may be correlated with unobserved heterogeneity in the outcome of interest, exhibiting 'pre-trends'. In this paper we propose a Synthetic IV estimator that combines the instrument and synthetic controls to account for the bias due to unobserved heterogeneity.

To clarify the setting we are interested in, consider an empirical example: the refugee crisis due to the Syrian civil war. Over the 2011-2017 period approximately 3.5 million Syrians entered Turkey as refugees. Using this shock we may be interested in studying the impact of new immigrants (refugees) on local labor market outcomes. However, comparing regions that received immigrants with regions that do not may not provide a causal estimate of the impact of immigration as immigrants choose *where* to locate endogenously. To account for this source of endogeneity we may want to use a shift-share instrumental variable (based on the distance from the Syrian border for example). The problem with this approach is that, as we show in Section 2 of the paper, the shift-share IV may be biased (even when including two way fixed effects) as regions close and further away from the Syrian border are on different economic trajectories. Our proposed method, the synthetic IV, creates a synthetic control unit for each region in the pre-intervention period and then debiases the outcomes of interest to account for the differential trends and corrects the bias in the two-stage least square estimator.

To motivate our method theoretically we derive consistency and asymptotic normality results in triangular panel designs with unmeasured confounding. We assume that the unobserved error term has two components: an idiosyncratic component that is orthogonal to the instrument and an unobserved heterogeneity component that follows a factor structure. If we could control for the unobserved factor structure the TSLS would be consistent, but we cannot do so directly. Our solution, the synthetic IV, proposes synthetic controls as a way to proxy for the unobserved confounding through interpolation. Under signal-to-noise restrictions and weak primitives assumptions we show that the synthetic IV is consistent and asymptotically normal when the number of units and time periods is large. Through finite sample bounds we highlight that the proposed estimator might be specially sensitive to the noise level and the weakness of the instrument. To account for this we propose empirical checks researchers might want to implement in practice, as well as a 'doubly robust' ensemble estimator that combines the synthetic IV with a projected synthetic IV that partials out the noise.

We show the applicability of our method in a simulation study and by re-visiting the Syrian refugee crisis example. The simulation study shows that the synthetic IV and ensemble estimators outperform the TSLS with two-way fixed effects in a variety of settings. Furthermore, the synthetic IV exhibits close to zero bias in settings with moderate and small levels of noise and unobserved heterogeneity, and the ensemble estimator is shown to be robust in settings with higher noise levels. In a study of the coverage of the synthetic IV estimator we find that it is good in cases in which the estimator exhibits small bias. Following the theoretical properties and the observed behavior under simulations we recommend that researchers implement three checks (in the spirit of the best practices detailed in Abadie and Vives-i-Bastida (2022)) when using the estimator: (1) ensuring that the instrument is not weak after the debiasing, (2) making sure that the estimator achieves good fit in the pre-treatment period and (3) implementing a back test to ensure the good fit is not due to over-fitting to the idiosyncratic noise. In re-visiting the Syrian refugee crisis we find that the while the TSLS with two-way fixed effects exhibits pre-trends, the synthetic IV estimator does not. More so, using the synthetic IV estimator yields different results when evaluating some of the main outcomes of interest: while the TSLS estimator cannot reject that there is no effect of immigration on natives' salaried employment, the synthetic IV estimator finds a statistically significant negative effect. For example, using TSLS we find that a 1 pp increase in refugee/native ratio is associated with a 0.01 pp *increase* in native salaried employment for low-skilled men, whereas using SIV shows that it causes a 0.16 pp *decrease*. This implies that for every 100 immigrants that arrived to Turkey, 16 low-skilled natives lost salaried jobs. These economically and statistically significant differences in the IV and SIV estimates highlight the role of unobserved confounders in the long-standing debate about the labor market effects of immigrants (Borjas, 2017; Peri and Yasenov, 2019).

This paper contributes to several strands of the literature. First, it complements the growing body of work on addressing unobserved confounding and 'pre-trends' in panel data settings by providing a method for the IV Diff-in-Diff case. Research in this area is built upon synthetic control based methods (Abadie et al., 2010, 2015; Ben-Michael et al., 2021), more general weighting methods such as the synthetic differences in differences (SDID , Arkhangelsky et al. (2021a)), as well as balancing methods (Hainmuller 2012), matrix completion methods (Agarwal et al., 2021; Athey et al., 2021) and factor model methods (Anatolyev and Mikusheva, 2022; Bai, 2009). Similarly, our paper complements related work on addressing and evaluating pre-trends in event-study designs, including Freyaldenhoven et al. (2019), Borusyak et al. (2023), Roth (2022) and Ham and Miratrix (2022) among others. A more closely related paper is Arkhangelsky and Korovkin (2023) which provide a novel weighting algorithm, similar to SDID, to address unobserved confounding in settings in which the exogenous variation comes from aggregate time series shocks. The authors propose a robust estimator that corrects the TSLS bias when the instrument has a product form and there are unobserved aggregate shocks that may affect different units differently. We see our method as complementary to Arkhangelsky and Korovkin (2023), and note that we consider a different setting in which the instrument need not have a product structure and the exogenous variation may come from the time or unit components.

Second, this paper is related to a growing literature studying and relaxing the identification assumptions embedded in shift-share designs. Goldsmith-Pinkham et al. (2020) show that the identification assumpions in SSIV designs are often based on the exogeneity of shares. Borusyak et al. (2022) relax this assumption and provide a framework in which identification can also come from the exogeneity of shifts, allowing shares to be endogenous. Adao et al. (2019) highlight an inference problem that arises from cross-regional correlation in the regression residuals due to similarity of sectoral shares in the US. In the immigration context, Jaeger et al. (2018b) show that past-settlement instruments in practice conflate both short-term and long-term adjustments to immigration shocks, which invalidates the exogeneity assumption. Our method provides an additional tool applied researchers can rely on to address unobserved confounders in the SSIV designs.

Lastly, our empirical example is related to a large literature studying the effects of immigration using refugee shocks (Card, 1990; Hunt, 1992; Friedberg, 2001; Angrist and Kugler, 2003; Lebow, 2022). More specifically, our focus on the effects of Syrian refugees on Turkish natives and the presence of unobserved confounders in Turkey follows Gulek (2023) closely. Whereas he focuses on the effects on the formal and informal labor markets, we focus on the overall impact on salaried employment and consider heterogeneity across men and women.

The paper proceeds as follows. Section 2 describes the setting and empirical example. Section 3 presents the synthetic IV estimator and two additional estimators. Section 4 discusses the theoretical results. Section 5 regards the simulation study. Section 6 revisits the empirical application and section 8 concludes.

2. Setting and the Syrian refugee shock

We are interested in a panel data setting in which some units of interest are exposed to a (potentially continuous) treatment and there are endogeneity concerns. The researcher may be worried about using a differences-in-differences design as the parallel trends assumption might not hold, but has access to an instrument that *partially* addresses the endogeneity concerns. More precisely, we consider J units indexed by $i = 1, \ldots, J$ that are observed for T periods of time with outcomes of interest Y_{it} and potential outcomes denoted by $Y_{it}(R_{it})$ for a random variable $R_{it} \in \mathbb{R}$. Throughout the paper we assume that the potential outcomes follow a linear triangular system as described by Assumption 1.

Assumption 1 *[Design] Potential outcomes follow a linear factor model*

$$Y_{it}(R_{it}) = \theta R_{it} + \mu_i' F_t + \epsilon_{it}$$

where μ_i is a vector of k unobserved factor loadings, F_t denotes the common factors and ϵ_{it} are unobserved error terms.
A1.1*: The treatment R_{it} follows*

$$R_{it} = \gamma Z_{it} + \eta_{it},$$

where λ_i are unobserved factor loadings, Z_{it} is the instrumental variable and η_{it} is the unobserved error term.
A1.2*: (Shock design) The treatment R_{it} instead follows*

$$R_{it} = \mathbf{1}\{t > T_0\}(\gamma Z_{it} + \eta_{it}).$$

Assumption 1 defines the triangular design we are interested in. A1.2 states that the intervention of interest affecting the units happens at time T_0 with the outcome of interest (R_{it}) being zero for $t < T_0$, as in shift-share designs with time series shocks. We make this simplification because we are interested in event study designs in our empirical application in which there are no refugees before T_0, however our theory holds more generally and one could equivalently re-write the design model as in A1.1. In the case in which there is no times series shock at T_0 one can think of choosing T_0 as a sample splitting procedure. More importantly, Assumption 1 imposes structure on the unobserved terms by assuming they follow a linear factor model. Linear factor model assumptions are common in the synthetic control literature (see Abadie et al. (2010), or Abadie and Vives-i-Bastida (2022) for a simulation review) and in the matrix completion literature (see Athey et al. (2021) for a review).

This functional form assumption allows us to separate the unobserved term into an omitted variable (the factor structure) and an unobserved error term. With this in mind, the parameter of interest is the expected marginal effect of R_{it}

$$\theta = \mathbb{E}\left[\frac{\partial Y_{it}(R_{it})}{\partial R_{it}}\right] = \frac{\partial Y_{it}(R_{it})}{\partial R_{it}}$$

for which the OLS estimator is unbiased under the independence assumption $R_{it} \perp \epsilon_{it}, \mu_i' F_t$. In many settings, however, R_{it} is likely correlated with the unobserved factor term $\mu_i' F_t$ or the unobserved error term ϵ_{it}. For example, in immigration settings refugees might take into account local labor market conditions when choosing where to re-locate or alternatively may relocate based on geographical distance. We distinguish between the two types of correlation by assuming we have an instrument Z_{it} that is valid for the unobserved endogeneity concern due to ϵ_{it}, but not for the omitted factor structure $\mu_i' F_t$.

Assumption 2 *[Partial instrument exogeneity] The following independence conditions hold* $\epsilon_{it}, \eta_{it} \perp Z_{it}$.

Assumption 2 states that the instrument is able to address the unobserved error part of the unmeasured confounding that is due to the correlation between η and ϵ. For example, common instruments in the immigration literature to address location choice endogeneity include past-settlement indicators or travel distance (Card, 2001; Angrist and Kugler, 2003).

This set up can be expressed in a simple DAG as shown in panel (a) of Figure 1. Assumption 1 and 2 are encoded in the DAG, and in principle there could also be an arrow between $U_{it} = \mu_i' F_t$ and R_{it}. Panel (c) shows a case in which the OLS would be valid and panel (b) shows the case in which instrumental variable approach would be valid.

Figure 1: Triangular designs

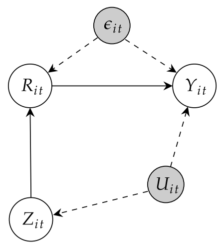

(a) General design

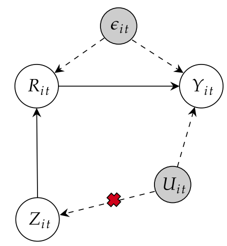

(b) IV design

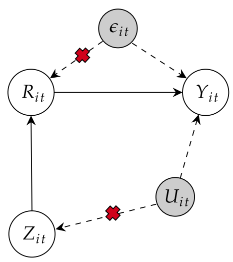

(c) OLS design

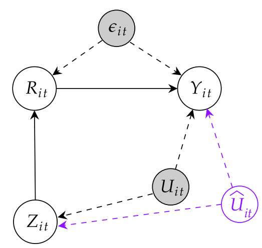

(d) Proxy design

We are interested in the general case in which neither OLS nor the IV provide unbiased estimates of the true effect θ. For example, this is likely in the immigration setting in cases in which regions that received immigrants in the past or were closer to the immigrants' origin are following different trends than the other regions. Our proposed approach can be thought of as a proxy control approach: if we could control for the omitted variable U_{it} (the differential trends in the immigration example) then the instrumental variable approach would be valid. This is represented in panel (d) of Figure 1. The statistical problem then becomes finding valid proxy controls for U_{it}, depending on whether we have additional data, or not, different strategies have been proposed in the proxy control literature (see Miao et al. (2018) or Deaner (2021)). Instead of relying on additional data, applied researchers often rely on controlling for linear trends (Wolfers, 2006), allowing year fixed effects to vary across groups of regions (Stephens Jr and Yang, 2014), or opting for a different design (for example synthetic control design if a valid donor pool exists, as in (Cengiz and Tekgüç, 2022)). An alternative approach based on synthetic differences in differences (Arkhangelsky et al., 2021b) has been proposed by Arkhangelsky and Korovkin (2023) in the case in which identification relies on aggregrate time series shocks. In this paper, we propose a strategy based on interpolation amongst units to control for U_{it} utilizing the popular synthetic control method.

To motivate why the setting described under Assumptions 1-2 and in Figure 1 is relevant to applied work, we consider an empirical example: the effect of the Syrian civil war on Turkey's local labor markets. The Syrian civil war started in March 2011 and by 2017, 6 million Syrians had sought shelter outside of Syria with 3.5 million locating in Turkey.[1] Figure 2 panel (a) shows the growth in the number of Syrian refugees in Turkey over time and panel (b) shows the geographic dispersion of the refugees. Given the structure of the Syrian refugee shock a natural approach to estimating the impact of refugees on local labor outcomes would be that of a shift-share instrumental variable design that exploits the exogenous time shock of the civil war and the differential impact across units.

To relate the Syrian example to our setting let R_{jt} denote the refugee/native ratio at province-year level and consider a travel distance shift-share instrument, as is common in the mass-immigration literature (Angrist and Kugler, 2003; Aksu et al., 2022).

$$Z_{jt} = \underbrace{\bar{H}_t}_{\text{shift}} \times \underbrace{Z_j}_{\text{share}},$$
$$Z_j = \sum_{s=1}^{13} \lambda_s \frac{1}{d_{j,s}}$$

where $\bar{H}_t$ is the number of refugees in Turkey in year t, $d_{j,s}$ is the travel distance between Turkish region p and Syrian governorate s, λ_s is the weight given to Syrian governorate s which we set it be proportional to the population share of s.[2] In panel (c) of Figure 2 we plot the first stage coefficients interacted with time dummies from the TWFE specification

$$R_{jt} = \sum_{j \neq 2010} \theta_j (\mathbb{1}\{t = j\} \times Z_j) + f_j + f_t + \eta_{jt}.$$

The first stage regression tests whether the instrument predicts refugees' location choice in every year. As expected, the distance is a strong predictor of the refugee treatment intensity. The F-stat of the shift-share first-stage (where we regress R_{jt} on Z_{jt} while controlling for region and time f.e.) is 125. The problem arises when one considers the reduced form of local wage-employment (salaried employment) of the natives that did not finish high-school

[1] Turkey hosts the largest number of refugees in the world (UNHCR, 2021).

[2] The idea is that all else equal, more Syrians would be expected to come from the more populous regions.

Figure 2: The Syrian refugee shock.

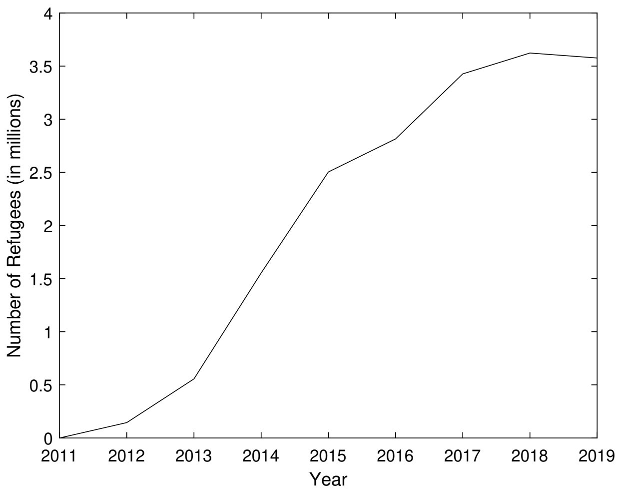

(a) Number of refugees

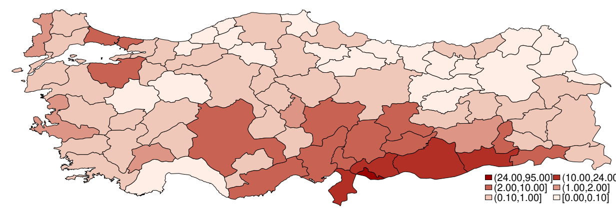

(b) Refugee shares

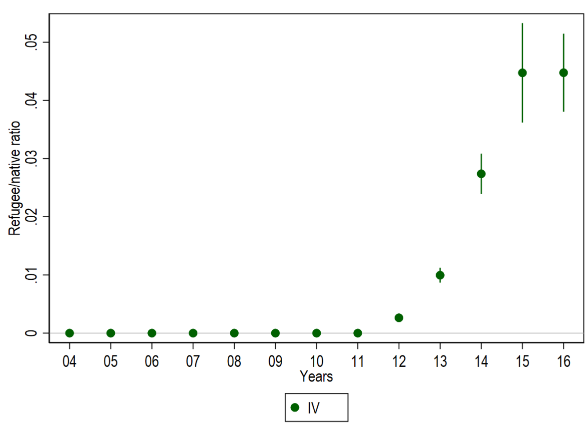

(c) First stage

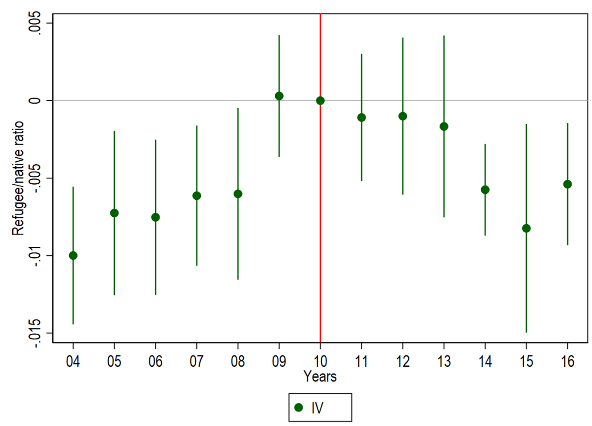

(d) Reduced form

Notes: In event-study designs the 95% confidence intervals are plotted. The F-stat of the main first-stage regression is 154. In Panels C and D, x axis shows the 2000s in 2 digit: e.g., 04 means 2004.

(low-skill)[3] on the instrument

$$Y_{jt} = \sum_{j \neq 2010} \beta_j (\mathbb{1}\{t = j\} \times Z_j) + f_j + f_t + \epsilon_{jt},$$

which is displayed in panel (d) of Figure 2. Between 2004–2010 (before the refugee crisis began), the provinces closer to the border observed employment gains compared to other regions. Being one standard deviation closer to the border predicts a wage-employment growth of 1 pp between 2004 and 2009. Put differently, there are "pre-trends" in the data. Given that the regions that are predicted by the instrument to receive immigrants were following different trends *before* the shock, it is likely that the IVDID design does not satisfy the parallel trend assumption. This is despite the fact that we have included region and year fixed effects in the regression specification. The appearance of pre-trends in similar designs is a common problem in practice (Wolfers, 2006; Stephens Jr and Yang, 2014; Gulek, 2023) and has been discussed extensively in the literature (Roth, 2022; Freyaldenhoven et al., 2019).

3. The synthetic estimator

The synthetic estimator consists of two steps. In the first step we find synthetic controls for each unit in a pre period ($t < T_0$) and generate counterfactual estimates for Y_{it}, R_{it} and Z_{it} for a post period. In the second step, as in the standard IV estimator, we use these counterfactual estimates to compute the first stage and reduced form estimates. To describe the procedure, consider J units indexed by $j = 1, \dots, J$ observed for T periods of time. We are interested in an outcome of interest Y_{it} with potential outcomes $Y_{it}(R_{it})$ indexed by random variable R_{it}.

Step 1: for each $j \in \{1, \dots, J\}$ we find the synthetic control weights w_j^* by solving the standard SC program for $t \in \{1, \dots, T_0\}$. For each j,

$$\hat{w}_j^{SC} \in \text{argmin}_{w \in \Delta^{J-1}} \| Y_j^{T_0} - Y_{-j}^{T_0}{}' w \|^2,$$

where Y^{T_0} is the $J \times T_0$ design matrix that includes pre-treatment outcomes Y_{jt} for $t < T_0$, with $Y_j^{T_0}$ denoting the predictors for unit j and $Y_{-j}^{T_0}$ the $J - 1 \times T_0$ matrix of predictors for the other units. In the case in which A1.1 holds instead of A1.2, we include R_{it} in the design

[3] This is the key outcome of interest because Syrian refugees were substantially less educated compared to the Turkish population, and hence constitute largely a low-skill immigration shock. We provide more details about the setting in the Appendix.

matrix to compute the synthetic control weights as well as Y_{it}.

Then, we define the following quantities for any t:

$$\hat{Y}_{it}^{SC} = \sum_{j \neq i} \hat{w}_{ij}^{SC} Y_{jt},$$
$$\hat{R}_{it}^{SC} = \sum_{j \neq i} \hat{w}_{ij}^{SC} R_{jt},$$
$$\hat{Z}_{it}^{SC} = \sum_{j \neq i} \hat{w}_{ij}^{SC} Z_{jt},$$

which we label the synthetic outcome, treatment level and instrument respectively. Similarly, we compute the *debiased* values for $t > T_0$:

$$\tilde{Y}_{it} = Y_{it} - \hat{Y}_{it}^{SC},$$
$$\tilde{R}_{it} = R_{it} - \hat{R}_{it}^{SC},$$
$$\tilde{Z}_{it} = Z_{it} - \hat{Z}_{it}^{SC}.$$

Step 2: Given $\{\tilde{Y}_{it}, \tilde{R}_{it}, \tilde{Z}_{it}\}_{t=T_0+1}^{T}$, we estimate the first stage and reduced form by OLS:

$$\tilde{\pi} \in \arg\min_{\beta} (\tilde{Y} - \tilde{Z}\beta)'(\tilde{Y} - \tilde{Z}\beta),$$
$$\tilde{\beta}_1 \in \arg\min_{\pi} (\tilde{R} - \tilde{Z}\pi)'(\tilde{R} - \tilde{Z}\pi).$$

where we may also includes observed covariates X in the regression. Then, the estimated average marginal effect is given by:

$$\tilde{\theta}^{SIV} = \frac{\tilde{\pi}}{\tilde{\beta}_1},$$

which we label the *synthetic* IV estimator.

Observe that this estimator is equivalent to the *synthetic TSLS* estimator given by

$$\tilde{\theta}^{TSLS} = \left(\sum_{it} \tilde{Z}_{it} \tilde{R}_{it} \right)^{-1} \sum_{it} \tilde{Z}_{it} \tilde{Y}_{it}.$$

Alternatively, an estimator that identifies the same parameter would be one using the instrument Z_{it} instead of the de-biased instrument $\tilde{Z}_{it}$

$$\tilde{\theta}_Z^{TSLS} = \left(\sum_{it} Z_{it}\tilde{R}_{it}\right)^{-1} \sum_{it} Z_{it}\tilde{Y}_{it}.$$

which is equivalent to the IV estimator using the reduced form and first stage with the instrument Z_{it} instead of $\tilde{Z}_{it}$. In the theory and simulation sections we show that while this estimator is similar to the proposed *synthetic* IV estimator, it may have worse finite sample properties, in particular in the case in which the instrument is not strong. For a discussion of identification in terms of potential outcomes of similar parameters in related IV difference-in-difference settings see Borusyak and Hull (2020).

3.1. Two additional estimators

The set up described under A1-A2 highlights a trade off between using the instrument variation to address the endogeneity bias due to the correlation between ϵ and η and incurring an omitted variable bias due to the instrument's correlation with $\mu_i' F_t$. Our synthetic IV estimator can address these biases asymptotically in regimes in which σ_ϵ is small as we highlight in sections 4 and 5. However, when σ_ϵ is large the endogeneity concern becomes more important than the omitted variable bias and, therefore, we might be able to design an estimator that addresses this bias directly. With this in mind, we provide two additional estimators that address this bias directly and propose an ensemble estimator as a 'doubly robust' alternative to the synthetic IV.

In cases in which the instrument also follows a factor structure, $Z_{it} = Z_i' G_t$, a natural estimator is to compute the synthetic control weights after projecting the outcome variable in the instrument space. The intuition for this estimator is that the outcome Y_{it} is noisy due to the unobserved error ϵ, but given our partial instrument validity assumption A2, after projecting the outcome in the instrument space we partial out the noise.

The **projected** synthetic estimator can be computed in the following steps

1. "de-noise" by projecting: $Y_{zt} = Z(Z'Z)^{-1}Z'Y_t$, where $Z = (Z_1, \ldots, Z_J)'$ and $Y_t = (Y_{1t}, \ldots, Y_{Jt})'$ are $J \times 1$ vectors.

2. Use the de-noised outcomes to compute the SC weights in the pre-period

$$w_j^P \in \text{argmin}_{w \in \Delta^{J-1}} \|Y_j^{T_0} - Y_{z,-j}^{T_0}{}' w\|^2.$$

3. Define the de-biased quantities $\tilde{Y}_{it}^P$, $\tilde{Z}_{it}^P$, $\tilde{R}_{it}^P$.

4. Estimate the synthetic TSLS projected estimator

$$\tilde{\theta}^{\text{Proj}} = \left(\sum_{it} \tilde{Z}_{it}^P \tilde{R}_{it}^P \right)^{-1} \sum_{it} \tilde{Z}_{it}^P \tilde{Y}_{it}^P.$$

Alternatively, in cases in which we believe the time series structure is not very important aggregating the outcome variables in the pre-period can reduce the dependence of the synthetic control match on the error terms ϵ. The intuition follows from a common rule of thumb for synthetic controls: aggregation can help reduce idiosyncratic noise.
The **aggregated** synthetic estimator can be computed as follows

1. Let $Q_i = \sum_{t<T_0} Z_{it} Y_{it}$.

2. Compute the SC weights in the pre-period using the aggregated quantities

$$w_j^{Agg} \in \text{argmin}_{w \in \Delta^{J-1}} \|Q_i - Q'_{-i} w\|^2.$$

3. Define the de-biased quantities $\tilde{Y}_{it}^A$, $\tilde{Z}_{it}^A$, $\tilde{R}_{it}^A$.

4. Estimate the synthetic TSLS aggregated estimator

$$\tilde{\theta}^{\text{Agg}} = \left(\sum_{it} \tilde{Z}_{it}^A \tilde{R}_{it}^A \right)^{-1} \sum_{it} \tilde{Z}_{it}^A \tilde{Y}_{it}^A.$$

The performance of each estimator will depend on the data generating process primitives. In cases in which the noise error term ϵ is more important than the factor term $\mu_i' F_t$ the projected estimator will perform favorably. On the other hand, if the factor structure (the signal) dominates, the projected estimator will perform worse than the synthetic IV as it will fit the factor structure worse. To account for the different biases in each case, we can construct an ensemble estimator that combines both estimators. For a hyper-parameter $\alpha \in [0, 1]$ let the **ensemble** estimator be defined as

$$\tilde{\theta}^E(\alpha) = \alpha \tilde{\theta}^{TSLS} + (1 - \alpha) \tilde{\theta}^P,$$

with an analogous definition for the aggregated synthetic IV. The α hyper-parameter can be chosen through cross-validation in the pre-period to optimize the mean squared error of the synthetic control estimator. The following steps detail how to compute the ensemble estimator

1. Split the pre-period into a training period $1, \ldots, T_v$ and a validation period $T_v + 1, \ldots, T_0$.
2. In the training period compute the synthetic control weights for each estimator, $\hat{w}^P$ and $\hat{w}$, and the debiased outcomes $\tilde{Y}_{it}^P$ and $\tilde{Y}_{it}$.
3. In the validation period choose α^* to minimize the mean squared error in the validation

$$\frac{1}{J(T_0 - T_v)} \|\alpha \tilde{Y}^{P,T_v} + (1-\alpha)\tilde{Y}^{T_v}\|_2^2,$$

   where $\tilde{Y}^{T_v}$ denotes the debiased outcomes for the validation period.

4. Compute the ensemble estimator in the post period as $\alpha^* \tilde{\theta}^{TSLS} + (1-\alpha^*)\tilde{\theta}^P$.

In the following two sections we discuss the theoretical properties of the synthetic IV estimator and finite sample properties through simulations. We highlight that in well behaved settings with low noise but significant correlation between the instrument and the unobserved factor structure, the aggregated, projected, and ensemble estimators will perform worse than the synthetic IV. However, in nosier settings the ensemble estimator appears 'doubly robust' and has lower finite sample bias than the synthetic IV estimator.

4. Theoretical Results

We start our theoretical discussion by noting that the standard TSLS estimator suffers from omitted variable bias under A1-A2. As expected, given the presence of the unobserved factor structure $\mu_i' F_t$, the estimator will be asymptotically biased. More precisely, let $T_1 = T - T_0$ and suppose for exposition that as $JT_1 \to \infty$

$$\frac{1}{JT_1} \sum_{it} Z_{it} \mu_{it}' F_t \xrightarrow{p} \xi,$$
$$\frac{1}{JT_1} \sum_{it} Z_{it}^2 \xrightarrow{p} Q > 0.$$

Then, as is standard,

$$\begin{aligned}\hat{\theta}^{TSLS} &= \theta + \left(\sum_{it} Z_{it}^2\right)^{-1} \sum_{it} Z_{it}\mu_{it}' F_t + o_p(1) \\ &= \theta + Q^{-1}\xi + o_p(1).\end{aligned}$$

The size of the bias will depend on the correlation between the instrument Z_{it} and the unobserved factors $\mu_i' F_t$, in the case, as in Figure 2, in which clear pre-trends are present the bias could be large. In a simulation study in Section 5 we investigate the size of the bias under different regimes and show that in many settings TSLS can be as biased as OLS.
To analyze our proposed *synthetic IV* estimator, we proceed similarly. In Section 3 we defined the *debiased* variables $\tilde{Y}_{it}$ and $\tilde{R}_{it}$, and under A1 it can be shown that

$$\begin{aligned}\tilde{Y}_{it} &= Y_{it} - \hat{Y}_{it}^{SC} \\ &= \theta R_{it} + \mu_i' F_t + \epsilon_{it} - \sum_{j\neq i} \hat{w}_{ij}^{SC} Y_{jt} \\ &= \theta \tilde{R}_{it} + (\mu_i - \sum_{j\neq i} \hat{w}_{ij}^{SC}\mu_j)' F_t + \epsilon_{it} - \sum_{j\neq i} \hat{w}^{SC}\epsilon_{jt}.\end{aligned}$$

It follows that the *synthetic* IV estimator for the regression of $\tilde{Y}$ on $\tilde{R}$ instrumented by $\tilde{Z}$ for $t > T_0$ recuperates θ up to two potential bias terms

$$\begin{aligned}\tilde{\theta}^{TSLS} &= \left(\sum_{it} \tilde{Z}_{it}\tilde{R}_{it}\right)^{-1} \sum_{it} \tilde{Z}_{it}\tilde{Y}_{it} \\ &= \theta + \left(\sum_{it} \tilde{Z}_{it}\tilde{R}_{it}\right)^{-1} \sum_{it} \tilde{Z}_{it} \left(\mu_i - \sum_{j\neq i} \hat{w}_{ij}^{SC}\mu_j\right)' F_t \\ &\quad + \left(\sum_{it} \tilde{Z}_{it}\tilde{R}_{it}\right)^{-1} \sum_{it} \tilde{Z}_{it} \left(\epsilon_{it} - \sum_{j\neq i} \hat{w}_{ij}^{SC}\epsilon_{jt}\right).\end{aligned}$$

The second term depends on the instrument and the unobserved error term ϵ and given our partial instrument validity assumption A2 can be shown to be zero in expectation when the error terms are i.i.d. In the appendix, we derive a finite sample bound in the case in which the errors are sub-gaussian. On the other hand, the first term depends on the unobserved factor structure $\mu_i' F_t$ and will not in general be zero in expectation or $o_p(1)$. Hence, in general the synthetic IV estimator will be biased in finite samples. Under weak conditions on the time series (β-mixing, or covariance stationarity), it may be possible to directly correct this bias using a cross-fitting procedure in the spirit of the method proposed in Chernozhukov et al. (2022) for the synthetic control framework. In this paper we take a different approach and provide general conditions on the data generating process such that the synthetic IV estimator is consistent when large pre and post periods are available. The following assumption imposes more structure on the design primitives.

Assumption 3 *[Model primitives] Assumptions on the factor structure, the error components and the instruments are as follows.*

- *The common factors are bounded such that for all t, $|F_{lt}| \leq \bar{F}$ for $l = 1, \dots, k$. Furthermore, the matrix $F_{T_0}' F_{T_0}$ has minimum eigenvalue ξ such that $\xi / T_0 > 0$. The factor loadings have bounded diameter such that $\mu_i \in \mathcal{M}$ with diam$(\mathcal{M}) = \sup\{\|t - s\| : \text{for } t, s \in \mathcal{M}\} \leq c_\mu$.*
- *The instruments have bounded diameter: $Z_{it} \in \mathcal{Z}$ such that diam$(\mathcal{Z}) = \sup\{\|t - s\| : \text{for } t, s \in \mathcal{Z}\} \leq c_z$.*
- *The instrument Z_{it} and the unobserved factor structure satisfy*

$$\frac{1}{JT} \sum_{it} Z_{it}^2 \xrightarrow{p} Q > 0,$$

  as $JT \to \infty$ and corr$(Z_{it}, \mu_i' F_t) < 1$ for all i, t.
- *ϵ_{it} and η_{it} are i.i.d mean zero subGaussian random variables with variance proxies σ_ϵ^2 and σ_η^2 respectively, finite covariance $\sigma_{\epsilon\eta} = \mathbb{E}[\epsilon_{it}\eta_{it}]$ and bounded fourth moments.*

Assumption 3 has three parts. First, we assume that the model primitives are bounded. This is a common assumption in papers analyzing the behavior of synthetic control estimators. Second, we assume that the instrument is strong and not perfectly correlated with the unobserved factor structure. This requirement avoids weak instrument problems and in the simulation discussion we highlight the importance of this assumption for the finite sample performance of the synthetic IV estimator. Finally, we assume that the unobserved error terms η and ϵ are i.i.d, but potentially correlated. This assumption can be weakened to allow for time series correlation, however in our main results the time series dependence is only present through the unobserved factor structure $\mu_i' F_t$. Under A1.2-A3 we can derive a bound on the unobserved factor term in the synthetic IV estimator.

Theorem 1 *[MAD bound] Under A1.2-A3, for $t > T_0$ the following bound holds for all J, T_1 and T_0*

$$\frac{1}{JT_1}\left|\mathbb{E}\left[\sum_{it}\tilde{Z}_{it}\tilde{\mu}_i' F_t\right]\right| \leq \left(\frac{\bar{F}^2 k}{\xi}\right) c_z \left(2\sqrt{\frac{J}{T_0}}\sigma_\epsilon + \mathbb{E}\left[\frac{1}{JT_0}\sum_{i,t\leq T_0}|\tilde{Y}_{jt}|\right]\right)$$

where all the terms are defined in A1-A3. Furthermore, under $A1 - A3$, as $JT_1 \to \infty$, $\mathbb{E}MAD(\tilde{Y}^{T_0}) \to 0$ and $\sqrt{\frac{J}{T_0}} \to 0$,

$$\frac{1}{JT_1}\sum_{it}\tilde{Z}_{it}\tilde{\mu}_i' F_t \xrightarrow{p} 0.$$

Theorem 1 states that the bias term that depends on the unobserved factor structure can be bounded above by the expected mean absolute deviation of the outcome variable in the pre-treatment period and a term that depends on the likelihood of pre-treatment "over-fitting". In the case in which A1.1 holds, the bound is similar and includes a mean absolute deviation term for $\tilde{R}_{it}$. This is a standard bound in papers evaluating the properties of synthetic control estimators (see Abadie et al. (2010) for the first example in the literature and Vives-i-Bastida (2022) for a example with covariates). It highlights the dependence of the estimator on good pre-treatment fit (see Ferman and Pinto (2021) for a discussion of synthetic controls with imperfect pre-treatment fit). In particular, the bound depends on the error noise level σ_ϵ and the ratio $\sqrt{J/T_0}$. In settings, in which the we have a small amount of pre-treatment periods, a large number of units, or in which the noise level is high, perfect interpolation of the noise is more likely, biasing the estimator. A discussion in Abadie and Vives-i-Bastida (2022) highlights the importance of pre-treatment fit and over-fitting for performance of synthetic control estimators through a simulation study. Similarly, we evaluate the performance of the synthetic IV estimator in simulations in section 5 and find that the estimator performs well even in settings with moderate $\sigma_\epsilon\sqrt{J/T_0}$.

To provide conditions under which $\mathbb{E}\text{MAD}(\tilde{Y}^{T_0}) \to 0$, we consider a relaxation of rank proposed by Rudelson and Vershynin (2007) that allows for *small* perturbations in the matrix.

Assumption 4 *[Numerical rank assumption]*

For all J and T_0 the design matrix has bounded numerical rank, $\frac{\|Y^{T_0\prime}\|_F^2}{\|Y^{T_0\prime}\|_2^2} \leq \bar{r}$, and its largest singular value is bounded above such that $\sigma_1(Y^{T_0}) \leq \bar{\sigma}_1$, where $\bar{r}$ and $\bar{\sigma}_1$ may depend on J and T_0.

The intuition behind Assumption 4 is better seen by considering the rank of the $J \times T_0$ design matrix Y^{T_0}. If the matrix had fixed rank $r < \min\{T_0, J\}$ all points would lie in a low dimensional manifold of the space and the pre-treatment fit error would grow proportional to r. Given that in our setting the error terms are $i.i.d$ shocks, this is not a reasonable assumption. Instead, we consider a bound on the numerical rank; the ratio between the Frobenius and 2-norm of a matrix. This notion of rank allows for points to lie "close" to a low dimensional manifold. Furthermore, for a matrix A it follows that

$$\frac{\|A\|_F^2}{\|A\|_2^2} \leq \text{rank}(A),$$

therefore the bounded numerical rank assumption is implied by a bounded rank assumption. Whether A4 is satisfied will depend on the model primitives. In particular, it will be satisfied when the signal to noise ratio is high. That is, when the factor structure $\mu_i' F_t$ dominates the noise term ϵ. In cases in which σ_ϵ is large relative to the factor term the numerical rank is likely to be large and the pre-treatment fit bad. Note that a bound on the fit of the outcome will also imply a bound on the fit of R_{it} under A1, therefore our bounds and results follow from A1.2 to A1.1 under A3 (with a worse rate). In section 5 we explore the performance of our estimator in a variety of settings and propose checks researchers can implement to evaluate whether their empirical setting is likely to satisfy this assumption.

Theorem 2 *[Factor term consistency] Under A1-A4, for $t > T_0$ the following bound holds for all J, T_1 and T_0*

$$\frac{1}{JT_1}\left|\mathbb{E}\left[\sum_{it} \tilde{Z}_{it}\tilde{\mu}_i' F_t\right]\right| \leq \left(\frac{\bar{F}^2 k}{\xi}\right) c_z \left(2\sqrt{\frac{J}{T_0}}\sigma_\epsilon + \bar{r}\bar{\sigma}_1\left[\frac{1}{\sqrt{T_0 J}} + \sqrt{\frac{J}{T_0}}\right]\right)$$

where all the terms are defined in A1-A4. Furthermore, as $JT_1 \to \infty$ and $\bar{r}\bar{\sigma}_1\sqrt{\frac{J}{T_0}} \to 0$,

$$\frac{1}{JT_1}\sum_{it} \tilde{Z}_{it}\tilde{\mu}_i' F_t \xrightarrow{p} 0.$$

Theorem 2 shows that the bias due to the factor term is $o_p(1)$ as long as $\bar{r}\bar{\sigma}_1\sqrt{\frac{J}{T_0}} \to 0$. For fixed J, this implies that we need $T_0, T_1 \to \infty$. The restrictions on $\bar{r}$ and $\bar{\sigma}_1$ are not uncommon in the matrix completion literature. Combining the consistency result with the additional assumptions on the instrument behavior we can show that the synthetic IV estimator is a consistent estimator of θ.

Theorem 3 *[Consistency] Under A1-A4 and a regularity condition, as $JT_1 \to \infty$ and $\bar{r}\bar{\sigma}_1\sqrt{\frac{J}{T_0}} \to 0$,*

$$\begin{aligned}
\tilde{\theta}^{TSLS} - \theta &\xrightarrow{p} 0, \\
\tilde{\theta}_Z^{TSLS} - \theta &\xrightarrow{p} 0, \\
\hat{\theta}^{TSLS} - \theta &\neq o_p(1).
\end{aligned}$$

Theorem 3 states that both the synthetic TSLS estimator $\tilde{\theta}^{TSLS}$ and the synthetic TSLS estimator for which we do not debias the instrument $\tilde{\theta}_Z^{TSLS}$ are consistent under our assumptions and the rate conditions of Theorem 2, while the standard TSLS estimator is not. As discussed, however, these estimators are biased in finite samples and the finite sample bias will depend on the signal to noise ratio, the length of the pre and post treatment periods in relation to the number of units J and, through the first stage, the correlation between Z_{it} and $\mu_i' F_t$. It is important to note that while debiasing the instrument does not affect the consistency of the estimator it may improve the finite sample properties of the estimator. In the appendix, we show under additional assumptions, that debiasing the instrument can lead to a stronger first stage and, therefore, better finite sample properties. We confirm this intuition in the simulation study by comparing $\tilde{\theta}^{TSLS}$ and $\tilde{\theta}_Z^{TSLS}$.

Under A1-A3 it is also possible to show that the synthetic IV estimator is asymptotically normal.

Theorem 4 *[Asymptotic normality] Under A1-A4, for $t > T_0$, if $\sqrt{\frac{T_1}{T_0}}(1+J)\bar{r}\bar{\sigma}_1 \to 0$ as $\sqrt{\frac{J}{T_1}} \to 0$, then*

$$\frac{\sqrt{JT_1}(\tilde{\theta}^{SW} - \theta)}{v_{JT_1}} \xrightarrow{d} N(0,1).$$

where $v_{JT_1} = (\mathbb{E}\tilde{Z}'\tilde{R})^{-2}\left(\mathbb{E}\left[\sum_{it} var(\tilde{Z}_{it}\tilde{\epsilon}_{it} \mid Z, w)\right]\right)$ and

$$\sum_{it} var(\tilde{Z}_{it}\tilde{\epsilon}_{it} \mid Z, w) = \sigma^2 \sum_{it}(1 + \|w_i\|^2)(Z_{it}^2 + \sum_{j \neq i} Z_{jt} w_{ij} - \sum_{i \neq j} Z_{jt} Z_{it} w_{ij}).$$

The asymptotic variance in Theorem 4, v_{JT_1}, has an additional term with respect to the TSLS variance, which accounts for the influence of the weights on the estimator distribution. The asymptotic normality result relies on the martingale characterization of matching estimators by Abadie and Imbens (2012) and can be used to provide confidence intervals using its sample counterpart $\hat{v}_{JT_1}$ where $\hat{\sigma}^2$ can be estimated from the regression residuals and plugged in $var(\tilde{Z}_{it}\tilde{\epsilon}_{it} \mid Z, w)$. In the heteroskedastic case, the quantity $\sum_{it} \sigma_{it}(1 + \|w_i\|^2)(Z_{it}^2 + \sum_{j\neq i} Z_{jt}w_{ij} - \sum_{i\neq j} Z_{jt}Z_{it}w_{ij})$ would need to be estimated. In the simulation section we study the coverage of the synthetic IV estimator in the homoskedastic case.

5. Simulation study

We consider a simulation design setting that resembles our empirical application. In particular, we impose the following design primitives:

$$
\begin{aligned}
Y_{it} &= \beta R_{it} + \mu_i' f_t + \epsilon_{it}, \\
R_{it} &= (\gamma Z_{it} + \eta_{it}) * \mathbb{1}(t \geq T_0), \\
Z_{it} &= Z_i' g_t * \mathbb{1}(t \geq T_0),
\end{aligned}
$$

with time series structure

$$
\begin{aligned}
f_t &= \kappa_f f_{t-1} + u_{ft}, \\
g_t &= \kappa_g g_{t-1} + u_{gt},
\end{aligned}
$$

and error structure

$$
\begin{aligned}
\begin{pmatrix} u_{ft} \\ g_{ft} \end{pmatrix} &\sim N\left(0, \begin{bmatrix} \sigma_f^2 & \rho_g\sigma_f\sigma_g \\ \rho_g\sigma_f\sigma_g & \sigma_g^2 \end{bmatrix}\right), \\
\begin{pmatrix} Z_i \\ \mu_i \end{pmatrix} &\sim N\left(0, \begin{bmatrix} \sigma_z^2 & \rho_z\sigma_z\sigma_\mu \\ \rho_z\sigma_z\sigma_\mu & \sigma_\mu^2 \end{bmatrix}\right), \\
\begin{pmatrix} \epsilon_{it} \\ \eta_{it} \end{pmatrix} &\sim N\left(0, \begin{bmatrix} \sigma_\epsilon^2 & \rho\sigma_\epsilon\sigma_\lambda \\ \rho\sigma_\epsilon\sigma_\lambda & \sigma_\lambda^2 \end{bmatrix}\right).
\end{aligned}
$$

The simulation design is rich and allows for a variety of empirical settings. The key parameters are ρ, that controls the degree of the *endogeneity* problem that can be addressed using the instrument, and ρ_z and ρ_g that control the *omitted variable bias* through the correlation between the instrument and the unobserved factors. In the case in which $\rho = \rho_z = \rho_g = 0$, there is no endogeneity nor omitted variable bias and the OLS and TSLS estimators are unbiased. As the ρs increase the OLS and TSLS estimators become biased and we can compare them to the proposed synthetic estimators. Figure 3 shows that the synthetic IV estimator is able to correct the bias present in the OLS and TSLS (with TWFE) estimators when there is moderate endogeneity and omitted variable bias. Panel (a) shows the case in which both estimators are consistent (the true parameter is 1) and the synthetic IV estimator performs similar to the OLS estimator. In panel (b) we increase the correlation between ϵ and η, creating an endogeneity problem that can be addressed using the instrument. As expected, the OLS estimator is now biased, while the TSLS and synthetic IV estimators remain consistent. In panel (c) we introduce correlation between the instrument and the unobserved factor structure, by setting $\rho = \rho_z = \rho_g = 0.5$, and the instrument becomes invalid leading to biased TSLS estimates despite adding TWFE in the specification. The synthetic IV on the other hand is approximately unbiased. To relate this simulation results to the 'pre-trends' discussion in Figure 2, panel (d) shows the event study coefficients (over 1000 simulations). Before the treatment starts at $T_0 = 20$, the coefficients should be close to zero, as the instrument is not active. As can be seen the TSLS estimator exhibits large deviations in the pre-period, while the synthetic IV exhibits substantially closer to zero "pre-trends".

As suggested by the theoretical properties in section 4, the performance of the synthetic IV estimator depends on the data generating process primitives, in particular σ_ϵ for the pre-treatment goodness of fit and the $corr(Z_{it}, \mu_i' f_t)$ for the strength of the first stage. Table 1 explores the sensitivity of the proposed estimators to varying ρ, ρ_z and ρ_g from 0.5 to 0.9 and σ_ϵ from 0.5 to 8, for the same simulation design as in Figure 3. The main takeaway from Table 1 is that the synthetic IV and the ensemble estimator which combines the synthetic IV and the projected estimator (shaded in gray in the figure) outperform the TWFE OLS and TSLS estimator in all settings in terms of both bias and mean squared error. Furthermore, both estimators exhibit close to zero bias in settings with small or moderate noise and correlation. In fact, the ensemble estimator appears to have low bias even in cases in which the noise level is 8 times the level of the signal ($\sigma_\epsilon = 8$ vs $\sigma_\mu = \sigma_\eta = 1$).

More concerning is the effect of a high $corr(Z_{it}, \mu_i' f_t)$. As can be seen in Table 1, for $\rho = \rho_z = \rho_g = 0.9$ while the synthetic IV exhibits significantly less bias than the TSLS; it remains biased with the bias increasing towards the TSLS as the noise level rises. This behavior mirrors the weak instrument problem: as the $corr(Z_{it}, \mu_i' f_t)$ increases to 1, the first stage of the debiased estimator becomes weaker as there is no variation left in the instrument that explains the outcome. This bias is not corrected by using Z instead the

Figure 3: Model comparison in simulations

Note: Panels (a)-(c) display kernel density plots for TWFE OLS, TWFE TSLS and the synthetic TSLS. Panel (d) shows simulated event study estimates as in Figure 2 panel (d) with 95% confidence bands for $\rho = \rho_z = \rho_g = 0.5$*. Simulations are done over 1000 iterations with the following parameters:* $\beta = \gamma = 1$*,* $k = 1$*,* $T = 30$*,* $T_0 = 20$*,* $J = 20$*,* $\sigma_\epsilon = 0.5$*,* $\kappa = 0.5$*,* $\sigma_\eta = \sigma_z = \sigma_g = 1$*.*

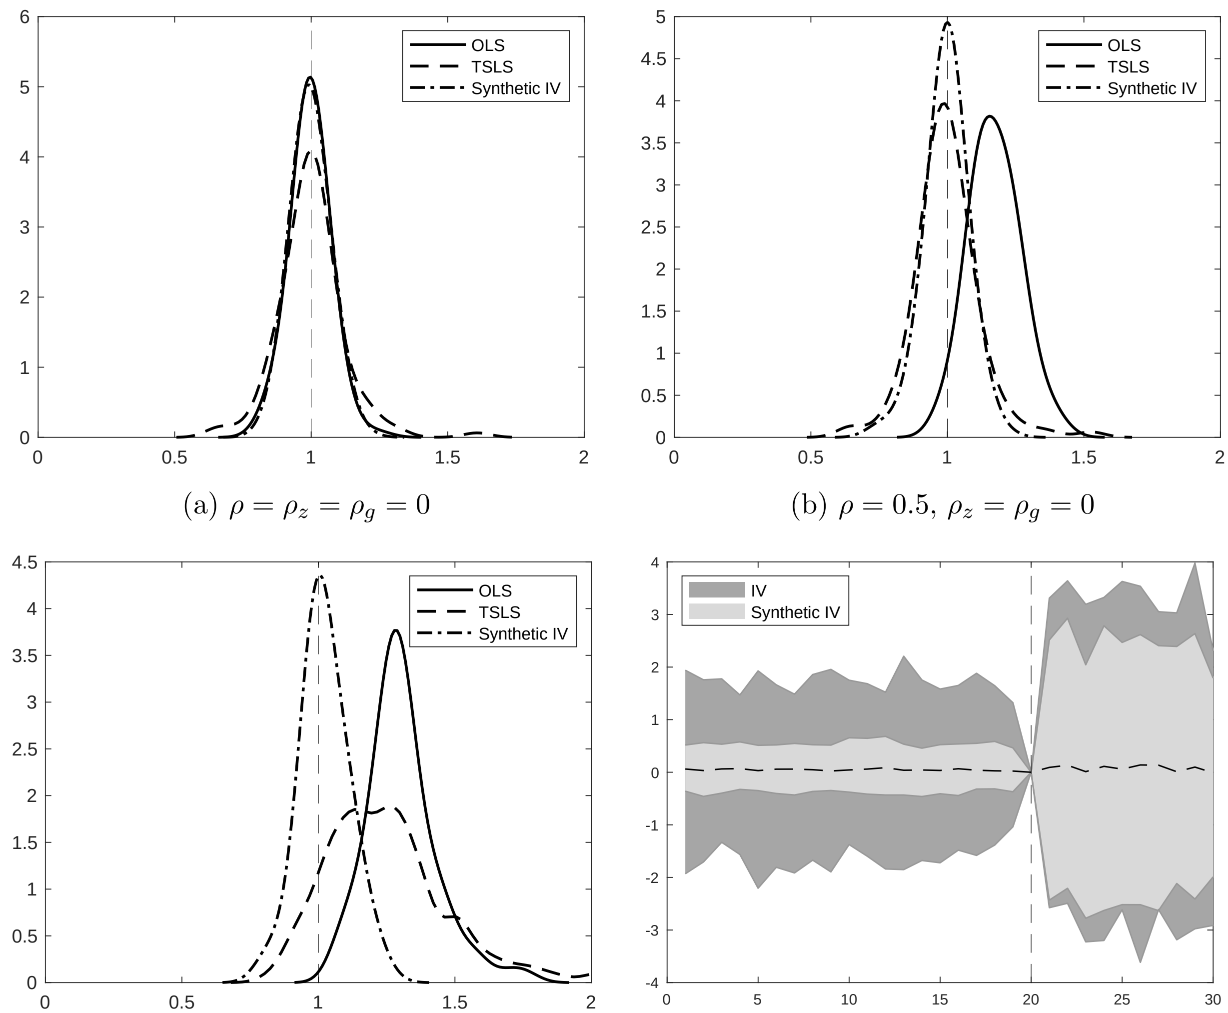

(a) $\rho = \rho_z = \rho_g = 0$

(b) $\rho = 0.5$, $\rho_z = \rho_g = 0$

(c) $\rho = \rho_z = \rho_g = 0.5$

(d) Event study estimates

debiased instrument $\tilde{Z}$, in fact it becomes worse as can be seen in the performance of the SIV Z estimator in Table 1. In the appendix, we provide a theoretical justification as for why debiasing the instrument might improve the first stage.

Another aspect highlighted by the theory and by reviews of best practices for synthetic control estimators (Abadie and Vives-i-Bastida 2021), is how the relative sizes of T_0, T_1 and J influence the behavior of the estimator. The consistency result requires requires that both JT_1 is large and $\sqrt{J/T_0}$ is small. In our baseline simulation we considered a setting with $JT_1 = 200$ and $\sqrt{J/T_0} = 1$. In the appendix, we replicate Table 1 in a setting with a short pre-treatment period $T_0 = 10, J = 20$. We find that while the performance of the estimators deteriorates (specially in settings with high noise and/or correlation) it is comparable to the results reported in Table 1 and the gains relative to OLS and TWFE are larger. Another consideration, given our asymptotic normality result, is the requirement that $\sqrt{\frac{T_1}{T_0}}(1+J)\bar{r}\bar{\sigma}_1 \to 0$ as $\sqrt{\frac{J}{T_1}} \to 0$. To see the effect of T_1 on the sample distribution of the estimator, in Table 2 we evaluate the coverage of the 95% confidence intervals for the synthetic IV using $\hat{v}_{JT_1}$ for different correlation settings and T_1s. We find that in settings in which the OLS and TSLS are unbiased the synthetic IV exhibits a slight over-coverage, in the well behaved settings with moderate noise and correlation the coverage is good, and in high correlation settings, as expected, we report under-coverage.

Table 1: Simulations for $\rho = \rho_z = \rho_g = r$ and different σ_ϵ.

	r=0.5				r=0.7				r=0.9			
	Mean	Var	Bias	MSE	Mean	Var	Bias	MSE	Mean	Var	Bias	MSE
					$\sigma_\epsilon = 0.5$							
OLS (TWFE)	1.31	0.02	0.31	0.11	1.50	0.02	0.50	0.27	1.73	0.01	0.73	0.55
TSLS (TWFE)	1.26	0.07	0.26	0.13	1.51	0.08	0.51	0.34	1.83	0.06	0.83	0.74
SIV	1.02	0.01	0.02	0.01	1.05	0.02	0.05	0.02	1.19	0.04	0.19	0.07
projected SIV	0.92	0.03	-0.08	0.04	0.95	0.04	-0.05	0.05	1.11	0.07	0.11	0.08
Agg. SIV	1.23	0.08	0.23	0.13	1.46	0.08	0.46	0.29	1.80	0.04	0.80	0.68
SIV + projected	1.01	0.01	0.01	0.01	1.03	0.02	0.03	0.02	1.15	0.04	0.15	0.06
SIV + Agg.	1.03	0.01	0.03	0.01	1.07	0.02	0.07	0.02	1.21	0.04	0.21	0.08
SIZ Z	1.07	0.02	0.07	0.02	1.15	0.03	0.15	0.05	1.43	0.03	0.43	0.21
					$\sigma_\epsilon = 1$							
OLS (TWFE)	1.38	0.02	0.38	0.16	1.60	0.02	0.60	0.38	1.86	0.02	0.86	0.76
TSLS (TWFE)	1.26	0.07	0.26	0.14	1.50	0.08	0.50	0.34	1.82	0.06	0.82	0.74
SIV	1.03	0.01	0.03	0.01	1.07	0.03	0.07	0.03	1.26	0.05	0.26	0.12
projected SIV	0.90	0.05	-0.10	0.06	0.94	0.06	-0.06	0.07	1.14	0.08	0.14	0.10
Agg. SIV	1.22	0.07	0.22	0.12	1.47	0.08	0.47	0.30	1.80	0.04	0.80	0.69
SIV + projected	1.01	0.01	0.01	0.01	1.03	0.03	0.03	0.03	1.21	0.05	0.21	0.10
SIV + Agg.	1.04	0.01	0.04	0.02	1.10	0.03	0.10	0.04	1.30	0.05	0.30	0.14
SIZ Z	1.08	0.02	0.08	0.03	1.19	0.03	0.19	0.07	1.50	0.03	0.50	0.28
					$\sigma_\epsilon = 2$							
OLS (TWFE)	1.48	0.02	0.48	0.26	1.74	0.03	0.74	0.58	2.05	0.02	1.05	1.12
TSLS (TWFE)	1.26	0.08	0.26	0.14	1.50	0.09	0.50	0.34	1.82	0.07	0.82	0.74
SIV	1.05	0.03	0.05	0.03	1.12	0.04	0.12	0.06	1.37	0.07	0.37	0.21
projected SIV	0.87	0.08	-0.13	0.09	0.93	0.10	-0.07	0.10	1.21	0.10	0.21	0.15
Agg. SIV	1.22	0.08	0.22	0.12	1.46	0.09	0.46	0.30	1.80	0.05	0.80	0.68
SIV + projected	1.01	0.03	0.01	0.03	1.06	0.05	0.06	0.05	1.29	0.08	0.29	0.16
SIV + Agg.	1.07	0.03	0.07	0.03	1.16	0.05	0.16	0.07	1.43	0.07	0.43	0.26
SIZ Z	1.10	0.03	0.10	0.04	1.24	0.05	0.24	0.10	1.58	0.04	0.58	0.38
					$\sigma_\epsilon = 4$							
OLS (TWFE)	1.63	0.04	0.63	0.43	1.95	0.04	0.95	0.94	2.31	0.04	1.31	1.75
TSLS (TWFE)	1.25	0.09	0.25	0.15	1.49	0.11	0.49	0.35	1.81	0.08	0.81	0.74
SIV	1.08	0.05	0.08	0.05	1.19	0.07	0.19	0.11	1.49	0.10	0.49	0.34
projected SIV	0.85	0.13	-0.15	0.15	0.95	0.15	-0.05	0.15	1.30	0.15	0.30	0.24
Agg. SIV	1.20	0.10	0.20	0.14	1.45	0.11	0.45	0.31	1.79	0.07	0.79	0.69
SIV + projected	1.01	0.05	0.01	0.05	1.10	0.08	0.10	0.09	1.41	0.11	0.41	0.28
SIV + Agg.	1.11	0.05	0.11	0.06	1.24	0.07	0.24	0.13	1.56	0.09	0.56	0.40
SIZ Z	1.13	0.06	0.13	0.07	1.30	0.07	0.30	0.16	1.65	0.06	0.65	0.48
					$\sigma_\epsilon = 8$							
OLS (TWFE)	1.83	0.06	0.83	0.75	2.24	0.08	1.24	1.61	2.68	0.09	1.68	2.91
TSLS (TWFE)	1.24	0.11	0.24	0.17	1.49	0.13	0.49	0.37	1.80	0.11	0.80	0.76
SIV	1.12	0.09	0.12	0.11	1.27	0.12	0.27	0.19	1.61	0.13	0.61	0.50
projected SIV	0.88	0.22	-0.12	0.24	1.01	0.24	0.01	0.24	1.42	0.25	0.42	0.42
Agg. SIV	1.19	0.15	0.19	0.19	1.43	0.15	0.43	0.34	1.78	0.11	0.78	0.72
SIV + projected	1.05	0.10	0.05	0.10	1.18	0.13	0.18	0.17	1.53	0.16	0.53	0.44
SIV + Agg.	1.15	0.10	0.15	0.12	1.32	0.12	0.32	0.23	1.66	0.11	0.66	0.55
SIZ Z	1.16	0.09	0.16	0.12	1.36	0.11	0.36	0.24	1.71	0.09	0.71	0.59

Table 2: $\mathrm{T}_0 = 20, J = 20, \sigma_\epsilon = 0.5, \sigma_z = 1, \sigma_{other} = 0.5, \kappa = 0.5$.

Coverage $\alpha = 0.05$			
	T=30	T=40	T=50
$\rho = \rho_g = \rho_z = 0.0$	0.981	0.962	0.952
$\rho = \rho_g = \rho_z = 0.3$	0.976	0.944	0.96
$\rho = \rho_g = \rho_z = 0.5$	0.960	0.945	0.923
$\rho = \rho_g = \rho_z = 0.7$	0.904	0.808	0.792

With the simulation results in mind we highlight three robustness checks that practitioners should implement when using synthetic IV or similar estimators:

1. **Check your first stage**: given that the debiasing procedure can lead to a weaker first stage, in cases with strong omitted variable bias if the synthetic IV estimator exhibits a weak first stage researchers may be worried using the synthetic estimator.

2. **Check your pre-treatment fit**: if the debiased outcomes exhibits large deviations in the pre-treatment period or an event study design reveals pre-trends, it is likely that the synthetic estimator will be biased. This bias however, may still be smaller than the TSLS bias.

3. **Back test**: given that the finite sample bias depends on the expected pre-treatment fit, back testing can reveal whether good pre-treatment fit was due to over-fitting (biasing the estimator) or not.

In the following section we implement these robustness checks when re-evaluating the effect of the Syrian refugee crisis using the synthetic IV.

6. Revisiting the Syrian refugee shock

With the SIV tool at hand, we now revisit our analysis of the impact of Syrian refugees on salaried employment of low-skill natives. We first solve the Synthetic Control problem using demeaned data between 2004–2010.[4] Following step 1 of our algorithm, we create synthetic regions with outcome Y^{SC}, treatment R^{SC} and instrument Z^{SC}. We then debias the data by subtracting the raw data with the synthetic data, generating $\tilde{Y}_{it}$, $\tilde{R}_{it}$ and $\tilde{Z}_{it}$.

[4]Demeaning the individual regions is an important detail in the Turkish setting due to the large heterogeneity in development rates across regions. For example, Istanbul is the most developed region with the highest employment rate in the country. No convex combination of other regions can match Istanbul on levels, but matching on trends is feasible.

Before estimating the treatment effect via TSLS on the debiased data following step 2 of the algorithm, we first check the quality of the matching in the pre-period. As discussed in the theory section, goodness of fit is a necessary component to get consistent estimates using SIV. One way to check for goodness of fit is to plot the debiased data. During the training period, the debiased data should fluctuate little (if any) around zero. We plot the debiased wage-employment data in Figure 4a, where black dashed lines belong to the less intensely treated regions that received less than 2% of refugees compared to their native population by 2016, and the green straight lines belong to the more intensely treated regions. During the training period 2004–2010, the debiased data fluctuate little around zero, which implies that we were able to match well on the trends.

The second check we perform is to look at the first-stage using the debiased data. The matching algorithm does not enforce maintaining a first-stage. For intuition, consider the case when the instrument is a binary indicator, and imagine that regions with $Z_i = 1$ follow one trend, and regions with $Z_i = 0$ follow another. Even if there is a first-stage in the raw data, $\tilde{Z}$ would be zero for the debiased data. We plot the first-stage estimates in Figure 4b. In our case, the debiased data maintains the strong first-stage. In a regression of $\tilde{R}$ on $\tilde{Z}$ while controlling for two-way fixed effects, the F-stat is 218.

The third check we perform is to look at the reduced-form using the debiased data. If the matching was successful, i.e., the donor pool had regions with similar trends for all the regions in the sample, then the event-study design on the debiased data should find estimates around zero in the pre-period. In particular, we estimate the following equations:

$$\begin{aligned} Y_{jt} &= \sum_{j \neq 2010} \beta_j (\mathbb{1}\{t = j\} \times Z_j) + f_j + f_t + \epsilon_{jt} \\ \tilde{Y}_{jt} &= \sum_{j \neq 2010} \tilde{\beta}_j (\mathbb{1}\{t = j\} \times \tilde{Z}_j) + g_j + g_t + \xi_{jt} \end{aligned} \tag{1}$$

where the first line is the reduced-form of the IV, and the second line is the reduced-form of the SIV. We plot the estimates of β and $\tilde{\beta}$ in Figure 4c. The results are striking. First, SIV trained using data between 2004–2010 completely eliminated the pre-trends. This means that for all regions the algorithm was able to find a convex combination of regions that had similar trends. Second, adjusting for pre-rends, SIV finds slightly stronger disemployment effects in the post-period.

One could argue that the non-existence of pre-trend in the reduced-form of SIV is mechanical. When the training is done over the entire pre-period, the reduced-form can give

Figure 4: Quality Checks

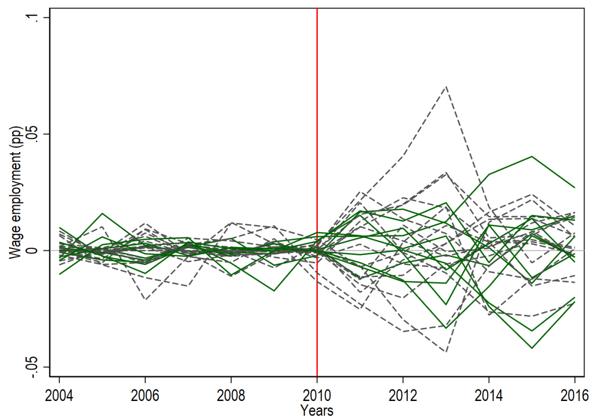

(a) Debiased Data

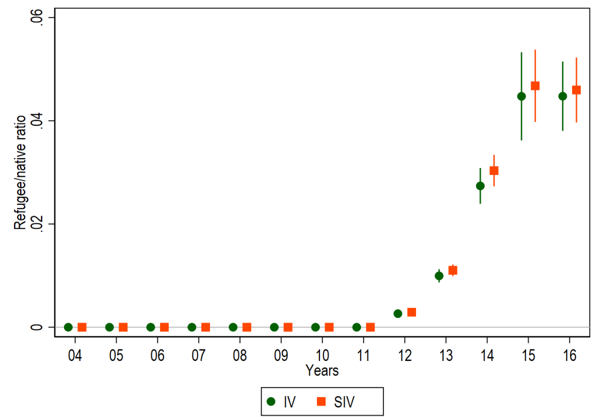

(b) Debiased First-Stage

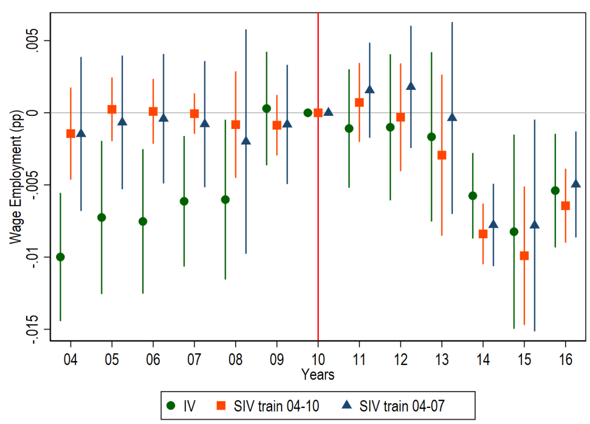

(c) Debiased Reduced-Form

Notes: Panel A uses the debiased data. The green solid lines belong to the intensely treated regions, the black dashed lines belong to the rest, and the cutoff is 2% refugee/native ratio. The first-stage using both raw and debiased data is plotted in Panel B. The F-stat in the main first-stage is 154 with the raw data and 218 with the debiased data. In Panel C, the reduced-form estimates come from the event-study design shown in equation 1. The outcome variable is the wage-employment rate of low-skill natives. Standard errors are clustered at the region level. The 95% confidence interval is plotted.

Figure 5: SIV Estimates

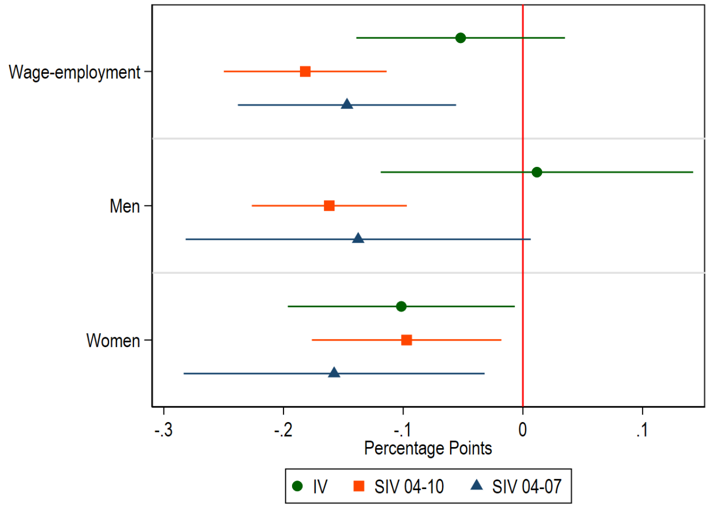

no pre-trends even when there is little signal in the data. To test for over-fitting bias, we perform back-testing. In particular, instead of using the entire pre-period in the matching, we solve for the SC weights using data between 2004-2007 and follow the rest of the algorithm as specified before. We plot the estimates in Figure 4c in blue. Despite the reduced amount of data that we match on, the reduced form does not find any placebo effect in the pre-period. All the estimates between 2004–2010 are both quantitatively close to zero and statistically insignificant, meaning synthetic distance is uncorrelated with the trends in the data in the pre-period. Notice that the largest change in the IV estimates in the pre-period occurs between 2008–2009. Despite being trained between 2004–2007, SIV is able to capture this change. This provides more evidence that the algorithm captures the signal in the data. Put differently, over-fitting is not a first-order concern in our setting.

It is worth discussing why IV and SIV estimates differ less in the post-period than the pre-period. It is unlikely to know for certain the nature of the unobserved confounder in the pre-period. If we knew it, we would add it as a control variable in the first place. However, there are some likely candidates that can explain the nature of the pre-trend. As explained in Gulek (2023), the regions close to the border are less-developed than the rest. Between 2004–2010, Turkey's GDP per capita grew by 75%. The data seems to suggest that the less developed south-east regions were "catching-up" to the rest of the Turkey. The employment rates were growing faster. This large growth rate did not last as Turkey entered into a recession starting from 2013. If economic growth in the pre-period was the main reason behind the pre-trends, it is likely that these pre-trends would not last in the post period. SIV seems to be capturing this change in the unobserved confounder(s) by not changing the post estimates by a huge margin.

Having seen how SIV *addresses* the pre-trend problem in the event-study design, we continue by implementing the second step of the algorithm: we apply 2SLS on the debiased data. We estimate the effect of Syrian refugees on natives' wage-employment. For heterogeneity, we also estimate the effects on men and women separately. We plot the estimates on Figure 5. As a benchmark, we first show the IV estimates. A researcher using IV would find no effect on men and negative effects on women. However, using SIV we find that Syrian refugees had dis-employment effects on both men and women. If anything, the effects on men were stronger. A 1 pp increase in refugee/native ratio decreases low-skill natives' salaried employment rate by 0.16 pp for men and 0.10 pp for women. As a robustness check, we also show the results that rely on estimated weights using the 2004–2007 data. The results remain quantitatively and qualitatively very similar. This is not surprising given that the event-study design showed no concern for over-fitting bias in Figure 4c.

It is worth highlighting how much our method impacts the economic conclusions in our setting. Turkey hosts the largest number of refugees in the world. The three most treated NUTS-2 level regions in Turkey observed an increase in labor supply of more than 10% in practically five years. Refugees, especially men, have a high propensity to work: 87% of prime age men are "employed" in Turkey (Turkish Red Crescent and WFP, 2019). Despite this *large* labor supply shock, in a short enough time period where spatial markets are unlikely to equilibrate (which would violate stable unit value treatment assumption embedded under the Spatial IV-DiD framework) and despite male refugees' having higher employment rates than females, IV finds no dis-employment effects for native men. Theoretically justifying this result would require either completely flat labor demand curves (Borjas, 2003) or refugees to provide a huge positive product demand shock (Borjas, 2014). There is very little empirical evidence for both, especially considering that Syrian refugees left most of their wealth behind while escaping a civil war. SIV reveals that this large labor supply shock has caused native dis-employment in the very short run, for both men and women, which is consistent with economic theory.

7. Conclusion

In this paper we provide a new method, the synthetic IV, to deal with unmeasured confounding in panel data settings in which researchers have access to an instrumental variable that is only partially valid. By assuming a factor structure on the unobserved confounding term we derive conditions under which a synthetic IV estimator, that combines synthetic controls and two-stage least squares, is consistent and asymptotically normal. Through a simulation study, we show that the estimator performs well in a variety of empirical settings and removes the bias in cases in which TSLS and OLS with two-way fixed effects do not. In an empirical application to the Syrian refugee crisis we showcase the utility of the synthetic IV in correcting worrisome pre-trends and highlight robustness checks researchers should implement when using the method.

References

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505.

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2):495–510.

Abadie, A. and Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132.

Abadie, A. and Imbens, G. W. (2012). A martingale representation for matching estimators. *Journal of the American Statistical Association*, 107(498):833–843.

Abadie, A. and Vives-i-Bastida, J. (2022). Synthetic controls in action.

Abadie, A. and Zhao, J. (2022). Synthetic controls for experimental design.

Adao, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010.

Agarwal, A., Dahleh, M., Shah, D., and Shen, D. (2021). Causal matrix completion.

Aksu, E., Erzan, R., and Kırdar, M. G. (2022). The impact of mass migration of syrians on the turkish labor market. *Labour Economics*, page 102183.

Anatolyev, S. and Mikusheva, A. (2022). Factor models with many assets: Strong factors, weak factors, and the two-pass procedure. *Journal of Econometrics*, 229(1):103–126.

Angrist, J. D. and Kugler, A. D. (2003). Protective or counter-productive? labour market institutions and the effect of immigration one unatives. *The Economic Journal*, 113(488):F302–F331.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021a). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021b). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

Arkhangelsky, D. and Korovkin, V. (2023). On policy evaluation with aggregate time-series shocks.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G., and Khosravi, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 116(536):1716–1730.

Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279.

Ben-Michael, E., Feller, A., and Rothstein, J. (2021). The augmented synthetic control method. *Journal of the American Statistical Association*. Forthcoming.

Borjas, G. J. (2003). The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *The quarterly journal of economics*, 118(4):1335–1374.

Borjas, G. J. (2014). Immigration economics. In *Immigration Economics*. Harvard University Press.

Borjas, G. J. (2017). The wage impact of the marielitos: A reappraisal. *ILR Review*, 70(5):1077–1110.

Borusyak, K. and Hull, P. (2020). Non-Random Exposure to Exogenous Shocks: Theory and Applications. NBER Working Papers 27845, National Bureau of Economic Research, Inc.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213.

Borusyak, K., Jaravel, X., and Spiess, J. (2023). Revisiting event study designs: Robust and efficient estimation.

Card, D. (1990). The impact of the mariel boatlift on the miami labor market. *ILR Review*, 43(2):245–257.

Card, D. (2001). Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64.

Card, D. and Krueger, A. B. (2000). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania: Reply. *The American Economic Review*, 90(5):1397–1420.

Cengiz, D. and Tekgüç, H. (2022). Is it merely a labor supply shock? impacts of syrian migrants on local economies in turkey. *ILR Review*, 75(3):741–768.

Chernozhukov, V., Wuthrich, K., and Zhu, Y. (2022). A t-test for synthetic controls.

Deaner, B. (2021). Proxy controls and panel data.

Ferman, B. and Pinto, C. (2021). Synthetic controls with imperfect pre-treatment fit. *Quantitative Economics*. Forthcoming.

Freyaldenhoven, S., Hansen, C., and Shapiro, J. M. (2019). Pre-event trends in the panel event-study design. *American Economic Review*, 109(9):3307–38.

Friedberg, R. M. (2001). The impact of mass migration on the israeli labor market. *The Quarterly Journal of Economics*, 116(4):1373–1408.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.

Gulek, A. (2023). Formal effects of informal labor supply: Evidence from the syrian refugees in turkey. Available at SSRN: https://ssrn.com/abstract=4264865.

Hall, P. and Heyde, C. (1980). *Martingale limit theory and its application*. Probability and mathematical statistics. Academic Press.

Ham, D. W. and Miratrix, L. (2022). Benefits and costs of matching prior to a difference in differences analysis when parallel trends does not hold.

Hunt, J. (1992). The impact of the 1962 repatriates from algeria on the french labor market. *ILR Review*, 45(3):556–572.

Jaeger, D. A., Ruist, J., and Stuhler, J. (2018a). Shift-share instruments and the impact of immigration. Working Paper 24285, National Bureau of Economic Research.

Jaeger, D. A., Ruist, J., and Stuhler, J. (2018b). Shift-share instruments and the impact of immigration. Technical report, National Bureau of Economic Research.

Lebow, J. (2022). The labor market effects of venezuelan migration to colombia: reconciling conflicting results. *IZA Journal of Development and Migration*, 13(1).

Miao, W., Geng, Z., and Tchetgen Tchetgen, E. (2018). Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*.

Peri, G. and Yasenov, V. (2019). The labor market effects of a refugee wave synthetic control method meets the mariel boatlift. *Journal of Human Resources*, 54(2):267–309.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–22.

Rudelson, M. and Vershynin, R. (2007). Sampling from large matrices: An approach through geometric functional analysis. *J. ACM*, 54(4):21–es.

Stephens Jr, M. and Yang, D.-Y. (2014). Compulsory education and the benefits of schooling. *American Economic Review*, 104(6):1777–1792.

Turkish Red Crescent and WFP (2019). Refugees in turkey: Livelihoods survey findings. https://reliefweb.int/report/turkey/refugees-turkey-livelihoods-survey-findings-2019-entr.

UNHCR (2021). Unhcr refugee statistics. https://www.unhcr.org/refugee-statistics/. Accessed: 2022-06-15.

Vives-i-Bastida, J. (2022). Predictor selection for synthetic controls.

Wolfers, J. (2006). Did unilateral divorce laws raise divorce rates? a reconciliation and new results. *American Economic Review*, 96(5):1802–1820.

A.1. Theory

Throughout the appendix we introduce the following notation to refer to the debiased quantities: $\tilde{\epsilon}_{it} \equiv \epsilon_{it} - \sum_{j\neq i} \hat{w}_{ij}^{SC} \epsilon_{jt}$, as well as dropping the 'SC' weight subscript for notational convenience. Furthermore, we use T to mean T_1. The appendix consists of the following sections:

1. Bound on $\sum_{it} \tilde{Z}_{it}\tilde{\epsilon}_{it}$.
2. Proof of Theorem 1.
3. Proof of Theorem 2.
4. Proof of Theorem 3.
5. First stage debiasing discussion.
6. Proof of Theorem 4.
7. Additional simulation table.

1.1. Bound on $\sum_{it} \tilde{Z}_{it}\tilde{\epsilon}_{it}$

Lemma A.1 *[Bound on $\tilde{Z}_{it}\tilde{\epsilon}_{it}$] Under A1-A3, for any $\delta > 0$,*

$$\mathbb{P}\left(\left|\sum_{it} \tilde{Z}_{it}\tilde{\epsilon}_{it}\right| \geq \delta\right) \lesssim 2\exp\left(-\frac{\delta^2}{2c_z^2 JT\sigma_\epsilon^2}\right).$$

Hence, as $JT \to \infty$, $\frac{1}{JT}\sum_{it} \tilde{Z}_{it}\tilde{\epsilon}_{it} \xrightarrow{p} 0$.

Proof. First we show that the term has zero expectation given Assumption 3 and independence of the error terms ϵ_{it}. The argument follows by noting that the SC weights depend only on ϵ_{it} for $t \leq T_0$,

$$w_j^{SC} \in \text{argmin}_{w\in\Delta^{J-1}} \|Y_j^{T_0} - Y_{-j}^{T_0}{}'w\|^2,$$

as only data from the pre-treatment period is used, here denoted as the $J+1 \times T_0$ matrix Y^{T_0}. Therefore, $w_j^{SC} \perp \epsilon_{it}$ for $t > T_0$. Recall that by the law of iterated expectations if random variables b is independent of z and a such that $\mathbb{E}[b|c] = 0$ a.e., then $\mathbb{E}[ab|z] = 0$. Using this fact, under A2 it follows that $\mathbb{E}[\epsilon_{it} w_{ij}^{SC}|Z_{it}] = 0$ for $t > T_0$. Similarly, for any injective function $h : \text{Supp}(w) \to \mathbb{R}$ it follows that $h\left(w_j^{SC}\right) \perp \epsilon_{it}$ for $t > T_0$ and, consequently, $\mathbb{E}[\epsilon_{it} h\left(w_{ij}^{SC}\right)|Z_{it}] = 0$. To apply these facts, we re-write the second term, dropping the 'SC' subscript for convenience

$$\begin{aligned}\mathbb{E}\left[\tilde{Z}_{it}\left(\epsilon_{it}-\sum_{j\neq i}\hat{w}^{SC}_{ij}\epsilon_{jt}\right)\right] &= \mathbb{E}\left[(Z_{it}-Z'_{-it}w_i)(\epsilon_{it}-\epsilon'_{-it}w_i)\right] \\ &= \mathbb{E}\left[Z_{it}\mathbb{E}[\epsilon_{it}-\epsilon'_{-it}w_i|Z_{it}]-Z'_{-it}\mathbb{E}[w_i(\epsilon_{it}-\epsilon'_{-it}w_i)|Z_{it}]\right],\end{aligned}$$

where the $-i$ subscripts denote $J\times 1$ vectors not including unit i and w_i denote the $J\times 1$ vector of weights for unit i. Given that $\mathbb{E}[\epsilon_{it}w_{ij}|Z_{it}]=0$ and $\mathbb{E}[\epsilon_{it}w^2_{ij}|Z_{it}]=0$, it follows that both conditional expectation terms are zero. Next, consider the following upper bound for the term of interest given by $Z_{it}-Z'_{-it}w\leq c_z$ for all $Z_{it}\in\mathcal{Z}$ when $w\in\Delta^{J-1}$,

$$\left|\sum_{it}\tilde{Z}_{it}\tilde{\epsilon}_{it}\right|\leq c_z\left|\sum_{it}\tilde{\epsilon}_{it}\right|.$$

Given that ϵ_{it} are subgaussian and for $t>T_0$ independent of w_i, it follows that $\epsilon_{it}-\sum_{j\neq i}\hat{w}^{SC}_{ij}\epsilon_{jt}$ is a linear combination of subgaussian random variables. The first term has variance proxy σ^2 and the second term has variance proxy $\sigma^2\|w_i\|^2\leq\sigma^2$ as for weights in the simplex $\|w_i\|^2\leq 1$. Therefore, $\tilde{\epsilon}_{it}$ is subgaussian with variance proxy $2\sigma^2$. The result then follows directly by Hoeffding's inequality for subgaussian random variables (Theorem 2.6.2 Vershynin 2018)

$$\begin{aligned}\mathbb{P}\left(\left|\sum_{it}\tilde{Z}_{it}\tilde{\epsilon}_{it}\right|\geq\delta\right) &\leq \mathbb{P}\left(\left|\sum_{it}\tilde{\epsilon}_{it}\right|\geq\delta/c_z\right) \\ &\lesssim 2\exp\left(-\frac{\delta^2}{2c_z^2JT\sigma^2_\epsilon}\right).\end{aligned}$$

□

1.2. Proof of Theorem 1

Proof. We start by re-writing the factor structure in terms of the outcome variable and in the pre-treatment period

$$\tilde{\mu}'_iF_t=\tilde{Y}_{it}(0)-\tilde{\epsilon}_{it}.$$

Using the projection trick, we can rewrite $\tilde{\mu}_i$ in terms of pre-treatment quantities:

$$\tilde{\mu}_i=(F_{T_0}F'_{T_0})^{-1}F_{T_0}(\tilde{Y}^{T_0}_i+\tilde{\epsilon}_{it}).$$

With this in mind, consider the object of interest

$$\begin{aligned}\left|\sum_{it} \tilde{Z}_{it}\tilde{\mu}_i' F_t\right| &= \left|\sum_{it} \tilde{Z}_{it} F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}(\tilde{Y}_i^{T_0} + \tilde{\epsilon}_{it})\right| \\ &\leq \sum_{it} |\tilde{Z}_{it} F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{\epsilon}_{it}| + \sum_{it} |\tilde{Z}_{it} F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{Y}_i^{T_0}| \\ &\leq c_z \left(\sum_{it} |F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{\epsilon}_{it}| + \sum_{it} |F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{Y}_i^{T_0}|\right).\end{aligned}$$

Where the inequalities follow from the triangle inequality and the bound for $\tilde{Z}_{it}$. For the first term bound we proceed as in Abadie and Zhao (2022) and apply the Cauchy-Schwarz inequality and the eigenvalue bound on the Rayleigh quotient to bound the factor terms for any t, s

$$(F_t'(F_{T_0}F_{T_0}')^{-1}F_s)^2 \leq \left(\frac{\bar{F}^2 k}{T_0 \underline{\xi}}\right)^2,$$

To bound these terms in expectation observe that $\bar{\epsilon}_{it} \equiv F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\epsilon_{iT_0}$ is a linear combination of subgaussian random variables and therefore it is itself a subgaussian random variable with variance proxy $\left(\frac{\bar{F}^2 k}{T_0 \underline{\xi}}\right)^2 \frac{\sigma_\epsilon^2}{T_0}$. Therefore,

$$\begin{aligned}|\mathbb{E}[(\bar{\epsilon}_{iT_0} - \bar{\epsilon}_{-iT_0}' w_i)]| &\leq \mathbb{E}[|(\bar{\epsilon}_{iT_0} - \bar{\epsilon}_{-iT_0}' w_i)|] \\ &\leq \mathbb{E}\left[\sum_j |\bar{\epsilon}_{iT_0}|\right] \\ &\leq \left(\mathbb{E}\left[\sum_j |\bar{\epsilon}_{iT_0}|^2\right]\right)^2 \\ &= \left(\sum_j \mathbb{E}\left[|\bar{\epsilon}_{iT_0}|^2\right]\right)^2 \\ &\leq 2\left(\frac{\bar{F}^2 k}{\underline{\xi}}\right)\sqrt{\frac{J}{T_0}}\sigma_\epsilon.\end{aligned}$$

The first inequality follows from Jensen's inequality. The second inequality follows by the triangle inequality and the absolute value and expectation operator inequality. The third follows from Holder's inequality with $q = 2$ and Jensen's inequality. Finally, the last inequality follows from Rigollet and Hutter 2019 (Lemma 1.4) which bounds absolute moments of sub-gaussian random variables. It follows that

$$
\begin{aligned}
\mathbb{E}[|\tilde{\mu}_i' F_t|] &= \mathbb{E}[|F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}(\epsilon_{iT_0} - \epsilon_{-iT_0}' w_i)|] \\
&= \mathbb{E}[|(\bar{\epsilon}_{iT_0} - \bar{\epsilon}_{-iT_0}' w_i)|] \\
&\leq 2\left(\frac{\bar{F}^2 k}{\underline{\xi}}\right)\sqrt{\frac{J}{T_0}}\sigma_\epsilon.
\end{aligned}
$$

For the second term observe that

$$
\begin{aligned}
\sum_{it} |F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{Y}_i^{T_0}| &\leq \left(\frac{\bar{F}^2 k}{\underline{\eta} T_0}\right)\sum_{it} |\sum_{t<T_0} \tilde{Y}_{it}| \\
&= \frac{T}{T_0}\left(\frac{\bar{F}^2 k}{\underline{\eta}}\right)\sum_{i,t<T_0} |\tilde{Y}_{it}|
\end{aligned}
$$

Dividing by TJ it follows that the term is bounded by

$$
\left(\frac{\bar{F}^2 k}{\underline{\eta}}\right)\frac{1}{JT_0}\sum_{i,t<T_0} |\tilde{Y}_{it}| = \left(\frac{\bar{F}^2 k}{\underline{\eta}}\right) MAD(\tilde{Y}^{T_0}).
$$

The bound then follows from the proof of Theorem 1 and by Jensen's inequality applied to the absolute value. Consistency follows by an application of Markov's inequality. □

1.3. Proof of Theorem 2

Proof. The proof follows the proof of Theorem 1 by bounding $\mathbb{E}\left[\frac{1}{JT_0}\sum_{i,t\leq T_0} |\tilde{Y}_{jt}|\right]$. It is useful to re-write the SC problem in matrix form. Let W be the $J \times J$ matrix of weights where each row sums to 1 and $diag(W) = 0$. Then the $J \times T_0$ matrix $\tilde{Y}^{T_0}$ can be re-written as $Y^{T_0} - \hat{W}Y^{T_0}$. It follows that the Frobenius norm of the matrix $\|\tilde{Y}^{T_0}\|_F^2 = \sum_{it<T_0} \tilde{Y}_{it}^2$ is bounded.

$$
\begin{aligned}
\|\tilde{Y}^{T_0}\|_F^2 = \|Y^{T_0} - \hat{W}Y^{T_0}\|_F^2 &\leq \|Y^{T_0}\|_F^2 + \|\hat{W}Y^{T_0}\|_F^2 \\
&\leq \|Y^{T_0}\|_F^2 + \|\hat{W}\|_F^2\|Y^{T_0}\|_F^2 \\
&\leq \|Y^{T_0}\|_F^2(1+J) \\
&\leq \bar{r}\bar{\sigma}_1(1+J).
\end{aligned}
$$

where the first inequality follows by the triangle inequality. The second by the bound on the frobenius norm of a matrix product. The third by noting that each row of W sums to 1 and $W_{ij} \in [0,1]$ and so $\|w_i\|^2 \leq 1$ which implies $\|\hat{W}\|_F^2 \leq J$. Finally, the last inequality follows from A4 as $\|Y^{T_0}\|_F^2 \leq \|Y^{T_0}\|_2^2\bar{r}$. Next, observe that $\sum_{it<T_0} |Y_{it}| = \|vec(Y^{T_0})\|_1^2$ and $\|vec(Y^{T_0})\|_2^2 = \|Y^{T_0}\|_F^2$.

So by the inequality between l_1 and l_2 norms,

$$\sum_{it<T_0} |Y_{it}| = \|vec(Y^{T_0})\|_1^2 \leq \sqrt{JT_0}\|vec(Y^{T_0})\|_2^2 = \sqrt{JT_0}\|Y^{T_0}\|_F^2.$$

Given the previous derivations we get the following bound

$$\frac{1}{JT_0}\sum_{i,t\leq T_0} |\tilde{Y}_{jt}| \leq \frac{\sqrt{JT_0}}{JT_0}\bar{r}\bar{\sigma}_1(1+J) = \bar{r}\bar{\sigma}_1\left(\frac{1}{\sqrt{JT_0}} + \sqrt{\frac{J}{T_0}}\right).$$

The result follows by noting that for a bounded random variable $|X| \leq c$, $\mathbb{E}|X| \leq c$. □

1.4. Proof of Theorem 3

Proof. Under A1-A3 Lemma A.1 and Theorem 3 show that both $\frac{1}{JT}\sum_{it}\tilde{Z}_{it}\tilde{\epsilon}_{it}$ and $\frac{1}{JT}\sum_{it}\tilde{Z}_{it}\tilde{\mu}_i'F_t$ are $o_p(1)$. It remains to be shown that the first stage term $\sum_{it}\tilde{R}_{it}\tilde{Z}_{it}$ is $O_p(1)$. Then the consistency results follows given that $O_p(1)o_p(1) = o_p(1)$. Observe that

$$\begin{aligned}
\sum_{it}\tilde{R}_{it}\tilde{Z}_{it} &= \sum_{it}(\gamma\tilde{Z}_{it} + \tilde{\eta}_{it})\tilde{Z}_{it} \\
&= \gamma\sum_{it}\tilde{Z}_{it}^2 + \sum_{it}\tilde{Z}_{it}\tilde{\eta}_{it} \\
&= \gamma\sum_{it}\tilde{Z}_{it}^2 + o_p(1) \\
&= \gamma\sum_{it}\tilde{Z}_{it}^2.
\end{aligned}$$

Furthermore, since the diameter of Z_{it} is bounded we can show that

$$\frac{1}{TJ}\mathbb{E}|\sum_{it}\tilde{Z}_{it}^2| \leq \frac{1}{TJ}\mathbb{E}|\sum_{it}c_z^2| = c_z^2.$$

It follows that the first stage term is $O_p(1)$. Indeed, for any $\delta > 0$ there exists an $M = c_z^2/\delta$ such that

$$P(\frac{1}{TJ}|\sum_{it}\tilde{Z}_{it}^2| > \delta) \leq \frac{\frac{1}{TJ}\mathbb{E}|\sum_{it}\tilde{Z}_{it}^2|}{\delta} \leq c_z^2/\delta = M.$$

where the inequality follows by Markov's inequality and if c_z depends on JT then the statement holds for large enough JT. The inconsistency of the TSLS estimator follows from the simple example in the main text, by noting that $\frac{1}{JT}\sum_{it}Z_{it}\mu_i'F_t$ is not $o_p(1)$ under our assumptions. □

1.5. First stage debiasing?

As noted in the main text debiasing the first stage is not a necessary condition for the identification of θ. In this section we note that it may have implications for the first stage strength.

Note that under A1-A3, $corr(Z_{it}, \mu_i' F_t) < 1$, then $cov(Z, \mu' F) \leq var(Z) var(\mu' F_t)$. In the proof of Theorem 3 we consider the case in which we debias the instrument; in the case in which we do not debias the instrument we get a similar derivation.

$$
\begin{aligned}
\sum_{it} \tilde{R}_{it} Z_{it} &= \sum_{it} (\gamma \tilde{Z}_{it} + \tilde{\eta}_{it}) Z_{it} \\
&= \gamma \sum_{it} \tilde{Z}_{it} Z_{it} + \sum_{it} Z_{it} \tilde{\eta}_{it} \\
&= \gamma \sum_{it} \tilde{Z}_{it} Z_{it} + o_p(1).
\end{aligned}
$$

Suppose that our instrument follows a particular structure: $Z_{it} = A_{it} + \mu_i' F_t$ such that A_{it} is mean zero $i.i.d$ and $A_{it} \perp \mu_i' F_t$ and $\mathbb{E}[A_{it}^2] = \sigma_A^2$. Then, it can be shown that if $\tilde{\mu}_i' F_t$ is $o_p(1)$ then

$$
\begin{aligned}
\frac{1}{JT} \sum_{it} \tilde{R}_{it} \tilde{Z}_{it} &\xrightarrow{p} \gamma \tilde{\xi}, \\
\frac{1}{JT} \sum_{it} \tilde{R}_{it} Z_{it} &\xrightarrow{p} \gamma \sigma_A^2,
\end{aligned}
$$

where $\sigma_A^2 \leq \tilde{\xi} \leq 2\sigma_A^2$. To see this observe that

$$
\begin{aligned}
\frac{1}{JT} \sum_{it} \tilde{R}_{it} \tilde{Z}_{it} &= \gamma \frac{1}{JT} \sum_{it} \tilde{A}_{it}^2 + o_p(1), \\
\frac{1}{JT} \sum_{it} \tilde{R}_{it} Z_{it} &= \gamma \frac{1}{JT} \sum_{it} \tilde{A}_{it} A_{it} + o_p(1),
\end{aligned}
$$

and,

$$
\begin{aligned}
\mathbb{E}\left[\frac{1}{JT} \sum_{it} \tilde{A}_{it} A_{it}\right] &= \frac{1}{JT} \sum_{it} \mathbb{E}[A_{it}^2] = \sigma_A^2, \\
\mathbb{E}\left[\frac{1}{JT} \sum_{it} \tilde{A}_{it}^2\right] &= \frac{1}{JT} \sum_{it} \mathbb{E}[A_{it}^2](1 + \|w_i\|_2^2) = \sigma_A^2 \frac{1}{JT} \sum_{it} (1 + \|w_i\|_2^2).
\end{aligned}
$$

So the result follows under an appropriate LLN and noting that $1/(J-1) \leq \|w_i\|_2^2 \leq 1$, given that the weights are in the simplex. It follows that under this simple instrument structure, debiasing the instrument will lead to a weakly larger first stage than not debiasing it. The intuition for this result is that in high correlation cases debiasing only one side (Y but not Z) leads to $\tilde{Y}$ being mostly uncorrelated with Z in the case in which the debiasing is successful in removing most of the correlation. This intuition is confirmed in the simulation exercise.

1.6. Proof of Theorem 4

Proof. The additional regularity condition we consider is $z^{min}(z^{min}+1)-z_{max}^2 > 0$, where z_{min}, z_{max} denote the supremum and infimum of the instrument set. We are interested in the following quantity for a given set of weights w:

$$\frac{\sqrt{JT}(\tilde{\theta}^{SW}-\theta)}{v_{JT}^w} = \left(\frac{1}{JT}\sum_{it}\tilde{Z}_{it}\tilde{R}_{it}\right)^{-1}\frac{1}{v_{JT}^w\sqrt{JT}}\sum_{it}\tilde{Z}_{it}\left(\mu_i - \sum_{j\neq i}\hat{w}_{ij}^{SC}\mu_j\right)' F_t$$
$$+\left(\frac{1}{JT}\sum_{it}\tilde{Z}_{it}\tilde{R}_{it}\right)^{-1}\frac{1}{v_{JT}^w\sqrt{JT}}\sum_{it}\tilde{Z}_{it}\left(\epsilon_{it} - \sum_{j\neq i}\hat{w}_{ij}^{SC}\epsilon_{jt}\right),$$

where the conditional variance is given by $v_{JT}^w = \sum_{it} var(\tilde{Z}_{it}\tilde{\epsilon}_{it} \mid Z, w)$.

First, we show that the bias term is $o_p(1)$. Note that from the proof of the consistency theorem that under A1-A3 we have that

$$\sum_{it}|F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{Y}_i^{T_0}| \leq \left(\frac{\bar{F}^2 k}{\underline{\eta} T_0}\right)\sum_{it}|\sum_{t<T_0}\tilde{Y}_{it}|$$
$$= \frac{T}{T_0}\left(\frac{\bar{F}^2 k}{\underline{\eta}}\right)\sum_{i,t<T_0}|\tilde{Y}_{it}|$$

So, dividing by $\sqrt{JT}$ and using the bound on the pre-treatment mean absolute deviation

$$\frac{1}{\sqrt{JT}}\sum_{it}|F_t'(F_{T_0}F_{T_0}')^{-1}F_{T_0}\tilde{Y}_i^{T_0}| \leq \frac{\sqrt{T}}{T_0\sqrt{J}}\left(\frac{\bar{F}^2 k}{\underline{\eta}}\right)\sum_{i,t<T_0}|\tilde{Y}_{it}|$$
$$\leq \frac{\sqrt{T}}{T_0\sqrt{J}}\left(\frac{\bar{F}^2 k}{\underline{\eta}}\right)\sqrt{JT_0}\bar{r}\bar{\sigma}_1(1+J)$$
$$= \sqrt{\frac{T}{T_0}}(1+J)\bar{r}\bar{\sigma}_1.$$

Therefore, the first term is $o_p(1)$ when $\sqrt{\frac{T}{T_0}}(1+J)\bar{r}\bar{\sigma}_1 \to 0$.

To show the second term is bounded in probability observe that conditional on Z and w

$$\begin{aligned} var\left(\sum_{it}\tilde{Z}_{it}\tilde{\epsilon}_{it}\right) &= \mathbb{E}\left(\sum_{it}\tilde{Z}_{it}\tilde{\epsilon}_{it}\right)^2 \\ &\leq c_z^2\mathbb{E}\left(\sum_{it}\tilde{\epsilon}_{it}\right)^2 \\ &= c_z^2\sigma^2 T(J + \sum_i \|w_i\|^2 + 2\sum_{i<j}\bar{W}_{ij}), \end{aligned}$$

where $\bar{W}_i = \sum_{k\neq i} w_{jk}w_{ik} - w_{ij} - w_{ji}$. Similarly, we can derive that

$$\sum_{it} var(\tilde{Z}_{it}\tilde{\epsilon}_{it} \mid Z, w) = \sigma^2\sum_{it}(1 + \|w_i\|^2)(Z_{it}^2 + \sum_{j\neq i} Z_{jt}w_{ij} - \sum_{i\neq j} Z_{jt}Z_{it}w_{ij}).$$

We consider a martingale representation as Abadie and Imbens (2012) do for matching estimators. Define the partial sums for a given time t as

$$S_{tJk} = \sum_{l=1}^{k}\tilde{Z}_{lt}\tilde{\epsilon}_{lt},$$

under our assumptions it follows that

$$\mathbb{E}[S_{tJk+1} \mid S_{tJ1}, \ldots, S_{tJk}] = \mathbb{E}[\tilde{Z}_{lt}\tilde{\epsilon}_{lt} \mid S_{tJ1}, \ldots, S_{tJk}] + S_{tJk} = S_{tJk},$$

where the condition expectation is zero given that conditional on the weights w under our error independence and partial instrument validity assumptions the instrument and the error term are uncorrelated when $t > T_0$ as shown in Lemma 1. Furthermore, define the martingale difference as

$$X_{tJk} = S_{tJk} - S_{tJk-1} = \tilde{Z}_{kt}\tilde{\epsilon}_{kt},$$

and the information set is given by the generated σ-algebra $\mathcal{F}_{tJk} = \sigma(\{Y_1^{T_0}, \ldots, Y_{k-1}^{T_0}, Z_{1t}, \ldots, Z_{k-1t}\})$ as the weights depend only on the outcome values in the pre-treatment period.

We want to apply the martingale CLT (Theorem 3.2, p. 59 from Hall and Heyde (1980)):

Theorem A.1 *[Martingale CLT] Let $\{S_{ni}, \mathcal{F}_{ni}, 1 \leq i \leq k_n, n \geq 1\}$ be a zero-mean, square-integrable martingale array with differences X_{ni} and let η^2 be an a.s. finite random variable. Suppose (1) a*

Lindeberg condition, for all $\varepsilon > 0$:

$$\sum_i E\left(X_{ni}^2 \mathbf{1}\{|X_{ni}| > \varepsilon\} |\mathcal{F}_{n,i-1}\right) \xrightarrow{p} 0,$$

(2):

$$V_{nk_n}^2 = \sum_i E\left(X_{ni}^2 |\mathcal{F}_{n,i-1}\right) \xrightarrow{p} \eta^2,$$

and (3) the σ-fields are nested $\mathcal{F}_{n,i} \subset \mathcal{F}_{n+1,i}$ for $1 \leq i \leq k_n, n \geq 1$. Then:

$$S_{nk_n} = \sum_i X_{ni} \xrightarrow{d} Z,$$

where the random variable Z has characteristic function $E \exp\left(-\frac{1}{2}\eta^2 t^2\right)$.

Condition (3) is easy to check given our definition of $\mathcal{F}_{tJk}$. We start with condition (1) and consider $\frac{1}{\sqrt{JT}} X_{tJk}$. We will show point-wise convergence. Note that conditional on $\mathcal{F}_{tJk}$ by applying Holders inequality,

$$\mathbb{E}\left[\frac{X_{tJk}^2}{JT} \mathbf{1}\{|X_{tJk}| > \varepsilon\}\right] \leq \frac{1}{JT} \mathbb{E}\left[X_{tJk}^4\right]^{1/2} P\left(\left\{|X_{tJk}| > \sqrt{JT}\varepsilon\right\}\right)^{1/2}.$$

The second term can be further bounded by applying Chebyshev's inequality and under A1-A3

$$\begin{aligned} P\left(\left\{|X_{tJk}| > \sqrt{JT}\varepsilon\right\}\right) &\leq \frac{var(\tilde{\epsilon}_{kt}|w) c_z^2}{JT\varepsilon^2} \\ &\leq \frac{\sigma^2 c_z^2}{T\varepsilon^2}, \end{aligned}$$

where the conditional variance is bounded above by the sum of the J variances. The first expectation term can be bounded by noting that the instruments are bounded and under the assumption of bounded fourth moments of the error term

$$\begin{aligned} \mathbb{E}\left[X_{tJk}^4\right] &\leq c_z^4 \mathbb{E}\left[\tilde{\epsilon}_{kt}^4\right] \\ &\leq c_z^4 \mathbb{E}\left[\sum_i \epsilon_{it}^4\right] \\ &= c_z^4 J \mathbb{E}\left[\epsilon_{it}^4\right]. \end{aligned}$$

Combining the two bounds we get that for $X_{tJk}/\sqrt{JT}$,

$$\sum_{kt} E\left((X_{tJk}/\sqrt{JT})^2 \mathbf{1}\left\{|X_{tJk}/\sqrt{JT}| > \varepsilon\right\} |\mathcal{F}_{tJk-1}\right) \leq JT \frac{\sigma\sqrt{m_4} c_z^3 \sqrt{J}}{JT\sqrt{T}\varepsilon} \lesssim \sqrt{\frac{J}{T}},$$

where $\mathbb{E}\left[\epsilon_{it}^4\right] \leq m_4$. Hence, as $\sqrt{\frac{J}{T}} \to 0$ Lindeberg's condition (1) is satisfied point-wise. Next, we show that the variance term in condition (2) is bounded in probability and not $o_p(1)$. We start by noting that under our assumptions

$$E\left(X_{tJi}^2 |\mathcal{F}_{tJi-1}\right) = \sigma^2 \Delta_{it},$$

where $\Delta_{it} = (1 + \|w_i\|^2)(Z_{it}^2 + \sum_{j\neq i} Z_{jt} w_{ij} - \sum_{i\neq j} Z_{jt} Z_{it} w_{ij})$. We proceed by bounding the Δ_{it}, recall that the instruments are bounded and $1/(J-1) \leq \|w_i\|^2 \leq 1$, therefore

$$(1 + 1/(J-1))(z^{min}(z^{min} + 1) - z_{max}^2) \leq \Delta_{it} \leq 2(z^{max}(z^{max} + 1)),$$

where we assume that $z^{min}(z^{min} + 1) - z_{max}^2 > 0$. It follows that

$$\sum_{it} E\left((X_{tJi}/\sqrt{JT})^2 |\mathcal{F}_{n,i-1}\right) = \frac{\sigma^2}{JT} \sum_{it} \Delta_{it} = O(1),$$

and as $\lim_{JT\to\infty} \frac{\sigma^2}{JT} \sum_{it} \Delta_{it} = c > 0$.

Hence, all conditions for the martingale Lindeberg CLT are satisfied and as noted under our strong instrument conditions the first stage term is $o_p(1)$ therefore the normality result follows. $\square$

1.7. Additional simulation table

Table A.1: Simulations for $T_0 = 10$

	r=0.5				r=0.7				r=0.9			
	Mean	Var	Bias	MSE	Mean	Var	Bias	MSE	Mean	Var	Bias	MSE
					$\sigma = 0.5$							
OLS (TWFE)	1.29	0.02	0.29	0.10	1.48	0.02	0.48	0.25	1.71	0.01	0.71	0.52
TSLS (TWFE)	1.23	0.05	0.23	0.11	1.46	0.07	0.46	0.29	1.78	0.06	0.78	0.68
SIV	1.01	0.01	0.01	0.01	1.07	0.02	0.07	0.02	1.25	0.04	0.25	0.10
projected SIV	0.93	0.04	-0.07	0.04	0.99	0.05	-0.01	0.05	1.24	0.36	0.24	0.41
Agg. SIV	1.23	0.06	0.23	0.12	1.47	0.07	0.47	0.30	1.79	0.05	0.79	0.68
SIV + projected	0.94	0.04	-0.06	0.04	0.99	0.05	-0.01	0.05	1.24	0.35	0.24	0.41
SIV + Agg.	1.23	0.06	0.23	0.12	1.47	0.07	0.47	0.29	1.79	0.05	0.79	0.67
SIZ Z	1.05	0.02	0.05	0.02	1.16	0.03	0.16	0.05	1.49	0.03	0.49	0.27
					$\sigma = 1$							
OLS (TWFE)	1.36	0.02	0.36	0.15	1.59	0.02	0.59	0.36	1.84	0.01	0.84	0.73
TSLS (TWFE)	1.22	0.06	0.22	0.11	1.46	0.07	0.46	0.29	1.78	0.06	0.78	0.67
SIV	1.02	0.02	0.02	0.02	1.11	0.03	0.11	0.04	1.35	0.05	0.35	0.17
projected SIV	0.93	0.05	-0.07	0.06	1.00	0.07	0.00	0.07	0.66	44.22	-0.34	44.29
Agg. SIV	1.24	0.06	0.24	0.12	1.47	0.07	0.47	0.29	1.80	0.05	0.80	0.69
SIV + projected	0.93	0.05	-0.07	0.06	1.00	0.07	0.00	0.07	0.67	43.34	-0.33	43.40
SIV + Agg.	1.24	0.06	0.24	0.12	1.47	0.07	0.47	0.29	1.79	0.05	0.79	0.68
SIZ Z	1.06	0.02	0.06	0.03	1.21	0.03	0.21	0.07	1.56	0.03	0.56	0.34
					$\sigma = 2$							
OLS (TWFE)	1.47	0.02	0.47	0.24	1.73	0.02	0.73	0.56	2.03	0.02	1.03	1.08
TSLS (TWFE)	1.22	0.06	0.22	0.11	1.46	0.08	0.46	0.29	1.78	0.06	0.78	0.67
SIV	1.04	0.03	0.04	0.03	1.17	0.05	0.17	0.07	1.47	0.06	0.47	0.28
projected SIV	0.93	0.08	-0.07	0.08	1.03	0.10	0.03	0.10	1.34	0.13	0.34	0.25
Agg. SIV	1.25	0.08	0.25	0.14	1.47	0.09	0.47	0.31	1.80	0.05	0.80	0.70
SIV + projected	0.93	0.08	-0.07	0.08	1.04	0.09	0.04	0.10	1.34	0.13	0.34	0.25
SIV + Agg.	1.24	0.07	0.24	0.13	1.47	0.09	0.47	0.31	1.80	0.05	0.80	0.69
SIZ Z	1.08	0.03	0.08	0.04	1.26	0.04	0.26	0.11	1.62	0.04	0.62	0.43
					$\sigma = 4$							
OLS (TWFE)	1.62	0.04	0.62	0.42	1.94	0.04	0.94	0.92	2.30	0.03	1.30	1.72
TSLS (TWFE)	1.21	0.08	0.21	0.12	1.45	0.09	0.45	0.29	1.77	0.07	0.77	0.67
SIV	1.06	0.05	0.06	0.06	1.23	0.07	0.23	0.12	1.58	0.08	0.58	0.42
projected SIV	0.94	0.13	-0.06	0.14	1.08	0.15	0.08	0.15	1.46	0.15	0.46	0.35
Agg. SIV	1.24	0.09	0.24	0.15	1.46	0.12	0.46	0.34	1.80	0.06	0.80	0.70
SIV + projected	0.95	0.13	-0.05	0.13	1.08	0.15	0.08	0.15	1.46	0.15	0.46	0.35
SIV + Agg.	1.24	0.09	0.24	0.15	1.46	0.12	0.46	0.33	1.80	0.06	0.80	0.70
SIZ Z	1.11	0.05	0.11	0.06	1.32	0.05	0.32	0.16	1.68	0.05	0.68	0.51
					$\sigma = 8$							
OLS (TWFE)	1.82	0.06	0.82	0.74	2.23	0.07	1.23	1.58	2.67	0.07	1.67	2.87
TSLS (TWFE)	1.20	0.11	0.20	0.15	1.44	0.12	0.44	0.31	1.77	0.09	0.77	0.68
SIV	1.08	0.09	0.08	0.10	1.29	0.11	0.29	0.19	1.66	0.11	0.66	0.55
projected SIV	0.96	0.20	-0.04	0.20	1.15	0.22	0.15	0.24	1.55	0.23	0.55	0.52
Agg. SIV	1.22	0.13	0.22	0.18	1.46	0.16	0.46	0.37	1.79	0.09	0.79	0.72
SIV + projected	0.97	0.19	-0.03	0.19	1.15	0.22	0.15	0.24	1.55	0.22	0.55	0.52
SIV + Agg.	1.22	0.13	0.22	0.17	1.46	0.16	0.46	0.37	1.79	0.09	0.79	0.72
SIZ Z	1.13	0.07	0.13	0.09	1.36	0.08	0.36	0.21	1.73	0.06	0.73	0.60

A.2. Data

Turkish Statistical Institute (Turkstat) defines employment under four categories: wage-employment (60.7%), self-employment (20.3%), unpaid family worker (13.2%) and employer (5.6%). Wage-employment, or salaried employment, refers to the type of jobs that are done as an exchange for monetary or non-monetary payment. Both fixed and hourly pay are considered wage-employment under this category. The reason why we focus on salaried employment as opposed to overall employment for the empirical section of the paper is that, as suggested by Gulek (2023), wage employment and non-wage employment (self-employment, employer, or unpaid family work) are driven by different economic forces. Whereas there has to be an employer willing to hire a worker for a particular wage for that worker to have a salaried job (i.e, we can think about a labor demand curve), self-employment is an individual labor-supply decision. Natives who lose their salaried jobs due to the labor supply shock may choose to search for a salaried job while remaining unemployed, or if self-employment is a feasible alternative, may choose to remain employed. Gulek (2023) shows that transition from salaried to non-salaried jobs is an important adjustment mechanism for Turkish men but not so for Turkish women. Whereas he finds similar effects for men and women in salaried employment, he finds opposing results for non-salaried employment. He further argues that the canonical labor demand framework is more appropriate to think about wage employment (as opposed to non-wage employment) in settings where self-employment is a feasible alternative.

Table B.2: Educational Attainment of Syrian refugees and Natives

Educational Attainment	Syrian migrants (age 18+)	Natives (Age: 18-64)
No degree	0.21	0.12
Primary school	0.42	0.33
Secondary school	0.20	0.16
High school	0.10	0.20
Some college and above	0.08	0.19

Source: Author's calculation using 2019 Household Labor Force Survey for natives, and Turkish Red Crescent and WFP (2019) for the Syrian refugees.

In the main text, we write that Syrian refugees are less educated than the Turkish natives. We show evidence for this on table B.2. We use Turkish Household Labor force Surveys to determine the educational attainment of natives, and use livelihood surveys that are conducted on Syrian refugees to determine their educational attainment. According to these surveys, 21% of Syrian refugees in Turkey do not have any degree, 63% have at most a primary school degree, and 83% do not have a high school diploma, whereas these numbers are 12%, 45%, and 61%, respectively for natives.