

Voting with their Money: Brexit and Outward Investment by UK Firms

Technical Appendix

Holger Breinlich, Elsa Leromain,
Dennis Novy and Thomas Sampson

Voting with their Money: Brexit and Outward Investment by UK Firms - Technical Appendix

Holger Breinlich Elsa Leromain Dennis Novy Thomas Sampson

February 2019

This technical appendix contains additional details on the synthetic control method and the data used in our briefing note, as well as additional robustness checks and results.

1 The Synthetic Control Method

1.1 Determination of Control Group Weights

The synthetic control method (SCM) provides a systematic way to choose comparison units in comparative case studies (Abadie and Gardeazabal, 2003; Abadie, Diamond and Hainmueller, 2010 and 2015). In our case, we are interested in constructing a counterfactual outcome for the UK-EU27 FDI flows that would have taken place in the absence of a leave vote in the Brexit referendum of 2016. The difficulty is that we cannot simply assume that some other bilateral FDI flow, or a simple average of other flows, would provide a good approximation to the counterfactual UK-EU27 flows.

The SCM proposes instead to use a weighted average of other FDI flows, with the weights chosen such that the resulting synthetic control resembles UK-EU27 FDI flows in the pre-referendum period as closely as possible, in a sense to be defined more precisely below. If a number of potentially suitable control group flows are available, as is the case in our setting, the SCM has the additional advantage that researchers do not have to make ad hoc decisions which of these flows to use; instead, the SCM provides a procedure that reduces discretion in the choice of control group by “letting the data speak”.

We now provide a technical description of the SCM, following the exposition in Abadie, Diamond and Hainmueller (2010) and Ferman, Pinto and Possebom (2018). Assume that we observe data for $J + 1$ units (here: country pairs) for T time periods (here: quarters). Unit 1 will be affected by an intervention (here: the outcome of the Brexit referendum) that is in force from period $T_0 + 1$ until period T . The remaining flows $j = 2, \dots, J + 1$ are not affected by the intervention and form the so-called donor pool from which the synthetic control will be constructed.

Let Y_{1t}^N be the outcome (here: the count of FDI transactions) that would be observed for unit 1 in the absence of the intervention and Y_{1t}^I the outcome in its presence. The effect of the intervention in period t is then measured by $\alpha_{1t} = Y_{1t}^I - Y_{1t}^N$. Of course, for the post-intervention period, $t \geq T_0 + 1$, we observe $Y_{1t}^I = Y_{1t}^N$ but not the counterfactual non-intervention

outcome (Y_{1t}^N). The goal of the synthetic control method is to construct an estimate for this counterfactual outcome as a weighted average of the outcomes (Y_{jt}) of the non-treated units:

$$\hat{Y}_{1t}^N = \sum_{j=2}^{J+1} \hat{w}_j Y_{jt},$$

where $\hat{w}_j \geq 0$ for $j = 2, \dots, J+1$ and $\sum_{j=2}^{J+1} \hat{w}_j = 1$. The weights \hat{w}_j are obtained as the solution to the following minimisation problem:

$$\hat{\mathbf{W}} \left(\hat{\mathbf{V}} \right) = \arg \min_{\mathbf{W} \in \mathcal{W}} (\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W})' \hat{\mathbf{V}} (\mathbf{X}_1 - \mathbf{X}_0 \mathbf{W}), \quad (1)$$

where \mathcal{W} is the set of all possible combinations of weights $\mathbf{W} = (w_2, \dots, w_{J+1})'$, \mathbf{X}_1 is an $F \times 1$ vector of pre-treatment observations of the treated unit and \mathbf{X}_0 is an $F \times J$ matrix of the corresponding observations for the donor pool. Note that \mathbf{X}_0 and \mathbf{X}_1 can include pre-intervention outcomes of the variable of interest (i.e., Y_{jt} for $t \leq T_0$) as well as other predictors of Y_{jt} . Thus, the approach underlying the SCM is to choose weights to minimise pre-intervention differences (in terms of FDI counts and additional determinants of these counts) between the treated unit and the synthetic control. Abadie, Diamond and Hainmueller (2010) show that if the synthetic control can match \mathbf{X}_1 , it provides a valid counterfactual for Y_{jt}^N in the sense that $\hat{Y}_{1,t}^N - Y_{jt}^N$ will be close to 0 for all $t \geq T_0 + 1$.¹

The weighting matrix $\hat{\mathbf{V}}$ in (1) is determined by minimising the distance between pre-treatment outcomes of unit 1 and the synthetic control:

$$\hat{\mathbf{V}} \left(\hat{\mathbf{W}} \right) = \arg \min_{\mathbf{V} \in \mathcal{V}} \left(\mathbf{Y}_1 - \mathbf{Y}_0 \hat{\mathbf{W}} \left(\mathbf{V} \right) \right)' \left(\mathbf{Y}_1 - \mathbf{Y}_0 \hat{\mathbf{W}} \left(\mathbf{V} \right) \right), \quad (2)$$

where \mathcal{V} is the set of diagonal positive semidefinite matrices of dimension $F \times F$. The SCM algorithm iterates between (1) and (2) until convergence is achieved.

In practice, however, a simpler and faster method of choosing $\hat{\mathbf{V}}$ often yields essentially identical results (see Kaul et al., 2018). For every $t \leq T_0$, this method regresses Y_{jt} ($j = 1, \dots, J+1$) on all predictors contained in X_t , yielding regression coefficients $\hat{\beta}_{1t}, \dots, \hat{\beta}_{Ft}$. The diagonal elements of $\hat{\mathbf{V}}$ corresponding to each predictor are then simply given by

$$\hat{v}_f = \frac{\sum_t \left(\hat{\beta}_{ft} \right)^2}{\sum_{k=1}^F \sum_t \left(\hat{\beta}_{kt} \right)^2}.$$

Intuitively, both approaches give more weight to variables with greater predictive power for the outcome of interest, Y_t . For computational reasons, we use the faster regression-based method

¹See Abadie et al. (2010, p.495) for details. This result is derived under the assumption that Y_{it}^N is given by a factor model, $Y_{it}^N = \delta_t + \theta_t Z_i + \lambda_t \mu_i + \varepsilon_{it}$, where Z_i is a vector of observed covariates, μ_i are time-invariant unobserved determinants of Y_{it}^N and ε_{it} are unobserved transitory shocks with mean zero. Intuitively, if the number of pre-intervention periods is large relative to the scale of the transitory shocks, the only way the synthetic control can match pre-intervention outcomes as well as the additional covariates is by fitting Z_i and μ_i exactly which in turn guarantees that $Y_{it}^N - \hat{Y}_{it}^N$ is close to zero. Note that the above data generating process generalises the traditional difference-in-difference model by allowing the effect of the unobserved confounders μ_i to vary with time.

to obtain most of our results, although we have checked that using the full nested procedure yields essentially identical estimates for control group weights.²

In our baseline specification, we try to match all quarterly UK-EU27 pre-referendum FDI counts since 2010. We start in 2010 only because the global financial crisis of 2008/2009 was associated with very substantial fluctuations in FDI activity, so that it is doubtful whether the data generating process for Y_{jt}^N remained stable over that time period.³

Given that we attempt to match the entire path of pre-intervention outcomes, both algorithms outlined above will give zero weight to additional co-variates and we do not include any in our baseline specification.⁴ As discussed by Ferman, Pinto and Possebom (2018), this reduces issues with specification searching among a large set of potential co-variates. Using pre-intervention outcomes only may also improve the SCM’s ability to capture unobserved determinants of FDI flows, albeit at the cost of potentially omitting relevant co-variates (Kaul et al., 2018). Even if such co-variates are omitted, however, Botosaru and Ferman (forthcoming) show that the synthetic control estimator will not necessarily be biased.⁵ Indeed, we show in our robustness checks that using a less-than-complete series of pre-intervention outcomes together with standard gravity predictors of FDI flows (bilateral distance and GDPs of the origin and destination countries) yields very similar results.

1.2 Statistical Inference

Abadie, Diamond and Hainmueller (2010) also propose a way of evaluating the statistical significance of the estimated treatment effect, $\hat{\alpha}_{1t} = Y_{1t}^I - \hat{Y}_{1t}^N$, based on the classic framework for permutation inference (see Abadie and Cattaneo, 2018). The idea is to sequentially reassign treatment to all units j in the donor pool and construct a new synthetic control in each case using all remaining units in that pool as well as the originally treated unit. For all $j = 2, \dots, J + 1$, we can then compute the corresponding treatment effects, $\hat{\alpha}_{jt} = Y_{jt}^I - \hat{Y}_{jt}^N$. Intuitively, this exercise allows us to examine whether or not the estimated effect of the Brexit referendum is large relative to the distribution of the effects estimated for the FDI flows not affected by the vote.

Given our estimates of all $\hat{\alpha}_{jt}$, we can evaluate statistical significance by computing a p-value associated with the Brexit referendum effect, $\hat{\alpha}_{1t}$. For this, we first compute the ratio of mean squared prediction errors in the post-intervention period relative to the pre-intervention period

²Computational concerns only play a role for the computation of significance levels using permutation methods (see below). It is here that we exclusively use the regression-based method. The estimated weights for the original treated flow (UK-EU27) are essentially identical for all our results, irrespective of whether we use the nested or the regression-based approach.

³In principle, our data allows us to go back to 2003. In practice, using the period 2003-2018 instead of 2010-2018 only leads to minor differences in the estimated treatment effects and significance levels (see below).

⁴See Kaul et al. (2018) for a formal proof. Intuitively, weights are chosen to match the pre-intervention path of the outcome variable of interest and the outcome at time t is of course fully explained by the outcome itself. So if the full set of pre-intervention outcomes is included in \mathbf{X}_0 , all additional co-variates will be assigned zero weight.

⁵Unbiasedness in this case requires an extension of the regularity conditions in Abadie, Diamond and Hainmueller (2010) from the unobserved to the observed determinants of Y_{it}^N .

for each of the $J + 1$ units:

$$R_j = \frac{RMSP E_{j,post}}{RMSP E_{j,pre}} = \frac{\sum_{t=T_0+1}^T (Y_{jt} - \hat{Y}_{jt}^N)^2 / (T - T_0)}{\sum_{t=1}^{T_0} (Y_{jt} - \hat{Y}_{jt}^N)^2 / T_0}.$$

We can then calculate a p-value by comparing the value of this statistic for unit 1 (R_1) to that of all other units:

$$p_1 = \frac{\sum_{j=1}^{J+1} 1(RMSP E_j \geq RMSP E_1)}{J + 1},$$

where $1(\cdot)$ denotes the indicator function.⁶

Using this procedure, we compute p-values for all SCM figures in our policy brief (Figures 2, 4, 5, 6a and 6b). With the exception of Figures 4 and 6a, all treatment effects are statistically significant at the 10% level (i.e., $p \leq 0.1$).⁷

2 Data Description

We use data on the count of greenfield and M&A transactions between the UK and the EU27 as well as between other non-EU OECD and EU27 countries. Greenfield investments are taken from the Financial Times' fDi Markets database and M&A transactions from Bureau van Dijk's Zephyr database. We briefly describe each database in turn.

2.1 fDi Markets

The fDi Markets database has been tracking cross-border greenfield investment since 2003, covering all sectors and countries worldwide. Our baseline specification using data from 2010 to 2018 contains around 59,500 greenfield investments between EU and non-EU OECD countries (96,000 for the full period 2003 to 2018).

fDi Markets obtains data on new greenfield transactions by searching over 8,000 information sources (newspapers, magazines, industry associations, company websites) in 23 languages on a daily basis. Each news article is then checked on the investing company's website, which also allows fDi Markets to gather additional information on the company and further details of the FDI project in question.

Whenever possible, fDi Markets also collects information on the capital investment and jobs associated with FDI projects based on announcements by the investing company. In practice, however, this information is only available for around a quarter of projects and has to be estimated by fDi Markets based on similar investments for the remaining cases. Even if data on jobs and capital expenditure are released, they are usually based on plans rather than realised outcomes. These are the principal reasons why we use count data throughout our analysis. In addition, job and capital investment data can be dominated by a small number of very large

⁶As Abadie, Diamond and Hainmueller (2010) discuss, this approach produces classical randomisation inference if the intervention is indeed randomly allocated across units. If this is not the case, the approach is best interpreted as a series of placebo checks that examine whether the estimated treatment effect is large compared to the placebo effects for other flows that we would not expect to be affected by the referendum.

⁷The exact values are $p = 0.032$ (Figure 2), $p = 0.056$ (Figure 5) and $p = 0.091$ (Figure 6b).

transactions, leading to much noisier time series than for counts. We do, however, make use of the available information about transaction values to compute a rough estimate of the changes in aggregate FDI investment caused by the referendum result (see below).

2.2 Zephyr

Zephyr is a database of deal information containing data on M&A, IPO, private equity and venture capital deals. It contains information on over 1,600,000 deals with up to 100,000 additional deals being added each year. Data on new transactions are obtained by searching a wide range of news publications, company press releases, stock exchange announcements, advisor submissions and websites in over 30 languages.

For our analysis, we focus on cross-border mergers and acquisitions between EU and non-EU OECD countries, yielding approximately 40,000 transactions for the period 2010-2018 (70,000 for the full period 2003-2018). For comparison with the greenfield investment data, we associate each transaction with its announcement date, although using completion dates yields very similar results.

Zephyr also provides information on deal values although this information is only available for around 40% of transactions in our sample. Similar to the greenfield data, aggregate bilateral M&A deal values are often dominated by a small number of large deals, so we again prefer to focus on counts of the number of deals.

2.3 Baseline Sample

For our baseline analysis of UK outward FDI transactions in the EU27, we include all bilateral FDI series between OECD and EU countries in the donor pool, aggregating all EU countries other than the UK into one aggregate group (EU27). We exclude all pairs that involve the UK from the donor pool since those series are potentially directly affected by the EU referendum and, therefore, would not be suitable for constructing the synthetic control. Finally, we drop all country pairs with five or fewer transactions over the entire period. This is because including too many units in the donor pool can lead to overfitting by matching the treated unit to idiosyncratic variation of a large number of control units (see Abadie, Diamond and Hainmueller, 2015). However, we show below that our results are not affected by including such pairs. With these restrictions, we end up with 124 country pairs in the donor pool. As discussed above, we focus on the period 2010-2018 for our main analysis although we show below that extending our sample to 2003-2018 does not affect our conclusions.

When we calculate the synthetic control for UK-EU27 FDI, the algorithm chooses the following bilateral series (with corresponding weights in parentheses): Switzerland-EU27 (52.1%), US-EU27 (37.6%), Japan-Mexico (8.9%), EU27-Switzerland (1.3%) and EU27-EU27 (0.1%). Thus, the most important series to construct the synthetic control is bilateral FDI from Switzerland into the EU27, followed by FDI from the US into the EU27. This makes intuitive sense as both Switzerland and the US – similar to the UK – have a close economic relationship with the EU and are important origin countries for FDI into the EU27.

3 Value of Additional FDI Outflows

To estimate the value of the additional outward FDI flows from the UK to the EU27 caused by the referendum, we use data on transaction values from the fDi Markets and Zephyr databases. From fDi Markets, we calculate that the mean capital investment value of UK-EU27 greenfield transactions in 2017 was \$19.0 million. This corresponds to £14.8 million based on the average 2017 exchange rate of 0.78 £/\$. From Zephyr, we obtain a mean value of €101.6 million for UK-EU27 M&A transactions in 2017, or £89.4 million based on the average 2017 exchange rate of 0.88 £/€.⁸

To use this information on mean values per transaction, we apply the synthetic control method to estimate the effect of the referendum on UK outward FDI to the EU27 separately for greenfield and M&A transactions. We find that the leave vote resulted in 160 additional greenfield transactions and 66 additional M&A transactions (cumulatively by 2018Q3).⁹ We then multiply these additional transactions by their respective mean values. This yields a total increase in FDI outflows from the UK to the EU27 due to the referendum of £8.3 billion by 2018Q3.

We can use a similar procedure to compute “lost investment” due to the reduction in FDI flows into the UK from the EU27 (Figure 5 in the policy brief). Here, however, we only find a significant impact of the referendum on greenfield investment, not on M&A transactions. Hence, we only need data on the average capital investment of greenfield EU27-UK transactions in 2017. We calculate a mean capital investment value of £25.8 million. Based on our estimate of 137 fewer greenfield transactions due to the referendum, this yields an estimated reduction in aggregate EU27-UK FDI flows of £3.5 billion by 2018Q3.

4 Robustness Checks and Additional Results

As our first set of robustness checks, we add to the donor pool bilateral pairs that involve the UK (Figure A.1) or pairs with five or fewer transactions (Figure A.2). The results are very similar to the baseline estimates from Figure 2 of the policy brief. Figure A.3 extends our sample period to include all quarters from 2003Q1 to 2018Q3. Again, results are essentially identical to the baseline.

In Figure A.4, we evaluate the robustness of the baseline results to using additional covariates to calculate the synthetic control. Specifically, we now include bilateral distance and the GDPs of the origin and destination countries as additional predictors for FDI flows. A simple regression of $\ln(FDIcounts)$ on the logs of these variables yields an R-squared of 71%, demonstrating their

⁸fDi Markets imputes value information for transactions where investment expenditure is not reported, using information from similar transactions. Likewise, we manually impute missing value information in Zephyr by using the mean value of other M&A transactions from the same country pair, year and 2-digit NACE industry code. If no such transactions are available, we successively widen the imputation comparison group to i) the same country pair and industry, ii) the same country pair, and iii) the same industry only. The effect of this imputation is to lower the mean transaction value compared to our raw data. We believe that this approach is superior to only using directly observable data. This is because value information tends to be more readily available for larger projects, implying the directly observable data is likely to overestimate average project value.

⁹See the robustness checks section below. We note that the estimated treatment effects for the M&A and greenfield sub-samples are statistically significant at the 10% and 5% level, respectively.

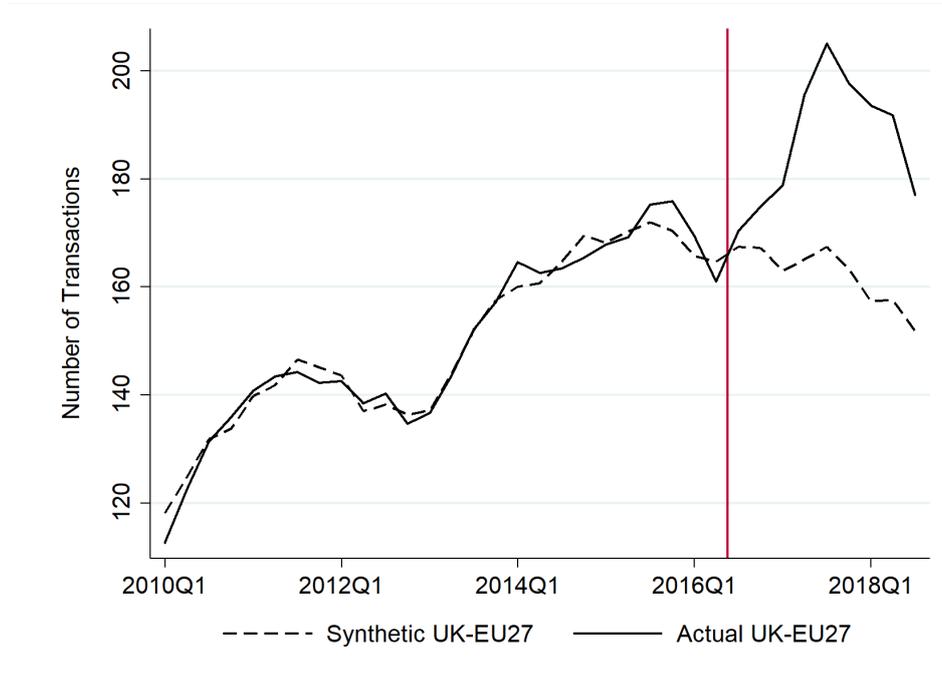
potential for explaining bilateral FDI counts. As discussed, for these additional co-variates to be given positive weights by the SCM algorithm, we have to exclude some pre-intervention outcomes. Figure A.4 plots a number of synthetic controls based on using only every second, fourth, eighth and sixteenth pre-intervention outcome, respectively. As seen, the trajectories of these additional synthetic controls look similar to before, yielding an estimated treatment effect very close to our baseline.

Finally, we run our baseline specification separately for M&A and greenfield transactions. As shown in Figures A.5 and A.6, we find a significant effect for both types of FDI flows although the impact on greenfield flows is more pronounced (a 20 percent increase compared to 10 percent for M&As).

References

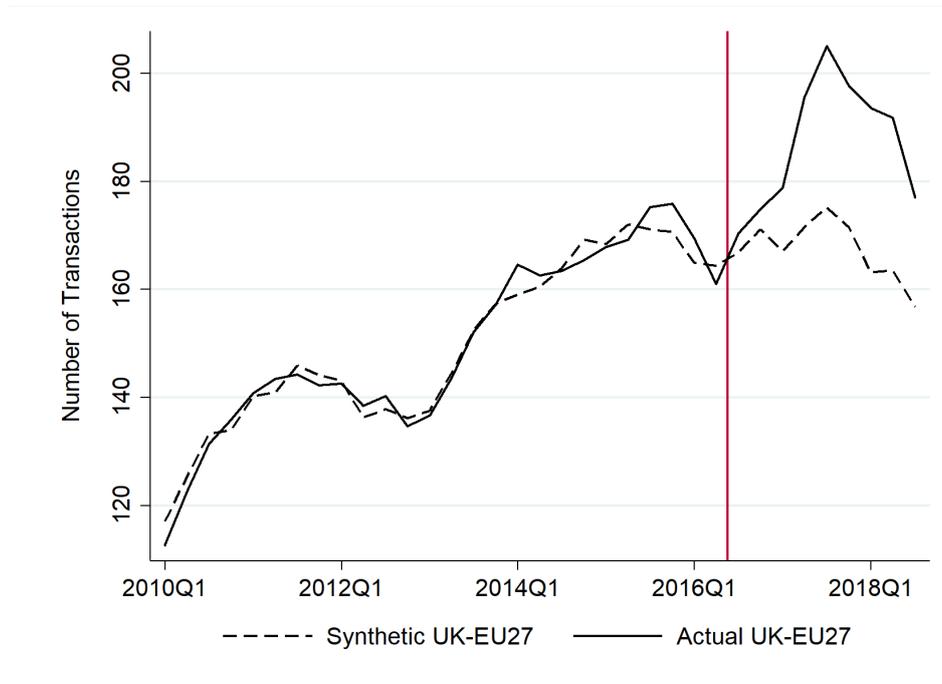
- [1] Abadie, Alberto, and Matias Cattaneo. 2018. "Econometric Methods for Program Evaluation." *Annual Review of Economics* 10:465-503.
- [2] Abadie, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *American Economic Review* 93(1): 113-32.
- [3] Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 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.
- [4] Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science*, 59(2): 495-510.
- [5] Botosaru, Irene, and Bruno Ferman. 2018. "On the Role of Covariates in the Synthetic Control Method." Forthcoming, *The Econometrics Journal*.
- [6] Ferman, Bruno, Cristine Pinto, and Vitor Possebom. 2017. "Cherry Picking with Synthetic Controls." MPRA Paper No. 85138.
- [7] Kaul, Ashok, Stefan Klößner, Gregor Pfeifer, and Manuel Schieler 2015. "Synthetic Control Methods: Never Use All Pre-Intervention Outcomes Together With Covariates." Working Paper, Saarland University.

Figure A.1: UK-EU27 FDI counts (actual vs synthetic control), full set of pairs



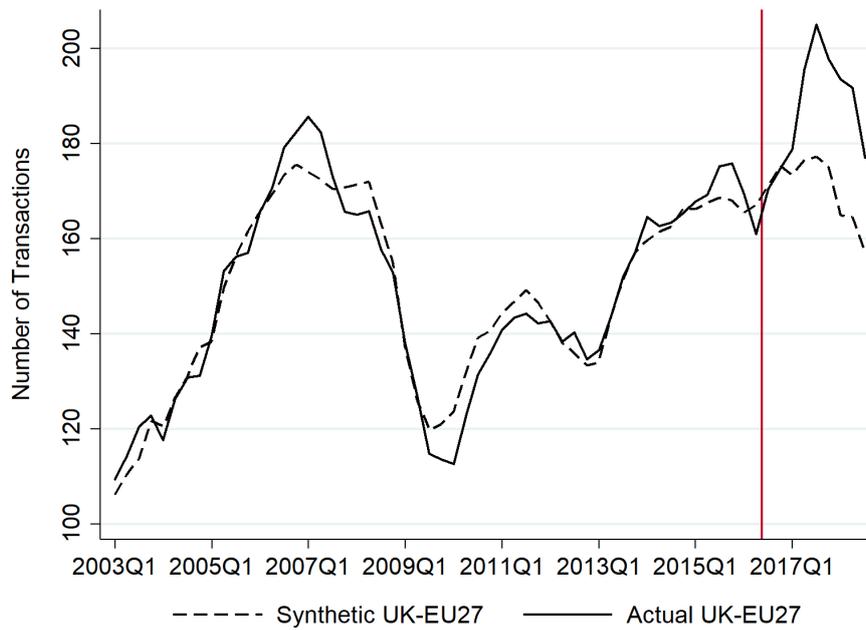
This figure plots the actual count of FDI transactions from the UK to the EU27 (solid line) and the corresponding synthetic control series (dashed line). The synthetic control series was computed using the full set of pairs in the donor group. Source: fDi Markets, Zephyr and authors' calculations.

Figure A.2: UK-EU27 FDI counts (actual vs synthetic control), including all transactions



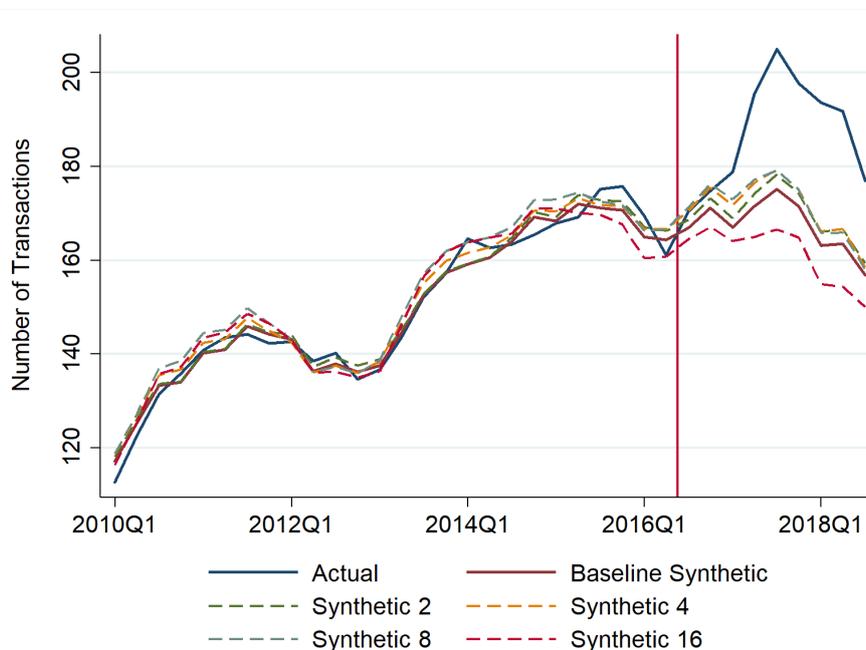
This figure plots the actual count of FDI transactions from the UK to the EU27 (solid line) and the corresponding synthetic control series (dashed line). Pairs with five or fewer transactions were included in the donor group. Source: fDi Markets, Zephyr and authors' calculations.

Figure A.3: UK-EU27 FDI counts (actual vs synthetic control), 2003Q1-2018Q3



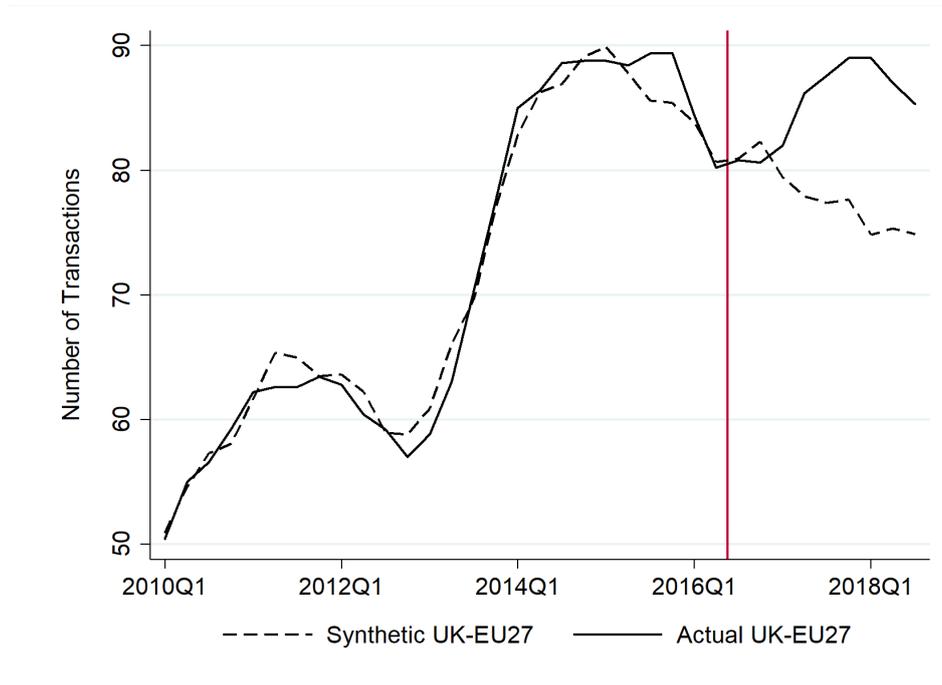
This figure plots the actual count of FDI transactions from the UK to the EU27 (solid line) and the corresponding synthetic control series (dashed line) for the period 2003Q1-2018Q3. Source: fDi Markets, Zephyr and authors' calculations.

Figure A.4: UK-EU27 FDI counts (actual vs synthetic control), including additional covariates



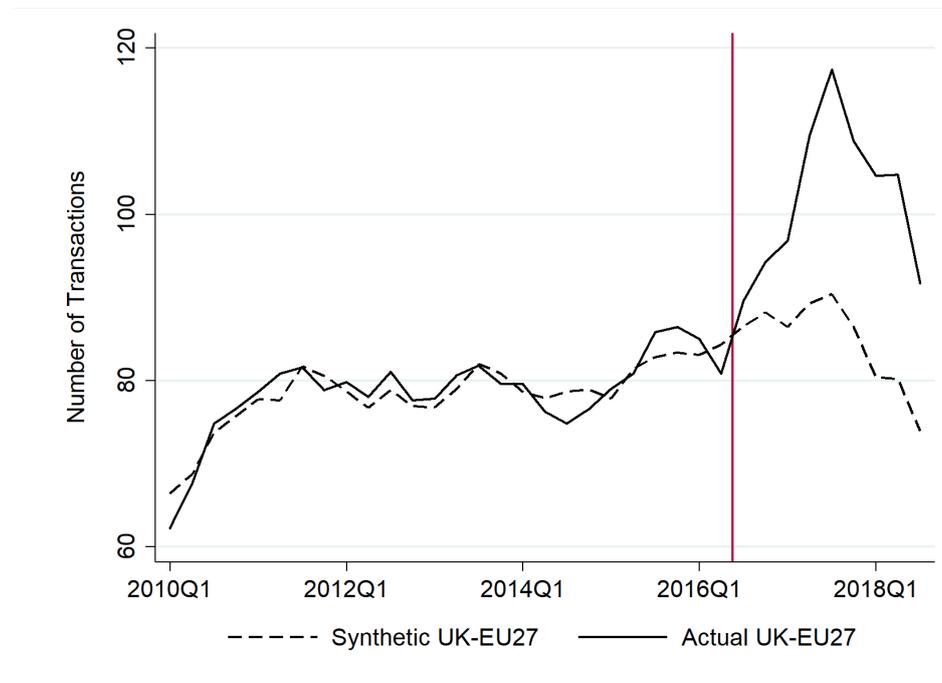
This figure plots the actual count of FDI transactions from the UK to the EU27 (solid line) and the corresponding synthetic control series (dashed lines). The dashed lines are based on using only every second, fourth, eighth or sixteenth pre-intervention outcome, respectively, to calculate the synthetic control. The additional covariates are bilateral distance and origin and destination GDP levels. Source: fDi Markets, Zephyr and authors' calculations.

Figure A.5: UK-EU27 M&A counts (actual vs synthetic control)



This figure plots the actual count of M&A transactions from the UK to the EU27 (solid line) and the corresponding synthetic control series (dashed line). Source: Zephyr and authors' calculations.

Figure A.6: UK-EU27 greenfield counts (actual vs synthetic control)



This figure plots the actual count of greenfield transactions from the UK to the EU27 (solid line) and the corresponding synthetic control series (dashed line). Source: fDi Markets and authors' calculations.