Economic costs of the War in Donbas for the affected Ukrainian regions

Frantisek Masek[*] Renan Serenini[†]

Sapienza University of Rome, Department of Economics and Law

Abstract

Using the Synthetic Difference-in-Differences (SDiD) estimator, we show that the economic effects of the war in Donbas on the Luhansk and Donetsk oblasts between 2014 and 2019 were enormous . We estimate the average treatment effect on per capita disposable income of 3,362 USD (constant 2011) and per capita gross regional product (GRP) of 4,853 USD. These effects correspond to a decline from the counterfactual by 53% and 60%, respectively. We also show a sharp increase in the unemployment rate of 5.56 percentage points. We also estimate the effect on investment activity, highlighting one of the possible underlying mechanisms through which the fall in income and GRP may have taken place. The impact on gross fixed capital formation was over 2.5 billion USD (65% lower than it would likely have been without the impact of war). We control for possible spillover effects on neighboring regions that could lead to the SUTVA violation. Using a spatial extension of DiD, we do not find any significant spillover effects.

Keywords: Ukraine, Economics of Conflict, Average Treatment Effect, Synthetic Difference-in-Differences

JEL classification: C21, C23, F51

[*]frantisek.masek@uniroma1.it.

[†]renan.serenini@uniroma1.it.

We gratefully thank the Center for Spatial Data Science (CSDS) at the University of Chicago for hosting us during the time of writing the article. We specially want to thank Luc Anselin and Julia Koschinsky for their support and enriching discussions. We also thank all participants of the study group at the CSDS for their valuable comments, especially Johannes Moser, Pedro Amaral, and Alexander Lehner. We are grateful to Kateryna Pavlyuk for her help with data collection. Thanks are also due to the State Statistics Service of Ukraine for their e-mail responses even during the Russian invasion of their country.

1 Introduction

The Russian invasion of eastern Ukraine in 2014 imposed a tremendous cost on the region. According to estimates published before the full-scale invasion began, at the end of February 2022, the war in Donbas alone had caused over 14,000 fatalities between 2014 and the end of 2021 (UN, 2022); roughly half on the Ukrainian side. The fighting has had immense economic consequences for Ukraine, especially for the directly affected regions. We strive to estimate the effect of the 2014 invasion of Donbas on the gross regional product per capita (GRP), disposable income per capita, and the unemployment rates of Donetsk and Luhansk.[1] We also highlight one of the potential core mechanisms driving the economic slump after 2014, by showing the effect on fixed capital investment activity in Donbas. We also investigate possible spillovers to neighboring regions.

To estimate the Average Treatment Effect (ATE) on the variables of interest, we deploy the Synthetic Difference-in-Differences estimator developed by Arkhangelsky et al. (2021). We find a sharp decline in economic activity measured by gross regional product and by disposable income (both per capita). Table 1 shows, the average effect of the Russian invasion and the continuing presence of Russian troops in Donbas between 2014 and 2019 was to lower per capita income in the region by 3,362 USD (constant 2011) and the gross regional product by 4,853 USD (respectively 53% and 60% lower than would likely have been in the absence of the conflict).[2] We also identify a sharp deterioration in labor market conditions in Donbas. The ATE on the unemployment rate is 5.56 percentage points. We further aim to identify one of the potential core mechanisms of the reduced economic activity - a drop in investment activity.[3] We document a very strongly negative ATE on investments, with losses of over 2.5 billion USD - investment activity was 65% lower than it would likely have been in the absence of the war. We also stress the persistence of the effect on investment activity. While in

[1] We do not analyze effects on Crimea and Sevastopol due to absence of data since Russian authorities took over the region.

[2] We end the period before the COVID-19 pandemic.

[3] Given the tremendous uncertainty about the future of the affected regions, it seems likely that investments in Donetsk and Luhansk would be heavily attenuated.

Donetsk, gross fixed capital formation reached its 2013 level once again only in 2019, Luhansk stayed below 2013 level through the entire period.

Table 1: The average treatment effects and standard errors for each variable of interest

	Income	*GRP*	*Unemployment*	*Investments*
Estimate	-3362 (USD)	-4853 (USD)	+5.56 (pp)	-2.558 (bn. USD)
Standard error	(164)	(467)	(0.77)	(0.236)

Source and note: Based on the databases of the State Statistics Service of Ukraine. Per capita GRP, per capita disposable income, and investments are in constant 2011 USD, and the unemployment rate is in percentage points. Following Arkhangelsky et al. (2021), we use the "placebo method" to compute the standard errors.

Additionally to the core results from the SDiD analysis, we provide sensitivity checks by utilizing conventional DiD and SCM estimators. The results are consistent with those generated by SDiD. In addition to estimating the ATE for Donbas as one unit, we disentangle the effects on Donetsk and Luhansk separately. We find highly similar magnitudes of the effects across disposable income per capita and GRP per capita. The labour market deterioration is more intense in Luhansk, and the weak investment activity is stronger in Donetsk.

We acknowledge a possible violation of the stable unit treatment value assumption (SUTVA), given that we use the rest of Ukrainian regions in our control set. Therefore, we couple the estimator developed by Arkhangelsky et al. (2021) with the approach proposed by Butts (2021) to cope with possible spatial spillovers to neighboring eastern regions of Ukraine. However, we do not observe any significant spillover in any of the outcomes of interest, so we consider that our estimation is resistant to possible interference of the treated and control units.

Related literature. Bluszcz and Valente (2022) quantify the overall effect on the Ukrainian economy using the Synthetic Control approach. They find that the 2014 invasion decreased the GDP per capita of Ukraine by 15.1% on average until 2017. The authors also estimate the effect on a regional level. They find a 43% negative average causal effect for Donetsk and 52% for Luhansk. Bluszcz and Valente (2022) focus their analysis more narrowly on the effects on GDP per capita. We extend the analysis by also showing the effect on disposable income per capita and on labor market conditions, and by using a different synthetic control variant of the method. Furthermore, we study one of the possible fundamental mechanisms behind the decrease in GRP by estimating the effect on investment activity. Generally, while Bluszcz and Valente (2022) focus their attention mainly on the aggregated effect on the whole of Ukraine and only briefly touch upon regional effects, we are primarily interested in a profound investigation of the effects of the war on the Donetsk and Luhansk regions.

Similarly, McCannon (2022) studies the impact of the 2014 invasion on the aggregate level. The author finds the average negative effect on GDP per capita to be 25.7% over the period of 2014-2020. McCannon (2022) uses the Synthetic Control Method (SCM) approach from Abadie and Gardeazabal (2003), Abadie et al. (2010) and Abadie et al. (2015), as did Bluszcz and Valente (2022).

Havlik et al. (2020) provides an intriguing summary of the economic costs of the battles for the Donbas area. Their careful systematic assessment of the costs of the conflict offers valuable insights into damages caused by the Russian invasion. The study estimates the minimum costs of reconstruction of Donbas when the war finally ends at 21.7 billion USD.[4]

Recently, many articles analyzing the effects of the full-scale Russian invasion in 2022 have appeared, but our work is not related to this strand of literature. Our interest is in investigating the consequences of the invasion from 2014. Though the economic effects of war in eastern Ukraine were considerable, literature carefully estimating regional effects is lacking.

As Murdoch and Sandler (2002), Murdoch and Sandler (2004), De Groot (2010), and Dunne and Tian (2015) show, the spillover effects of war and conflicts may be substantial. The destructive economic consequences are present not solely in areas of

[4]The authors conditioned their estimates on a hypothetical near ceasefire at the time the article was published.

direct conflict, but can be of considerable size even in neighbouring areas that are not directly involved. For this reason, we analyze the possible spillover effects on neighboring regions in eastern Ukraine.

2 Data and Methodology

Database. We use a database from the State Statistics Service of Ukraine from 2003 to 2019. We estimate the effect of conflict on disposable income per capita, gross regional product per capita, gross fixed capital formation (all in constant 2011 USD), and on the unemployment rate (% of total population aged 15-70, ILO estimates). We have observations across the whole period for the first and the third variables. However, we work with a shorter dataset that begins in 2004 for GRP. To estimate the effects on the unemployment rate, we shorten the database even more beginning in 2008, due to a lack of data.

We consider 24 of the 27 Ukrainian regions.[5] We omit the city of Kyiv, given its unique economic status in the economy. We also cannot incorporate data from the regions of Crimea and Sevastopol, which were seized by Russian authorities in 2014 and have not publish any data since. Therefore, our analysis is focused only on the regions of Donbas - Luhansk and Donetsk.

Synthetic Difference-in-Differences. To estimate the causal effect of the Russian invasion in 2014, we use a new method in the comparative case study literature that is an intersection of two well-established methods - Difference-in-Differences (DiD) and Synthetic Control (SCM). By exploiting appealing features from both methods, the Synthetic Difference in Differences (SDiD) estimator developed by Arkhangelsky et al. (2021) offers a new approach that may be competitive, or possibly superior, in situations in which SCM or DiD would normally be deployed individually. The SDiD re-

[5]Specifically, Ukraine has 24 so-called oblasts, 1 autonomous republic (Crimea), and 2 cities with special status (Kyiv and Sevastopol). We call all of them regions for the sake of simplification. However, we sometimes also use the term 'region' to signify the whole Donbas when we refer to both Luhansk and Donetsk together.

weights control units in the donor pool and uses them within a two-way fixed setting to run weighted least squares. The re-weighting relaxes the reliance on the parallel trend assumption, as is the case in the SCM literature. However, SDiD is still invariant to additive unit-level shifts, given that remains in the two-way fixed effect framework. Moreover, unlike in ordinary SCM, the re-weighting is not present only with respect to units in the donor pool, but also in regards to the pre-treatment periods similarity to the post-treatment behavior of the control units.

Arkhangelsky et al. (2021) show that their estimator may be written as the following optimization problem:

$$\left(\widehat{\tau^{sdid}}, \hat{\mu}, \hat{\alpha}, \hat{\beta}\right) = \arg\min_{\mu,\alpha,\beta,\tau} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left[Y_{it} - \left(\mu + \alpha_i + \beta_t + \tau D_{it}\right) \right]^2 \hat{\omega}_i^{sdid} \hat{\lambda}_t^{sdid} \right\}, \quad (1)$$

which boils down to a two-way fixed effects regression combined with nonuniform unit weights $\hat{\omega}_i^{sdid}$ and time weights $\hat{\lambda}_t^{sdid}$. The notation is standard, α_i and β_t stand for unit and time fixed effects, respectively. D_{it} is a binary variable for having been treated and τ determines the ATE.

We do not alter the way Arkhangelsky et al. (2021) obtain $\hat{\omega}_i^{sdid}$ and $\hat{\lambda}_t^{sdid}$. Thus, in the case of the former, we have:

$$(\hat{\omega}_0, \hat{\omega}_i^{sdid}) = \arg\min_{\omega_0 \in \mathbb{R}, \omega \in \Omega} \ell_{unit}(\omega_0, \omega), \quad (2)$$

where

$$\ell_{unit}(\omega_0, \omega) = \sum_{t=1}^{T_{pre}} \left(\omega_0 + \sum_{i=1}^{N_{co}} \omega_i Y_{it} - \frac{1}{N_{tr}} \sum_{i=N_{co}+1}^{N} Y_{it} \right)^2 + \zeta^2 T_{pre} ||\omega||_2^2, \quad (3)$$

$$\Omega = \left\{ \omega \in \mathbb{R}_+^N : \sum_{i=1}^{N_{co}} \omega_i = 1, \quad \omega_i = N_{tr}^{-1} \ \forall i = N_{co}+1, ..., N \right\}, \quad (4)$$

with $\mathbb{R}_+^N$ being the positive real line, and we pin down the regularization parameter as do Arkhangelsky et al. (2021). The problem differs from the usual SCM setting with two fundamental deviations. Firstly, inserting the intercept ω_0 results in obtaining the unit weights by fitting trends instead of levels. Second, the regularization causes the synthetic unit to be less sparse, because the weights are more dispersed compared to those in ordinary synthetic control methods, which often rely on large weights of only a few units.[6]

[6] Penalization can also help to establish the uniqueness of the weights.

The time weights problem is isomorphic with one change. The only alteration is that we do not include the regularization parameter. The problem follows:

$$(\hat{\lambda}_0, \hat{\lambda}_i^{sdid}) = \arg\min_{\lambda_0 \in \mathbb{R}, \lambda \in \Lambda} \ell_{time}(\lambda_0, \lambda), \tag{5}$$

where

$$\ell_{unit}(\lambda_0, \lambda) = \sum_{i=1}^{N_{co}} \left(\lambda_0 + \sum_{t=1}^{T_{pre}} \lambda_t Y_{it} - \frac{1}{T_{post}} \sum_{t=T_{pre}+1}^{T} Y_{it} \right)^2, \tag{6}$$

$$\Lambda = \left\{ \lambda \in \mathbb{R}_+^T : \sum_{t=1}^{T_{pre}} \lambda_t = 1, \quad \lambda_t = T_{post}^{-1} \ \forall t = T_{pre} + 1, ..., T \right\}, \tag{7}$$

The idea behind time re-weighting is to match the average post-treatment outcomes for control units (up to a constant) in the pre-treatment period. Only a subset of pre-treatment periods is taken into account in the final weighted regression.

It is straightforward that, when using uniform weights for units and time periods in equation 1, we are back to the standard DiD. When the intercept in equation 6 is omitted and uniform weighting for time periods are used, the estimator boils down to SCM. Thus, SDiD can be seen as a general version of these two cases. We utilize both in the sensitivity analysis in section 3 to show that our results do not change significantly, regardless of the method we apply.

Possible SUTVA violation. Given that we use the rest of Ukrainian regions in the donor pool, there could be concerns regarding the possibility of spillovers to other regions, especially in the eastern part of the country. Therefore, we extend the estimator developed by Arkhangelsky et al. (2021) by utilizing Butts (2021) spatial DID approach. By incorporating Butts (2021) into the SDiD framework, we can detect if there are any significant spillovers to neighboring regions that might have biased the results.[7]

[7]Note that Butts (2021) method is closely related to Delgado and Florax (2015). However, the latter works with the spatial weights matrix used in the spatial econometrics literature (Anselin, 1988). To use Butts (2021) within the approach stemming from the spatial econometrics literature, that uses the spatial weights matrix, is a question of assigning different values in the matrix. If we do not row-standardized the spatial weights matrix and define the vector of indirectly (through spillovers) treated units only as a binary indicator yielding one if a given unit is within a defined distance, the outcome is equivalent to Butts (2021) and Delgado and Florax (2015). In other words, Butts (2021) ‘ring method‘ can effortlessly be nested in the "spatial weights matrix approach" from Delgado and Florax (2015). We prefer to frame

Merging the estimators from Butts (2021) and Arkhangelsky et al. (2021) results in the following minimization problem:

$$\left(\hat{\tau},\hat{\mu},\hat{\alpha},\hat{\beta},\hat{\tau}_s\right) = \arg\min_{\mu,\alpha,\beta,\tau,\tau_s} \left\{ \sum_{i=1}^{N}\sum_{t=1}^{T} \left[Y_{it} - \left(TWFE + \tau D_{it} + \tau^{s}(1-D_{it})S_{it}\right)\right]^2 \hat{\omega}_i \hat{\lambda}_t \right\}, \tag{8}$$

where $TWFE = \mu + \alpha_i + \beta_t$ holds for the sake of clarity. An additional dummy variable S_{it} indicates whether a given unit is among those suspect of receiving spillovers.[8] We are interested only in the possibility of spillovers from the directly treated regions to untreated ones.[9] Thus, we do not distinguish spillover effects from one treated unit on another one and vice versa.[10] The algorithms for obtaining $\hat{\omega}_i$ $\hat{\lambda}_t$ are not altered in any way.[11]

Heterogeneous Treatment Effects. Considering that we estimate the ATE on the whole of Donbas, which consists of two regions, Luhansk and Donetsk, there may be some degree heterogeneity of the effects on the two regions. We check for differences

---

our estimation as does Butts (2021) given that we work only with a few possible indirectly treated units. We consider Delgado and Florax (2015) to be better suited for cases with substantially higher number of units obtaining spillovers.

[8] Butts (2021) defines S_{it} based on distance from the treated units. We use neighboring regions of Luhansk and Donetsk as regions susceptible to spillovers.

[9] Note that the usual approach to suspecting possible spillovers would be to restrict the donor pool by omitting problematic units. We do this, because the donor pool includes only regions that are neither affected by direct treatment nor indirectly by spillover. We do try to quantify the possible spillovers which is not done when restriction of the control set is the sole approach used.

[10] If we were interested in estimating it, the optimization problem would change into:

$$\left(\hat{\tau},\hat{\mu},\hat{\alpha},\hat{\beta},\hat{\tau}_s\right) = \arg\min_{\mu,\alpha,\beta,\tau,\tau_s} \left\{ \sum_{i=1}^{N}\sum_{t=1}^{T} \left[Y_{it} - \left(TWFE + \tau D_{it} + \tau_0^{s}(1-D_{it})S_{it} + \tau_1^{s} D_{it}S_{it}\right)\right]^2 \hat{\omega}_i \hat{\lambda}_t \right\} \tag{9}$$

However, this is beyond the scope of our paper. We run the spatial extension only to detect possibility of spillover effects on neighboring untreated units, to validate our core results.

[11] This implies that we assume a close similarity of the directly and indirectly treated units, given that we use the same weights within the weighted regression. Although this assumption may be ambiguous for some particular situations, we believe that it is a reasonable assumption in our case, due to the fact that we allow for the possibility of spillovers only in neighboring regions, which should be similar in terms of economic characteristics. Nevertheless, we run robustness checks by obtaining the weights for both directly and indirectly treated units, and the results do not change in any important way.

in the effects on each of them in section 3 by estimating the effect for each separately, omitting the other region from the dataset.

Chaisemartin and D'Haultfoeuille (2020) and Chaisemartin and D'Haultfoeuille (2022) show that using the Two-Way Fixed Effects (TWFE) estimator in the situation of non-constant treatment effect (TE) across units or time can result in violation of the "no-sign reversal" property. Considering that the assumption of constant TEs in each group and period is highly implausible in real data examples, their ascertainment must be taken into account when working with more treated units or when adopting staggered treatment. Put differently, the TWFE regression does not necessarily identify the ATE as a convex combination of units treatment effects. Given that some units can obtain negative weights, using the mere weighted mean to compute the total ATE may result in an opposite sign. Even though our results in section 3 are in line with the expected intuition of the direction of the effect, we do formally check for the weights of all unit-time effects, utilizing the approach of Chaisemartin and D'Haultfoeuille (2020).

3 Findings

Decomposition of the treatment effect weights. We decompose weights for the specific ATEs of units that we obtain in the process of running the TWFE regression for Luhansk and Donetsk. Denoting τ_{it} as the coefficient of D_{it}, the treatment for a given unit i at time t, from equation 1, we can write the ATE as follows:

$$\mathbb{E}[\tau_{it}] = \mathbb{E}\left[\sum_{(i,t):D_{it\neq 0}} \Omega_{it} TE_{it}\right], \tag{10}$$

where TE_{it} is the treatment effect for a given unit i at time period t and Ω_{it} stands for the weights for each unit-time effect. Naturally, if some of the weights are negative, the resulting estimation of the ATE generated by the TWFE can contain a bias, and possibly even lead to the estimation with the opposite sign than the actual effects on the treated unit. We refer to Chaisemartin and D'Haultfoeuille (2020) and Chaisemartin and D'Haultfoeuille (2022) for a profound discussion of when such a situation can happen, as well as for various consequences and relationships that can result.

We apply the decomposition developed by Chaisemartin and D'Haultfoeuille (2020) for our case of two treated units to check whether they satisfy the "no-sign reversal" feature. Given that we have six post-treatment periods, we estimate twelve treatment effects, with non showing negative weights. We obtains the ATE below as a convex combination of unit-time effects. Therefore, the ATE cannot be of a different sign than these effects. We repeat the exercise for all four variables of interest, and the results hold for all of them. Consequently, we believe that the TWFE approach is not problematic from this perspective in this case.[12]

The main results. We display all the core results in figures 1 and 2. As shown in table 1, the average treatment effects are all negative. For disposable income per capita, the difference compared to the counterfactual shown in figure 1 is 3,362 USD (constant 2011), 53% lower than it would like have been in the absence of conflict. The decrease in the gross regional product per capita is 4,853 USD, which is 60% lower than the control unit. Due to the effects of the invasion, the unemployment rate is higher by 5.56 percentage points on average. The effect on gross fixed capital formation is more than 2.5 billion USD (over 65% lower than the control). Table 1 also reveals that using the "placebo method" to compute the standard errors, all results are comfortably significant using conventional significance levels.

Figures 1 and 2 show that our synthetic control units fit the pre-treatment behavior of the variables of interest. What is more, although the SDiD does not strive to fit levels, but instead only trends, the algorithm actually matches levels of three of the

[12]We would see the absence of negative weights as the expected result in our application if we used only the DiD approach, given that we work with just two treated units and six post-treatment time periods. The outcome stemming from Chaisemartin and D'Haultfoeuille (2020) and Chaisemartin and D'Haultfoeuille (2022) is that the weights will probably be positive if neither of the treated units is treated for most of the observed period, and there are no periods across which most of the units would be treated simultaneously. Our case fits this description. Formally, the weights cannot be negative if $D_{i.} + D_{.t} \leq 1$ holds. $D_{i.}$ stands for the average treatment of group i over all periods, while $D_{.t}$ denotes the average treatment at period t across all units. Hence, we decide to undertake the weights decomposition because because we believe that the SDiD estimator is possibly even more susceptible to negative weights than the standard TWFE. The reason is that usually only a low number of pre-treatment periods is taken into account during the final weighted regression in Arkhangelsky et al. (2021) estimator.

Figure 1: Disposable Income and Gross Regional Product (per capita; constant 2011 USD)

(a) Disposable Income (b) GRP

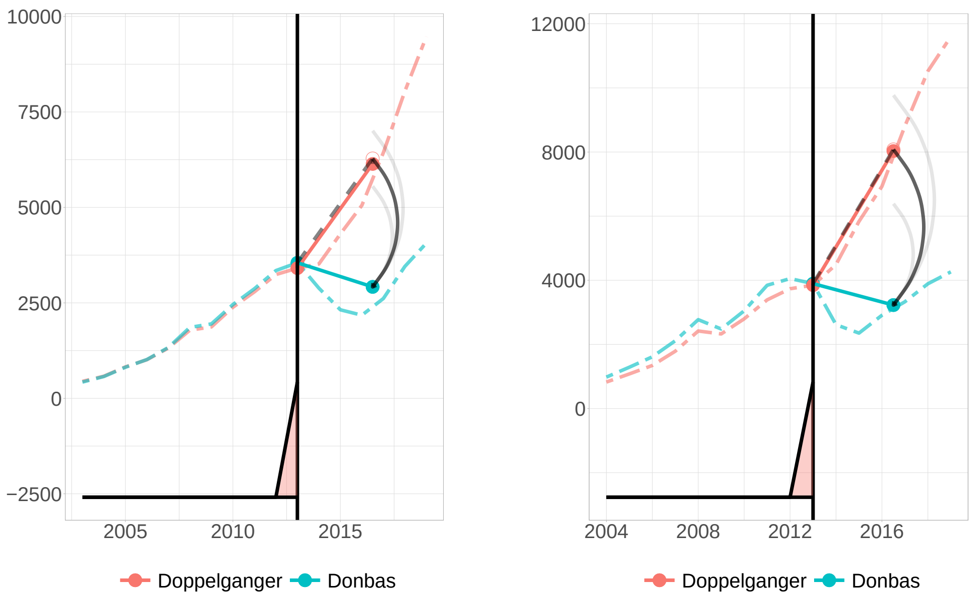

Source and note: Based on the databases of the State Statistics Service of Ukraine. The ATE on disposable income per capita and gross regional product per capita (both in constant 2011 USD). The red curve shows the synthetic control. The black curved line shows the magnitude of the ATE. The dash-dotted lines represent the observed and synthetic units, while the solid lines lead to the post-treatment averages depicted by the colored dots. The bottom shows the time weights. In these cases, the last period receives the full weight.

four variables that we are interested in. The remaining unemployment rate shows very close trends between the treated and control units, though the levels differ. Interestingly, the estimator assigns large values to the time periods just before the invasion. This is a common characteristic of the SDiD, given the absence of regularization in the time re-weighting algorithm.[13] However, we also provide robustness checks with uniform time weights. Table 9 in the appendix clearly shows that the results do not differ, aside from small alterations.

Figure 2: Unemployment Rate (%) and Investments (constant 2011 USD)

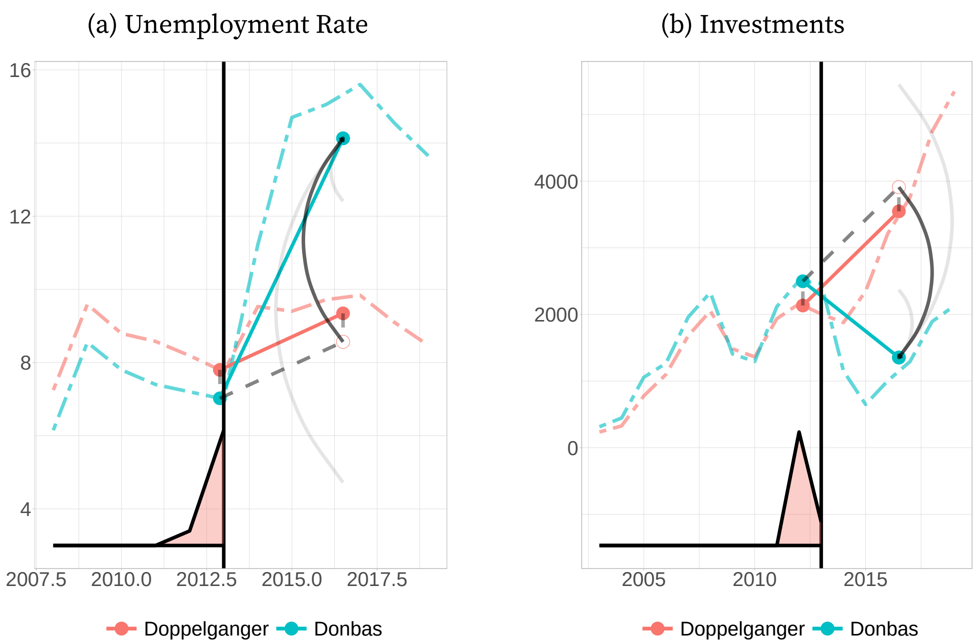

Source and note: Based on the databases of the State Statistics Service of Ukraine. The ATE on unemployment rate (%) and investments (in million constant 2011 USD)

We show region-by-region outcome difference adjusted by re-weighting in figure 4. Each dot represents the difference between the treated unit and the given control (the dot) re-scaled by the unit weight while also using the time weights. The size of the dots indicates the weight obtained in the synthetic unit algorithm, i.e., the importance

[13]Note that Arkhangelsky et al. (2021) discuss the reason. They purposefully want to allow for correlated observations within time periods for a given unit, while they want to get rid of correlation across units within a given time period. However, the decision to completely avoid the regularization in the time re-weighting may be seen as a rather discretionary one. They acknowledge this even through the fact that in a new Stata package, the regularization is also present for the time weights. Nevertheless, the strength o regularization is substantially lower than the one of the unit weights algorithm, because the main motivation is satisfying the uniqueness of the weights. See Clarke et al. (2023) for details.

of the control unit. This can be viewed as a sort of DiD estimation for each given control unit when allowing for time re-weighting.[14] If we divide re-scaled differences back by unit weights, we get DiD effects for each unit in the donor pool.[15]

However, the core idea stemming from Arkhangelsky et al. (2021) is to show that we are able to obtain the synthetic unit without relying on only a few units with large weights (no giant dots), and that the variance of the differences is not overly high compared to the ATE (middle horizontal line). The former is often the case when using SCM, while the latter appears in DiD estimator (Arkhangelsky et al., 2021).

Sensitivity to the method change. To ascertain that our results are relatively insensitive to the choice of SDiD estimator, we also run the estimation using conventional DiD and SCM. Thus, we repeat the results shown in table 1, but this time we apply all of the methods at once.

Table 2 shows that the direction of the effects does not differ regardless of the method we use, and the results remain comfortably within significant range for all methods. However, differences in the estimated effects of income, GRP, and investments are more distinct when we use DiD. While SDiD and SCM generate roughly the same effect, DiD shows smaller negative effects for all three variables. While DiD relies on the common trend assumption, both SCM and SDiD relax it given the re-weighting leading towards a sort of local approximation.[16]

The closer similarity of SDiD and SCM pertains to both estimation of the effects and to precision. As is clear from table 2, standard errors are mostly higher (with the exception of the unemployment rate) for DiD than for SDiD and SCM. Arkhangelsky et al. (2021) state that this feature may also be the result of the local characteristic of the synthetic control estimators.

We show comparisons of the results under different estimators in the ap-

[14]We also calculate results using conventional DiD and SCM methods as robustness checks. Table 2 documents that the two standard methods yield results consistent with our benchmark findings.

[15]The weights for each unit and time period for every estimation appear in the appendix.

[16]'Local' in the sense of comparing the treated unit only to a subset of more similar units from the donor pool. In the case of SDiD, the re-weighting is also present across pre-treatment time periods.

Figure 3: Region-by-region differences - Disposable Income and GRP

(a) Disposable Income (b) GRP

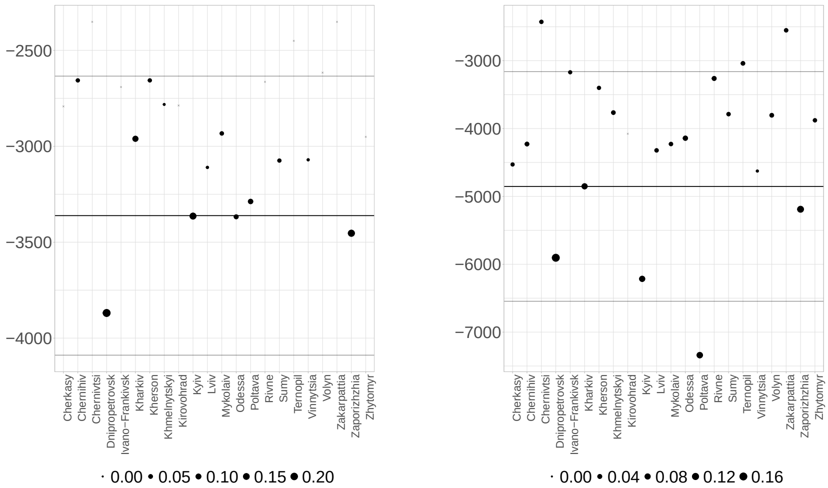

Figure 4: Region-by-region differences - Unemployment Rate and Investments

(a) Unemployment Rate (b) Investments

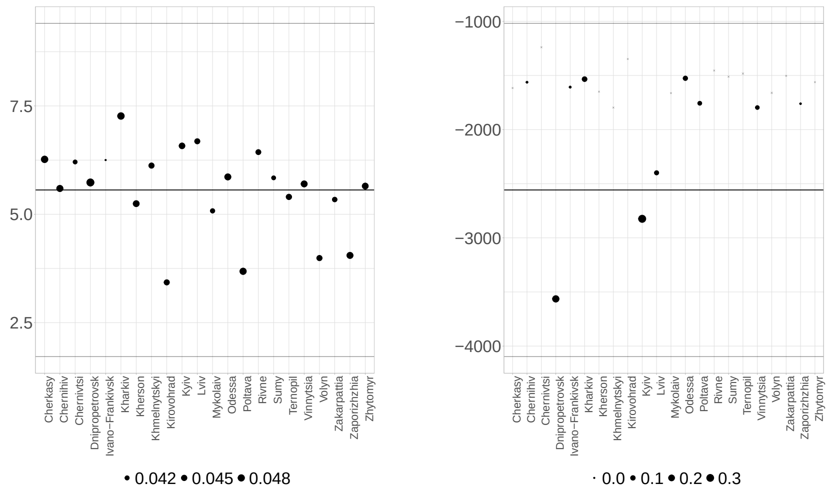

Source and note: Based on the databases of the State Statistics Service of Ukraine. Each dot indicates the difference between a given region and the treated unit (in % for unemployment rate, in USD (constant 2011) for disposable income and GRP, and in million USD (constant 2011) for investments). The size of the dots is related to the weight obtained in the synthetic control algorithm. The middle horizontal line displays the ATE.

Table 2: The average treatment effects and standard errors for each variable of interest

	Income	*GRP*	*Unemployment*	*Investments*
SDiD				
Estimate	-3362 (USD)	-4853 (USD)	+5.56 (pp)	-2.558 (bn. USD)
Standard error	(164)	(467)	(0.77)	(0.236)
SCM				
Estimate	-3497 (USD)	-5557 (USD)	+5.81 (pp)	-2.671 (bn. USD)
Standard error	(167)	(554)	(0.98)	(0.174)
DiD				
Estimate	-2569 (USD)	-4198 (USD)	+5.81 (pp)	-1.081 (bn. USD)
Standard error	(430)	(1219)	(0.86)	(0.578)

Source and note: Based on the databases of the State Statistics Service of Ukraine. Per capita GRP, per capita disposable income, and investments are in constant 2011 USD, while unemployment rate is in percentage points. Following Arkhangelsky et al. (2021), we use the "placebo method" to compute the standard errors.

The results are presented for three different methods - Synthetic Difference-in-Differences (SDiD), Synthetic Control Method (SCM), and Difference-in-Differences (DiD).

pendix (figures 6,7,8,9). We generate a separate figure for all four variables which each figure displaying the effects for DiD, SCM, and SDiD.

Separated effects. Thus far, we have worked with the ATE on the whole Donbas. In this section, we disentangle the effects for Luhansk and Donetsk separately. Our identification strategy does not differ, we only adjust the database to omit the other region and run the analysis with a single region. The results are in table 3 below.

The negative effects on income and GRP are equally dispersed across the two regions. On the other hand, the unemployment rate spikes more in Luhansk, while the

Table 3: The average treatment effects and standard errors for each variable of interest

	Income	*GRP*	*Unemployment*	*Investments*
Luhansk				
Estimate	-3416 (USD)	-4727 (USD)	+6.98 (pp)	-1.381 (bn. USD)
Standard error	(208)	(580)	(1.15)	(0.299)
Donetsk				
Estimate	-3315 (USD)	-4992 (USD)	+4.14 (pp)	-3.654 (bn. USD)
Standard error	(196)	(756)	(1.10)	(0.358)

Source and note: Based on the databases of the State Statistics Service of Ukraine. Per capita GRP, per capita disposable income, and investments are in constant 2011 USD , while the unemployment rate is in percentage points. The effects are estimated separately for the Luhansk and Donetsk regions instead of averaging them across the entire Donbas. Following Arkhangelsky et al. (2021), we use the "placebo method" to compute standard errors.

decrease in effect on investment activity is more profound in Donetsk. All the results are statistically significant using the conventional significance level, as in the case of the ATE on the whole Donbas.

The possibility of the SUTVA violation due to spillovers. Lastly, we repeat the whole estimation using the extended estimator shown in equation 8. We first define where we suspect possible spillovers. We limit our attention to the closest neighbors of the affected regions.[17] Then we conduct the whole estimation by restricting the donor pool to only pure control units, i.e., neither directly nor indirectly treated ones.

Figure 5: Map of Ukrainian Regions

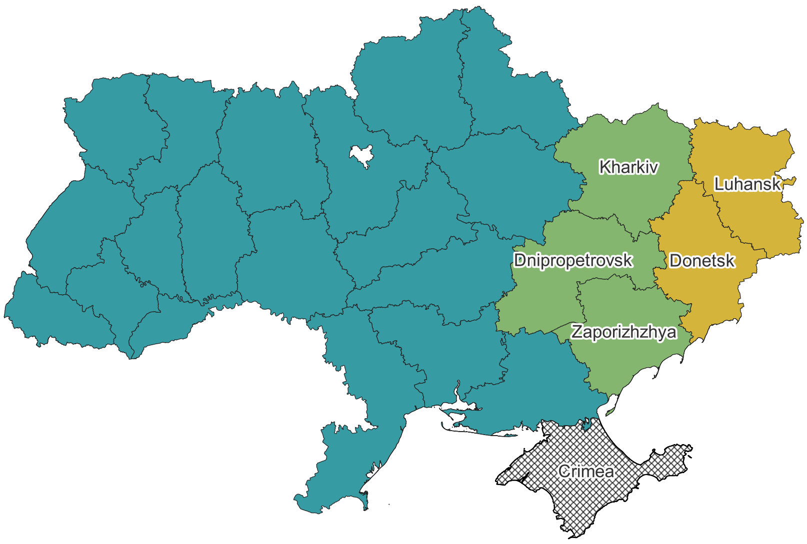

Source and note: Based on the databases of the State Statistics Service of Ukraine. The map of Ukrainian regions. The yellow color denotes directly affected areas, and the green regions are those where we suspect spillovers.

We restrict our attention to a sub-sample of five regions in eastern Ukraine (see map 5). Besides Luhansk and Donetsk, which are directly affected by the invasion, we have three neighbors where spillovers may have occurred - Kharkiv, Dnipropetrovsk, and Zaporizhzhya.[18]

[17] We also tested an extension of the possible spillover structure by including second-order neighbors, where spillovers would already have a impact relatively far. The results do not change substantially in any way. We omit this additional check for the sake of conciseness.

[18] Note in table 8 in the appendix the large values that these three units obtain in the construction of the synthetic unit. Therefore, the presence of spillovers would lead to biased results.

Table 4 below shows that the spatial extension of the SDiD estimator does not substantially change the results for the direct effect.[19][20] Further, the indirect treatment effects (spillovers) are not significant for any of the variables of interest.[21][22]

Table 4: Direct and indirect effects and standard errors for each variable of interest

	Income	*GRP*	*Unemployment*	*Investments*
Direct	-3187 (USD)	-5012 (USD)	+ 5.51 (pp)	-2.386 (bn. USD)
Standard error	(299)	(402)	(0.32)	(0.330)
Indirect (W)	240 (USD)	303 (USD)	-0.16 (pp)	-0.021 (bn. USD)
Standard error	(323)	(435)	(0.34)	(0.356)

Source and note: Based on the databases of the State Statistics Service of Ukraine. Per capita GRP, per capita disposable income, and investments are in constant 2011 USD while the unemployment rate is in percentage points. Unlike in previous estimations, we use the standard errors of the sample mean, because the presence of spillovers complicates the "placebo method" implementation.

[19]We also perform the Chaisemartin and D'Haultfoeuille (2020) decomposition check also for all estimations using the indirectly treated units presented in this subsection. The results do not show any presence of negative weights when we compute the indirect treatment effects.

[20]The unit weights obtained via the estimation are in the appendix. The time weights do not change compared to the original estimation because the donor pool remains the same.

[21]We repeat the whole estimation with the weights obtained for the average of Kharkiv, Dnipropetrovsk, and Zaporizhzhya instead of the average of Luhansk and Donetsk. The results do not differ in any important aspect. We include them in the appendix in table 10. We again add the table with the weights used in the weighted regression, as we do in the previous estimations.

[22]We are aware that our results could still be biased, due to the use of the rest of Ukraine as the control set. Even though we formally check that there are no spillovers to neighboring units (as they generally receive the highest weights in the synthetic control unit), it is possible that the effect on the whole country beyond Donbas might be present; it may simply not be more nuanced in the neighboring regions in the eastern part than in the rest of the country. However, we argue that if such a bias does exist, its magnitude is probably relatively small and it goes against the direction of the effect. Thus, our estimation can be seen as a lower bound of the effect.

4 Conclusion

We estimate the economic costs of the Russian invasion of eastern Ukraine in 2014. In the absence of data from Crimea and Sevastopol, we study the effects on the Luhansk and Donetsk regions. Our results show enormous economic hardship measured by the gross regional product and disposable income per capita. The average treatment effects on these variables from 2014 to 2019 are negative in the magnitude of 3,362 and 4,853 USD, respectively. Furthermore, we document a steep increase in the unemployment rate: the average effect on the unemployment rate in Donbas is negative 5.56 percentage points. We highlight the tremendous effect on gross fixed capital formation in the affected regions. We argue that this may be one of the key channels of the persistent economic difficulties that Luhansk and Donetsk faced during the 2014-2019 conflict. We check for the possibility of the SUTVA violation by controlling for potential spillovers of the effects to neighboring regions. We find no evidence of such spillovers.

Bibliography

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.

Anselin, L. (1988), 'Spatial econometrics: Methods and models', *Kluwer Academic, Dordrecht* .

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

Bluszcz, J. and Valente, M. (2022), 'The economic costs of hybrid wars: The case of ukraine', *Defence and Peace Economics* **33(1)**, 1–25.

Butts, K. (2021), 'Difference-in-differences estimation with spatial spillovers', *Papers 2105.03737, arXiv.org* .

Chaisemartin, C. D. and D'Haultfoeuille, X. (2020), 'Two-way fixed effects estimators with heterogeneous treatment effects', *American Economic Review* **110(9)**, 2964–2996.

Chaisemartin, C. D. and D'Haultfoeuille, X. (2022), 'Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey'.

Clarke, D., Pailañir, D., Athey, S. and Imbens, G. (2023), 'Synthetic difference in differences estimation'.

De Groot, O. J. (2010), 'The spill-over effects of conflict on economic growth in neighbouring countries in africa', *Defence and Peace Economics* **21(2)**, 149–164.

Delgado, M. S. and Florax, R. J. (2015), 'Difference-in-differences techniques for spatial data: Local autocorrelation and spatial interaction', *Economics Letters* **137**, 123–126.

Dunne, J. P. and Tian, N. (2015), 'Military expenditure, economic growth and heterogeneity', *Defence and Peace Economics* **26(1)**, 15–31.

Havlik, P., Kochnev, A. and Pindyuk, O. (2020), 'Economic challenges and costs of reintegrating the donbas region in ukraine', *The Vienna Institute for International Economic Studies: Research Report 447* .

McCannon, B. C. (2022), 'The economic costs of conflict: Ukraine, 2014-2020'.

Murdoch, J. C. and Sandler, T. (2002), 'Economic growth, civil wars, and spatial spillovers', *Journal of Conflict Resolution* **46(1)**, 91–110.

Murdoch, J. C. and Sandler, T. (2004), 'Civil wars and economic growth: Spatial dispersion', *American Journal of Political Science* **48(1)**, 138–151.

UN (2022), 'Conflict-related civilian casualties in ukraine', *Office of the United Nations High Commissioner for Human Rights* .

Appendix

Table 5: Weights of regions for every variable of interest

Region	*Income*	*GRP*	*Unempl.*	*Investm.*
Cherkasy	00.00%	02.24%	04.86%	00.00%
Chernihiv	03.48%	03.60%	04.76%	00.43%
Chernivtsi	00.00%	02.65%	04.16%	00.00%
Dnipropetrovsk	23.66%	16.42%	05.09%	24.91%
Ivano-Frankivsk	00.00%	02.16%	03.91%	00.58%
Kharkiv	11.44%	08.14%	04.88%	11.54%
Kherson	03.48%	02.36%	04.68%	00.00%
Khmelnytskyi	00.59%	03.32%	04.46%	00.00%
Kirovohrad	00.00%	00.00%	04.45%	00.00%
Kyiv	16.45%	08.33%	04.60%	32.32%
Lviv	01.07%	02.84%	04.43%	08.13%
Mykolaiv	04.08%	02.86%	04.24%	00.00%
Odessa	05.95%	05.07%	04.74%	09.15%
Poltava	07.83%	09.25%	04.80%	06.59%
Rivne	00.00%	04.14%	04.39%	00.00%
Sumy	03.45%	02.80%	04.16%	00.00%
Ternopil	00.00%	03.27%	04.49%	00.09%
Vinnytsia	01.00%	00.61%	04.71%	05.96%
Volyn	00.00%	03.47%	04.46%	00.08%
Zakarpattia	00.00%	02.84%	04.30%	00.00%
Zaporizhzhia	17.55%	11.04%	04.73%	00.21%
Zhytomyr	00.00%	02.60%	04.67%	00.00%

Source and note: Based on the databases of the State Statistics Service of Ukraine. The unit weights of synthetic control units for variable of interest.

Table 6: Weights of pre-treatment time periods

Year	*Income*	*GRP*	*Unempl.*	*Investm.*
2003	00.00%	———–	———–	00.00%
2004	00.00%	00.00%	———–	00.00%
2005	00.00%	00.00%	———–	00.00%
2006	00.00%	00.00%	———–	00.00%
2007	00.00%	00.00%	———–	00.00%
2008	00.00%	00.00%	00.00%	00.00%
2009	00.00%	00.00%	00.00%	00.00%
2010	00.00%	00.00%	00.00%	00.00%
2011	00.00%	00.00%	00.00%	00.00%
2012	00.00%	00.00%	11.24%	82.94%
2013	100.00%	100.00%	88.76%	17.06%

Source and note: Based on the databases of the State Statistics Service of Ukraine. The pre-treatment time weights of synthetic control units for each variable of interest.

Figure 6: Comparison of SDiD to DiD and SCM - income per capita (in constant 2011 USD)

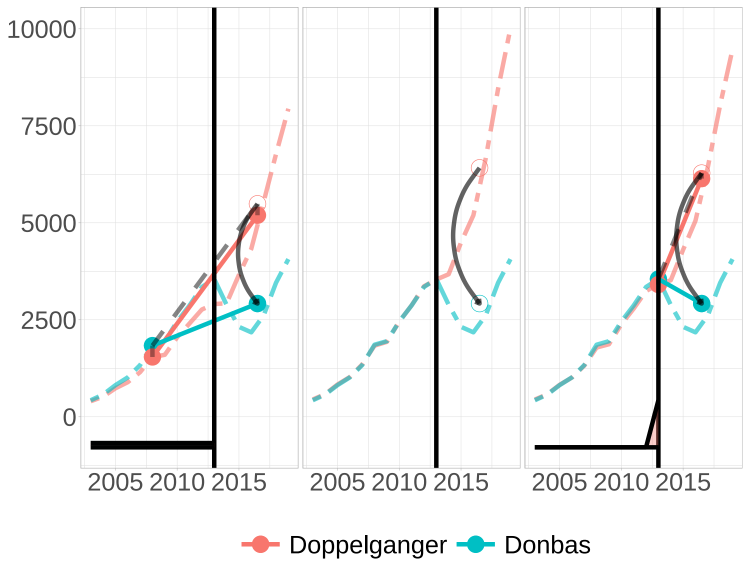

Source and note: Based on the databases of the State Statistics Service of Ukraine. The ATE on disposable income per capita (in constant USD 2011) using different identification strategies. Starting from the left, we have Difference-in-Differences, Synthetic Control Method, and Synthetic-Difference-in-Differences.

Figure 7: Comparison of SDiD to DiD and SCM - GRP per capita (in constant 2011 USD)

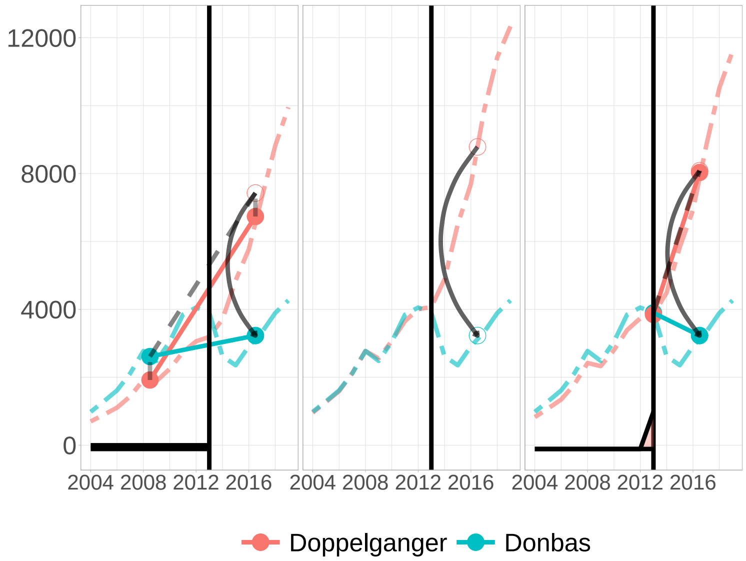

Source and note: Based on the databases of the State Statistics Service of Ukraine. The ATE on GRP per capita (in constant 2011 USD) using different identification strategies. Starting from the left, we have Difference-in-Differences, Synthetic Control Method, and Synthetic-Difference-in-Differences.

Figure 8: Comparison of SDiD to DiD and SCM - unemployment rate (in %)

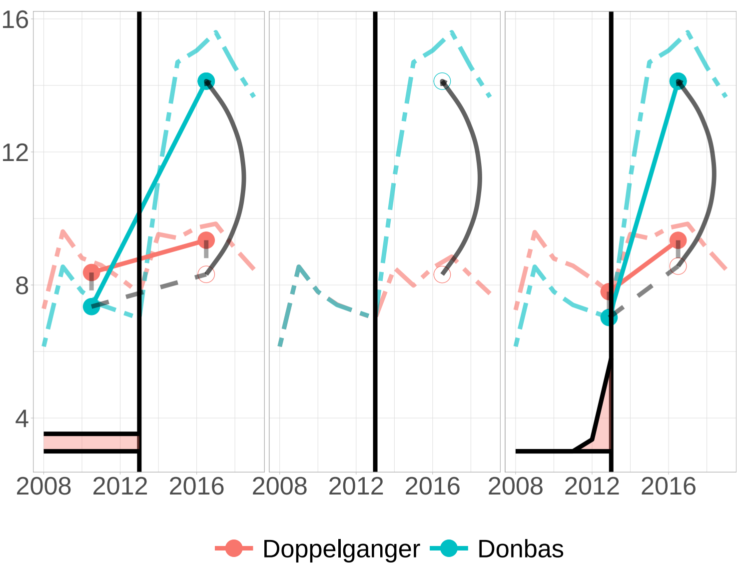

Source and note: Based on the databases of the State Statistics Service of Ukraine. The ATE on unemployment rate (in %) using different identification strategies. Starting from the left, we have Difference-in-Differences, Synthetic Control Method, and Synthetic-Difference-in-Differences.

Table 7: Weights for the spatial extension

Region	*Income*	*GRP*	*Unempl.*	*Investm.*
Cherkasy	00.00%	04.38%	05.65%	00.00%
Chernihiv	07.07%	04.71%	05.67%	00.00%
Chernivtsi	00.00%	01.40%	04.81%	00.00%
Ivano-Frankivsk	00.00%	00.99%	04.60%	00.00%
Kherson	07.07%	02.61%	05.50%	00.00%
Khmelnytskyi	00.13%	03.43%	05.25%	00.00%
Kirovohrad	00.00%	00.00%	05.18%	00.00%
Kyiv	38.53%	18.59%	05.35%	51.69%
Lviv	00.91%	03.89%	05.17%	16.11%
Mykolaiv	08.70%	06.53%	04.90%	00.00%
Odessa	12.42%	14.45%	05.59%	16.02%
Poltava	17.08%	20.06%	05.69%	09.61%
Rivne	00.00%	04.81%	05.07%	00.00%
Sumy	06.74%	04.34%	04.84%	00.00%
Ternopil	00.00%	02.05%	05.28%	00.00%
Vinnytsia	01.35%	00.00%	05.61%	06.57%
Volyn	00.00%	03.98%	05.17%	00.00%
Zakarpattia	00.00%	01.55%	05.18%	00.00%
Zhytomyr	00.00%	02.20%	05.51%	00.00%

Source and note: Based on the databases of the State Statistics Service of Ukraine. The unit weights of synthetic control units for each variable of interest for the spatial extension estimation.

Table 8: Weights using the indirectly treated units to construct the synthetic control

Region	*Income*	*GRP*	*Unempl.*	*Investm.*
Cherkasy	00.00%	06.88%	05.41%	02.95%
Chernihiv	07.63%	02.22%	06.09%	00.83%
Chernivtsi	00.00%	00.14%	04.49%	00.00%
Ivano-Frankivsk	00.00%	00.99%	04.69%	04.54%
Kherson	07.64%	00.59%	05.54%	01.78%
Khmelnytskyi	01.82%	00.00%	05.31%	01.67%
Kirovohrad	00.00%	01.73%	05.00%	02.94%
Kyiv	34.06%	24.88%	05.12%	25.73%
Lviv	02.31%	03.51%	05.02%	18.35%
Mykolaiv	09.13%	10.14%	04.56%	03.58%
Odessa	11.38%	21.94%	05.76%	16.58%
Poltava	15.72%	23.87%	05.98%	09.68%
Rivne	00.00%	00.00%	04.69%	00.31%
Sumy	07.27%	04.09%	04.64%	01.61%
Ternopil	00.00%	00.00%	05.37%	00.97%
Vinnytsia	03.03%	00.00%	06.01%	05.26%
Volyn	00.00%	00.00%	04.89%	00.79%
Zakarpattia	00.00%	00.00%	05.81%	01.49%
Zhytomyr	00.00%	00.00%	05.64%	00.95%

Source and note: Based on the databases of the State Statistics Service of Ukraine. The unit weights of synthetic control units for each variable of interest when we construct the synthetic unit using the indirectly treated units as robustness check for the spatial extension estimation.

Figure 9: Comparison of SDiD to DiD and SCM - investments (in constant 2011 USD)

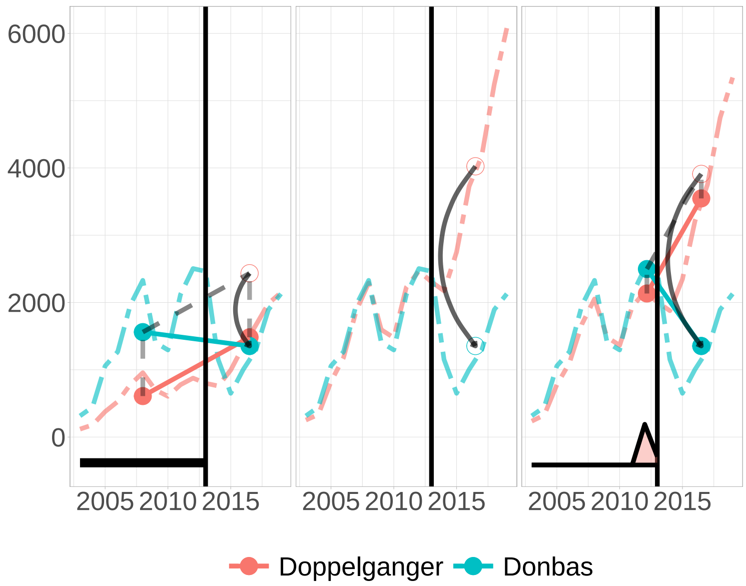

Source and note: Based on the databases of the State Statistics Service of Ukraine. The ATE on investments (in million USD constant 2011) using different identification strategies. Starting from the left, we have Difference-in-Differences, Synthetic Control Method, and Synthetic-Difference-in-Differences.

Table 9: The ATEs and SEs for each variable of interest using uniform time weights

	Income	*GRP*	*Unemployment*	*Investments*
Estimate	-3268 (USD)	-5058 (USD)	+5.79 (pp)	-2.379 (bn. USD)
Standard error	(197)	(519)	(0.81)	(0.221)

Source and note: Based on the databases of the State Statistics Service of Ukraine. Per capita GRP, per capita disposable income, and investments are in constant 2011 USD while the unemployment rate is in percentage points. The uniform time weights are used, i.e., the reweighting is deployed only at unit levels. Following Arkhangelsky et al. (2021), we use the "placebo method" to compute the standard errors.

Table 10: Direct and indirect effects for the weights from indirectly treated regions

	Income	*GRP*	*Unemployment*	*Investments*
Direct	-3159 (USD)	-5443 (USD)	+ 5.50 (pp)	-2.090 (bn. USD)
Standard error	(297)	(478)	(0.32)	(0.210)
Indirect (W)	269 (USD)	-128 (USD)	-0.17 (pp)	0.274 (bn. USD)
Standard error	(321)	(517)	(0.34)	(0.223)

Source and note: Based on the databases of the State Statistics Service of Ukraine. Per capita GRP, per capita disposable income, and investments are in constant 2011 USD while the unemployment rate is in percentage points. Unlike in the previous cases, we use the standard errors of the sample mean because the presence of spillovers complicates the "placebo method" implementation.