
WORKING PAPERS

Retrospective causal inference via
matrix completion, with an evaluation
of the effect of European integration on
crossborder employment

Jason **POULOS**^{1,2}
Andrea **ALBANESE**³
Andrea **MERCATANTI**⁴
Fan **LI**¹

¹ Duke University, United States of America

² Statistical and Applied Mathematical Sciences Institute (SAMSI), United States of
America

³ Luxembourg Institute of Socio-Economic Research (LISER), Luxembourg

⁴ Bank of Italy, Italy

LISER Working Papers are intended to make research findings available and stimulate comments and discussion. They have been approved for circulation but are to be considered preliminary. They have not been edited and have not been subject to any peer review.

The views expressed in this paper are those of the author(s) and do not necessarily reflect views of LISER. Errors and omissions are the sole responsibility of the author(s).

Retrospective causal inference via matrix completion, with an evaluation of the effect of European integration on cross-border employment

Jason Poulos[†].

Duke University and Statistical and Applied Mathematical Sciences Institute (SAMSI), Durham, NC, USA.

Andrea Albanese

Luxembourg Institute of Socio-Economic Research (LISER), Esch-sur-Alzette, Luxembourg.

Andrea Mercatanti

Bank of Italy, Rome, Italy.

Fan Li

Duke University, Durham, NC, USA.

Summary. We propose a method of retrospective counterfactual imputation in panel data settings with later-treated and always-treated units, but no never-treated units. We use the observed outcomes to impute the counterfactual outcomes of the later-treated using a matrix completion estimator. We propose a novel propensity-score and elapsed-time weighting of the estimator's objective function to correct for differences in the observed covariate and unobserved fixed effects distributions, and elapsed time since treatment between groups. Our methodology is motivated by studying the effect of two milestones of European integration—the Free Movement of persons and the Schengen Agreement—on the share of cross-border workers in sending border regions. We apply the proposed method to the European Labour Force Survey (ELFS) data and provide evidence that opening the border almost doubled the probability of working beyond the border in Eastern European regions.

Keywords: Causal inference; cross-border employment; European integration; matrix completion; panel data.

1. Introduction

We consider the problem of estimating the causal effects of a policy intervention in panel data settings where there exists units exposed to the policy after a given time (*later-treated*) and units that are exposed to the policy throughout the study period (*always-treated*). This is distinct from the standard setting in comparative case studies which consist of later-treated and never-treated units, but no always-treated units.

[†]*Address for correspondence:* Duke University Box 90251, Durham, NC 27708. *E-mail:* jason.poulos@duke.edu.

Motivated by the setting with later-treated and always-treated units, we formulate a new identification strategy of *retrospective* counterfactual imputation to predict potential outcomes under treatment for the later-treated in the pre-treatment period. We treat these potential outcomes as missing and impute them via the matrix completion method (Athey et al., 2017). We extend the matrix completion estimator two ways: first, by propensity- and elapsed-time weighting the objective function to correct for differences in the observed covariate and unobserved fixed effects distributions and elapsed time since treatment between treated and control groups; second, we explicitly include time-varying covariates in the outcome model and propose a procedure for imputing endogenous covariate values when reconstructing the missing potential outcomes.

A standard method for causal inference with panel data is difference-in-differences. This method hinges on the *parallel trends* assumption, which states that in the absence of treatment, the average outcomes of treated and control units would have followed parallel paths. Under parallel trends, difference-in-differences identifies causal effects by contrasting the change in outcomes pre- and post-treatment, among the treated and control groups (Ashenfelter, 1978; Ashenfelter and Card, 1985; Card and Krueger, 1994; Heckman et al., 1997; Callaway and Sant’Anna, 2018). Estimation in difference-in-differences is traditionally tied with an additive fixed-effects regression model (Angrist and Pischke, 2009). Less model-dependent semiparametric approaches (Abadie, 2005; Ding and Li, 2019), as well as matching and weighting methods (Hazlett and Xu, 2018; Kim and Wang, 2018; Strezhnev, 2018) have also been developed. However, the parallel trends assumption is generally invalid in the presence of unobserved time-varying confounders. In the retrospective setting, difference-in-differences identifies the effect of the treatment for later-treated in the pre-treatment period, but under a parallel trends assumption for the potential outcome under treatment. The matrix completion estimator relaxes this parallel trends assumption by allowing the presence of unobserved time-varying confounders.

A popular alternative method of handling unobserved time-varying confounders in panel data is the synthetic control method (Abadie and Gardeazabal, 2003; Abadie et al., 2010, 2015), which constructs a “synthetic control unit” as the counterfactual for single treated unit by reweighting the control units given their pre-treatment covariates or lagged outcomes. The idea is that conditioning on pre-treatment observables may help balance unobserved time-varying confounding between treatment and control groups. Recent methods have been developed to combine features of difference-in-differences and synthetic control (Arkhangelsky et al., 2019; Ben-Michael et al., 2018), or perform dimension reduction prior to reweighting the control units using lasso regression or matrix factorization (Belloni et al., 2017; Carvalho et al., 2018; Amjad et al., 2018). Doudchenko and Imbens (2016) propose a generalization of that relaxes the convexity restriction of the synthetic control method, which precludes negative and nonlinear interactions between control unit outcomes.

Xu (2017) further generalize by allowing for multiple treated units and treatment periods using an interactive fixed effects (i.e. latent factor) model (Bai, 2003, 2009; Gobillon and Magnac, 2016). The advantages of matrix completion over the interactive fixed effects models are three-fold: first, matrix completion does not require fixing the rank (i.e., number of unobserved factors) of the underlying data; second, it works in settings with staggered adoption of treatment (Athey and Imbens, 2018); third, matrix

completion uses all observed data to estimate unobserved factors, while interactive fixed effects models use pre-treatment data in the standard setting. An implication of the latter advantage is that matrix completion works well under weaker conditions; i.e., when there are a number of weak patterns over time that are common to units.

Our proposed method is motivated by studying the effects of two milestones of the European integration, namely the Schengen Agreement and the Free Movement of persons, on the labour market of bordering regions. The Schengen Agreement introduced a common visa policy and abolished border controls between countries with the goal of boosting economic integration by facilitating crossing of borders. Removing border controls may encourage cross-border employment by decreasing travel costs in terms of average and the variance of commuting time. The Freedom of Movement abolished the requirement of a work permit within the European Union (EU) as well as the requirement to prioritize local job seekers over other Europeans. This policy is expected to encourage job search across the border and widen job market opportunities of the poorer areas.

Our identification strategy relies on the staggered entry in the Schengen and the Freedom of Movement of European countries. Using European Labour Force Survey (ELFS) data aggregated at the regional-level for the period 2005-2019, we estimate their effect on the share of cross-border workers in sending border regions. We identify treatment effects by comparing sending regions to other border regions that maintained a stable status during the observational period. As most of the available control regions in the ELFS have already entered Schengen and the Freedom of Movement in 2005 when we start the analysis, we focus on these border regions, which make our control group composed of always-treated units. We provide evidence that opening the border almost doubled the probability of working beyond the border in Eastern European regions. The finding is consistent with previous studies that find opening the border increases the share of non-resident workers in host countries (e.g., Dustmann et al., 2016; Beerli et al., 2020).

The paper is structured as follows. We first provide background on the Schengen and the Freedom of Movement Agreements in Section 2. Section 3 introduces the setup of the retrospective causal estimation problem and the general matrix completion method. Section 4 discusses the specific modeling and estimation procedure adapted to the empirical application. Section 5 describes the data and presents the results for the empirical application. Section 6 concludes.

2. Institutions

Since its foundation, the European Union (EU) has been based on the integration of different economic systems through the removal of obstacles to the movement of people, goods, capital, and services. Geographical mobility of people has been a precondition to the creation of the European common currency, the Euro, because an optimal monetary area requires free movement of inputs to absorb regional shocks (e.g., Mundell, 1961), smooth spatial mismatches (e.g., Kain, 1968; Gobillon et al., 2007), and attenuate differences in the regional labour markets (e.g., Borjas, 2001; Niebuhr et al., 2012).

In this section, we firstly describe the institutional formation and contemporary status of two milestones of the European integration, the Schengen Agreement and the Freedom of Movement (FoM). Secondly, we mention the findings of the previous empirical

literature on the effect of these institutions.

2.1. Schengen Agreement

The Schengen Agreement was signed in 1985 by the Benelux countries, France and West Germany. This agreement had the goal of gradually abolishing the passport controls, harmonizing the visa policy and enhancing security coordination between countries. The crossing of the border was facilitated, and border controls were replaced by visual checks on private vehicles at a reduced speed. The full abolition of the border controls in the internal EU border and the creation of a single external border came into force in 1995, when also Spain and Portugal joined the Area. To access the Schengen Area, the countries had to meet specific requirements in the terms of external border controls, issuing visas, police cooperation and personal data protection.

The Schengen Agreement operated independently from the EU legislation until 1999 when it was incorporated in the Treaty of Amsterdam. Most of the EU countries progressively joined the Schengen Area, while four EU countries, namely Bulgaria, Croatia, Cyprus, and Romania, are not yet part of it. Today, the Schengen Area covers about 425 million people and 4.3 million square kilometre area and includes twenty-two EU countries and four European Free Trade Association (EFTA) members, including Iceland, Liechtenstein, Norway, and Switzerland.

2.2. Freedom of Movement (FoM)

The Freedom of Movement (FoM) in Europe abolished the requirement of a work permit within the EU as well as the requirement to prioritize local job seekers over other Europeans. Differently from the Schengen Agreement, the FoM was always an integral part of the EU integration and therefore applied automatically to all member states. EU enlargements were usually followed by temporary restrictions from some countries such as the ten eastern European countries entering the EU in 2004 (the so-called A10). While most of the A10 countries granted access to their labour market between 2004 and 2007, Austria and Germany, which shared the border with five A10 countries, maintained their system of quotas and work permits until 2011, which we exploit in our identification strategy. Switzerland, a non-EU country, progressively granted labour market access to EU residents. In 2004, full labour market access was granted to cross-border workers working in the regions of Switzerland close to the border. In 2007, this was extended to all regions of Switzerland and the FoM was granted also to EU migrants, with the abolition of the annual quotas.

2.3. Related literature

According to the economic theory, European institutions such as the Schengen Agreement and the FoM may affect cross-border commuting. By decreasing travel costs in terms of average and the variance of commuting time, the Schengen Agreement may encourage crossing of borders. Similarly, abolishing the requirement of a work permit as well as the requirement to prioritize local job seekers over other Europeans, the FoM may widen job market opportunities of individuals living in poorer areas beyond the border.

However, studies on the effects of these institutions on cross-border employment of sending regions are scarce. Most of the literature has focused on their effects on migration (e.g., Elsner, 2013; Beine et al., 2019), trade integration (Chen and Novy, 2011; Davis and Gift, 2014; Felbermayr et al., 2018), cross-regional labour market correlations (Bartz and Fuchs-Schündeln, 2012), labour market outcomes of natives in the host regions (Naguib, 2019; Basten and Siegenthaler, 2019; Aepli and Kuhn, 2019; Åslund and Engdahl, 2019; Beerli et al., 2020), and criminality (Sandner and Wassmann, 2018).

An exception in the literature is Parenti and Tealdi (2019), who assess the effect of the Schengen Agreement on the cross-border incidence in the nearby sending regions. The authors rely on ELFS data and restrict their sample on individuals working in another region than the one of residence. They implement a standard difference-in-differences estimator by comparing the incidence of cross-border employment out of total commuters in Italian, French and German regions who shared a border with Switzerland to a control group of regions from the same countries which instead shared the border with other countries that had already accessed Schengen at the beginning of the analysis (always-treated). The authors find an increase in the probability of cross-border employment (conditional on commuting) following the entry of Switzerland in the Schengen Area, while finding no effect attributable to the FoM.

The empirical literature has also shown that other socio-economic factors may affect commuting across the border. Mathä and Wintr (2009) state that in the Greater Region of Luxembourg the border discourages commuting flows due to language differences. Cross-border commuting is instead more likely to occur the wider the asymmetries in the local labour markets and the easier the accessibility. Similar results are also found by Ahrens et al. (2020) on the Irish border, which suggests that socio-cultural factors may discourage cross-border commuting despite the absence of custom and border checks. This is also found for commuting across internal regions in Belgium, which confirms that language and cultural differences can discourage commuting (Persyn and Torfs, 2015). Egger and Lassmann (2015) show that internal language borders in Switzerland can also affect trade. Finally, Bello (2019) shows that in Switzerland variations in the exchange rate may also influence the extensive and intensive margin of cross-border labour supply.

2.4. Staggered implementation

To assess the effects of these European institutions, we rely on their staggered implementation for the different regions of the EU. We focus on two clusters of sending regions which entered these institutions between 2007 and 2011 — namely, the regions of Eastern Europe and Switzerland. The first group is composed of external regions in Poland, Czech Republic, Slovakia, Slovenia and Hungary, which share a border with Germany or Austria and accessed Schengen in the 1st quarter of 2008 and the FoM with Germany and Austria in the first quarter of 2011. Second, the regional cluster around Switzerland (Italy, France, Germany and Austria) fully entered the FoM in the second quarter of 2007 and accessed Schengen in the first quarter of 2009.

To identify treatments effects, we compare these two sets of regions (later-treated) to the other border regions that maintained a stable status during the observational period. Differently from the previous literature, in our evaluation setting there is virtually no never-treated unit since most of the regions have already entered the treatments at the

beginning of the analysis. The regions of Romania and Bulgaria could compose the never-treated group for the Schengen Agreement. However, their treatment status is also not stable since they accessed the EU (and the FoM) during the analysis. Furthermore, their cross-border flow to bordering countries is close to zero. Therefore, we focus on these border regions, which make our control group composed of always-treated units. This identification strategy is formalized in the following section.

3. Retrospective causal estimation and the general matrix completion strategy

In our empirical application, we have later-treated units — i.e., Eastern cluster and Swiss border regions — and always-treated units — i.e., regions that maintained a stable treatment status throughout the study period — and we are interested in estimating the effect of a policy intervention (i.e., European integration) on the later-treated units.

To formalize our setting, consider a sample of N units, each of which was observed in T periods of time. The units are divided into two groups: the always-treated (AT) group, where N_{AT} are treated throughout the study period, and the later-treated (LT) group, where N_{LT} units are exposed to the treatment at the initial treatment period, $T_0 (< T)$; and $N = N_{AT} + N_{LT}$. For each unit i and at each time t , we observe an outcome Y_{it} , a set of covariates X_{it} , and the treatment indicator W_{it} . For the always-treated units, $W_{it} = 1$ for all t ; for the later-treated units, $W_{it} = 0$ for $t < T_0$ (denoted by “pre” period) and $W_{it} = 1$ for $T_0 \leq t \leq T$ (denoted by “post” period).

Under the Neyman-Rubin potential outcomes framework (Neyman, 1923; Rubin, 1974), for each unit i and time t , there exists a pair of potential outcomes, $Y(1)_{it}$ and $Y(0)_{it}$, corresponding to potential outcomes under the treatment and control condition, respectively. The potential outcomes framework implicitly assumes treatment is well-defined to ensure that each unit has the same number of potential outcomes. It also excludes interference between units, which would undermine the framework by creating more than two potential outcomes per unit, depending on the treatment status of other units. The fundamental problem is that we can only observe $Y(1)_{it}$ or $Y(0)_{it}$ for each (i, t) pair of indices.

In the retrospective setting, the causal estimand is the average treatment effect for the later-treated units and its average:

$$\tau_t^{LT,pre} = E[Y(1)_{it} - Y(0)_{it}], \quad \forall i \in LT, t < T_0; \quad \text{and} \quad \tau^{LT,pre} = E_t[\tau_t^{LT,pre}]. \quad (1)$$

In the above estimand, the potential outcome under control in the pre-treatment period is observed, but the counterfactual outcome of the later-treated units under treatment in the pre-treatment period is unobserved and has to be estimated from the observed quantities. In our empirical application, the estimand can be interpreted as the causal effect of European integration on the later-treated regions in the years prior to the implementation of Schengen and FoM, as if they had been always-treated.

We now introduce the matrix notation of the above setup. Let $\mathbf{Y} \in \mathbb{R}^{N \times T}$ be the matrix of outcomes with typical element Y_{it} , and $\mathbf{W} \in \mathbb{R}^{N \times T}$ be the treatment assignment matrix with typical element $W_{it} \in \{0, 1\}$. In the retrospective setting, we partition $\mathbf{Y}(1)$, the $N \times T$ matrix of the potential outcome under treatment, by treatment group and periods. In order to identify the causal estimand, $\tau_t^{LT,pre}$, we need to impute the

missing quadrant using information from the observed part of the matrix:

$$\mathbf{Y}(1) = \left(\begin{array}{c|c} \mathbf{Y}(1)_{AT, \text{pre}} & \mathbf{Y}(1)_{AT, \text{post}} \\ \hline ? & \mathbf{Y}(1)_{LT, \text{post}} \end{array} \right), \quad (2)$$

where the question mark corresponds to the missing potential outcomes, $\mathbf{Y}(1)_{LT, \text{pre}}$.

We illustrate the idea of retrospective causal estimation by exploiting the similarity between the matrix completion and interactive fixed effects models, which are both specified in terms of time-varying factors and unit-specific factor loadings. Following the steps of the interactive fixed effects estimation procedure is effective to clarify how the retrospective analysis works. We first write the potential outcomes under treatment $\mathbf{Y}(1)$ as:

$$Y(1)_{it} = \sum_{r=1}^R \gamma_{ir} \delta_{tr} + \epsilon_{it},$$

or $\mathbf{Y}(1) = \mathbf{U}\mathbf{V}^\top + \boldsymbol{\epsilon},$ (3)

where $\mathbf{U}_{N \times R}$ is a matrix of factor loadings (i.e., unit-specific intercepts) with typical element γ_{ir} , $\mathbf{V}_{T \times R}$ is a matrix of factors (i.e., time-varying coefficients) δ_{tr} , and $\boldsymbol{\epsilon}_{N \times T}$ is a matrix of random noise ϵ_{it} . In the literature on interactive fixed effects models (e.g., Bai, 2003, 2009; Gobillon and Magnac, 2016; Xu, 2017), estimating \mathbf{U} and \mathbf{V} relies fixing the number of factors, R . In the matrix completion literature (e.g., Candès and Recht, 2009; Mazumder et al., 2010; Recht, 2011), it is only assumed $\mathbf{U}\mathbf{V}^\top$ is low rank.

In the retrospective setting, we partition \mathbf{U} and \mathbf{V} analogous to Eq. (2),

$$\mathbf{U}\mathbf{V}^\top = \begin{pmatrix} \mathbf{U}_{AT} \\ \mathbf{U}_{LT} \end{pmatrix} \begin{pmatrix} \mathbf{V}_{\text{pre}}^\top \\ \mathbf{V}_{\text{post}}^\top \end{pmatrix}. \quad (4)$$

Using data from the always-treated in both periods, and post-treatment data for the later-treated, we can then model the observed part of the $\mathbf{Y}(1)$ matrix as:

$$\mathbf{Y}(1)_{AT} = \mathbf{U}_{AT} \begin{pmatrix} \mathbf{V}_{\text{pre}}^\top \\ \mathbf{V}_{\text{post}}^\top \end{pmatrix} + \boldsymbol{\epsilon}_{AT}, \quad \text{and} \quad \mathbf{Y}(1)_{LT, \text{post}} = \mathbf{U}_{LT} \mathbf{V}_{\text{post}}^\top + \boldsymbol{\epsilon}_{LT, \text{post}}. \quad (5)$$

Xu (2017) proposes a two-step method, where the left-hand-side equation in (5) is an estimator for \mathbf{U}_{AT} , \mathbf{V}_{pre} , and \mathbf{V}_{post} , and the right-hand-side equation is an estimator for \mathbf{U}_{LT} given $\hat{\mathbf{V}}_{\text{post}}$, both by regularized least squares,

$$\min_{\mathbf{U}_{AT}, \mathbf{V}_{\text{pre}}, \mathbf{V}_{\text{post}}} \left\| \mathbf{Y}_{AT} - \mathbf{U}_{AT} \begin{pmatrix} \mathbf{V}_{\text{pre}}^\top \\ \mathbf{V}_{\text{post}}^\top \end{pmatrix} \right\|_0 \quad \text{and} \quad \min_{\mathbf{U}_{LT}} \left\| \mathbf{Y}_{LT, \text{post}} - \mathbf{U}_{LT} \hat{\mathbf{V}}_{\text{post}} \right\|_0, \quad (6)$$

where $\|\cdot\|_0 = \sum_i \mathbf{1}_{\sigma_i(\cdot) > 0}$ is the rank norm. Athey et al. (2017) note that the left-hand-side equation in (6) is not an efficient estimator for \mathbf{V} because it only uses information from the always-treated and does not use $\mathbf{Y}_{LT, \text{post}}(1)$ when estimating unobserved factors.

Solving Eq. (6) yields $\hat{\mathbf{U}}_{\text{LT}}$ and $\hat{\mathbf{V}}_{\text{pre}}$, which is used to predict the missing potential outcomes:

$$\hat{\mathbf{Y}}(1)_{\text{LT, pre}} = \hat{\mathbf{U}}_{\text{LT}} \hat{\mathbf{V}}_{\text{pre}}^{\top}, \quad (7)$$

making feasible to estimate, $\hat{\tau}_t^{LT,pre} = \hat{\mathbf{Y}}(1)_{\text{LT, pre}} - \mathbf{Y}(0)_{\text{LT, pre}}$. The estimate $\hat{\mathbf{Y}}(1)_{\text{LT, pre}}$ can be thought as a *retrospective imputation*.

4. Estimation

In this section, we describe our estimation strategy of estimating retrospective causal effects and extensions of the matrix completion method proposed by Athey et al. (2017) to accommodate the retrospective setting.

4.1. Outcome model

In the empirical application, the outcome of interest is the share of residents working in another country that shares the border with the region of residence, conditional or unconditional on employment. We model the share of cross-border workers under treatment as:

$$Y(1)_{it} = L_{it} + X_{it}\beta_t + \gamma_i + \delta_t + \epsilon_{it}, \quad (8)$$

where \mathbf{L} is an unknown matrix to be estimated, with typical element L_{it} . Eq. (8) is identical to the outcome model of Athey et al. (2017), except we allow for a unit- and time-varying covariate matrix \mathbf{X} with typical element X_{it} and estimate a T -length vector of time-varying covariate coefficients, β_t . In the application, \mathbf{X} takes the form of regional GDP divided by the highest GDP of the nearby foreign regions. Unit-specific fixed effects, $\gamma = \{\gamma_1, \dots, \gamma_N\}$, and time-specific fixed effects, $\delta = \{\delta_1, \dots, \delta_T\}$ control for unit-specific (time-invariant) and time-specific (unit-invariant) unmeasured confounding. The identifying assumption is that the errors ϵ_{it} are independent across i , are conditional mean zero and are conditionally independent on W_{it} (Xu, 2017):

$$\mathbb{E}(\epsilon_{it} | L_{it}, X_{it}, \gamma_i, \delta_t) = \mathbb{E}(\epsilon_{it} | W_{it}, L_{it}, X_{it}, \gamma_i, \delta_t) = 0. \quad (9)$$

This exogeneity assumption is analogous to the unconfoundedness assumption in the literature on semiparametric difference-in-differences estimators (Abadie, 2005); namely, that unconfoundedness is conditional on covariates and fixed effects, the latter of which are imposed to capture unmeasured confounding (Angrist and Pischke, 2009).

The estimation problem is to recover the entire matrix \mathbf{L} despite the missing part of the matrix, $\mathbf{Y}_{\text{LT, post}}(1)$. It is assumed that \mathbf{L} can be approximated by a low-rank matrix, $\mathbf{L}_{N \times T} \approx \mathbf{U}_{N \times R} \mathbf{V}_{T \times R}^{\top}$, where R is the number of non-zero singular values $\sigma_i(\mathbf{L})$ recovered from a singular value decomposition of the matrix, $\hat{\mathbf{L}}_{N \times T} = \mathbf{S}_{N \times N} \mathbf{\Sigma}_{N \times T} \mathbf{R}_{T \times T}^{\top}$, where \mathbf{S} and \mathbf{R} are unitary matrices and $\mathbf{\Sigma}$ is a rectangular diagonal matrix with $\sigma_i(\mathbf{L})$ on the diagonal.

Let \mathcal{O} denote the set of observed values; i.e., the values for which $W_{it} = 1$. Athey et al. (2017) propose a nuclear norm estimator for \mathbf{L} in the standard “forward-looking”

approach that we reformulate below for the retrospective analysis:

$$\arg \min_{\mathbf{L}, \beta, \gamma, \delta} \left[\frac{1}{|\mathcal{O}|} \sum_{(i,t) \in \mathcal{O}} \frac{1 - \hat{w}_{it}}{\hat{w}_{it}} \left(Y_{it}(1) - L_{it} - X_{it}\beta_t - \gamma_i - \delta_t \right)^2 + \lambda_L \|\mathbf{L}\|_{\star} + \lambda_{\beta} \|\beta\|_1 \right], \quad (10)$$

where \hat{w}_{it} is the individual predicted probability to be treated at time t , and the nuclear norm, or sum of singular values, $\|\cdot\|_{\star} = \sum_i \sigma_i(\cdot)$, is used to yield a low-rank solution for \mathbf{L} by decomposing $\mathbf{Y}(1)$ in factors and factor loadings. The vector ℓ_1 norm, $\|\cdot\|_1$, is used to prevent overfitting by shrinking the values of β_t toward zero. The values of the penalty hyperparameters λ_L and λ_{β} are chosen by cross-validation.

The exogeneity assumption (9) and the goal to impute the missing potential outcomes, $\mathbf{Y}(1)$, $(i, t) \in \bar{\mathcal{O}}$, creates an analogy with a weighting estimation procedure under unconfoundedness (Hirano and Imbens, 2001; Li et al., 2018). Specifically, the weights $\frac{1 - \hat{w}_{it}}{\hat{w}_{it}}$ in (10) balance the observed covariate and unobserved fixed effects distributions between different groups when the target population is the untreated units. These weights place more emphasis on the loss, in (10), for the unit-time elements $(i, t) \in \mathcal{O}$ most similar to the elements $(i, t) \in \bar{\mathcal{O}}$ in terms of both observed covariates and unobserved fixed effects.

To quantify \hat{w}_{it} , we model treatment as a linear function of a low-rank representation of treatment, \mathbf{E} , covariates X_{it} , time-varying coefficients ϕ_t , and fixed effects ξ_i and ψ_t :

$$W_{it} = E_{it} + [X_{it}I(W_{it} = 1) + \hat{X}_{it}I(W_{it} = 0)]\phi_t + \xi_i + \psi_t + v_{it}, \quad (11)$$

where \hat{X}_{it} is an imputed value for X_{it} to correct for endogeneity of covariates, as detailed in the following subsection, and $I(\cdot)$ denotes the indicator function. This expression is similar to the treatment model proposed by Athey et al. (2017).

We then estimate the matrix of the treatment assignments, \mathbf{E} , analogously to \mathbf{L} in the outcome model (8), by nuclear norm regularized least squares:

$$\arg \min_{\mathbf{E}, \phi, \xi, \psi} \left[\sum_{(i,t)} \frac{1}{NT} \left(W_{it} - E_{it} - [X_{it}I(W_{i,t} = 1) + \hat{X}_{i,t}I(W_{i,t} = 0)]\phi_t - \xi_i - \psi_t \right)^2 + \lambda_L \|\mathbf{E}\|_{\star} + \lambda_{\phi} \|\phi\|_1 \right], \quad (12)$$

This way \mathbf{E} is the product of a matrix containing R unobserved treatment assignment factors (unit-invariant), $\mathbf{H}_{T \times R}$, and a matrix of unit (time-invariant) factor loadings, $\mathbf{\Lambda}_{N \times R}$, so that any unit loads each treatment assignment factors differently. Once $L_{it}, \beta_t, \gamma_i, \delta_t$ have been estimated, given assumption (9), we can predict $\hat{Y}(1)_{it}$ for the later-treated in the pre-treatment period by

$$\hat{Y}(1)_{it} = \hat{L}_{it} + \hat{X}_{it}\hat{\beta}_t + \hat{\gamma}_i + \hat{\delta}_t,$$

for (i, t) such that $i \in LT$ and $t < T_0$.

4.2. Imputing endogenous covariate values

The minimization problem in Eq. (10) requires unit covariate values that are observed both at the pre- and post-treatment periods. Including unit characteristics that are

possibly associated with the dynamics of the outcome helps to avoid a strict parallel trends assumption, such as in difference-in-differences analyses. On the other hand, the outcome model (8) imposes restrictions on the choice of the covariates. Like in standard difference-in-differences analyses, the model requires the so-called exogeneity condition for the covariates, which imposes that they are not affected by the treatment: $X_{it} = X(0)_{it} = X(1)_{it}$, where $X(w)_{it}$ is the potential quantity for the covariate when $W_{it} = w$, $w = \{0, 1\}$ (Lechner, 2010). However, in situations where the covariates are endogenous, so that $X(0)_{it} \neq X(1)_{it}$, the observed covariates are $X_{it} = X(1)_{it}I(W_{i,t} = 1) + X(0)_{it}I(W_{i,t} = 0)$, making

$$Y(1)_{it} = L_{it} + X(1)_{it}\beta_t + \gamma_i + \delta_t + \epsilon_{it} \neq L_{it} + X_{it}\beta_t + \gamma_i + \delta_t + \epsilon_{it}, \quad (13)$$

for (i, t) such that $i \in LT$ and $t < T_0$.

We deal with the problem of endogenous covariate values by treating endogenous values of the covariates the same way we consider the missing potential outcomes, namely as intentional missing data, and impute the endogenous covariate values as follows. We perform the matrix completion procedure in Eqs. (8) and (10), but without covariates and instead considering the covariate matrix \mathbf{X} as the outcome. We make the values of the covariate for the later-treated before treatment intentionally missing and replace the missing values with the corresponding values, \hat{X}_{it} , from the reconstructed matrix.

4.3. Elapsed-time weighted loss function

A possible complication for our retrospective analysis arises from the fact that we observe the outcomes under treatment, $\mathbf{Y}(1)$, for the later-treated and always-treated groups in the same calendar time t but the elapsed time since T_0 , z_t , is different for the two groups. This means that the observed evolution of $\mathbf{Y}(1)$ of the two groups might not be just linked to the different effect of t , but also to z_t . The latter is particularly important if the treatment effect takes some time before showing its full potential. According to economic theory, it is likely that the treatment effect stabilizes after some periods in a new "steady state" equilibrium. For example, it is possible that workers take time before adjusting their expectations of the job possibilities in the other countries and spreading the news to the other workers. The influence of z_t may be an issue in the retrospective setting, where we predict $\mathbf{Y}(1)_{LT,pre}$ by elaborating the information from the three time-series $\mathbf{Y}(1)_{LT,post}$, $\mathbf{Y}(1)_{AT,post}$, and $\mathbf{Y}(1)_{AT,pre}$, if only the two latter series are in a steady state equilibrium.

We address this problem in the retrospective analysis by weighting the objective function so that more weight is placed on the loss for observed values with higher values of z_t . More specifically, we use the product of \hat{w}_{it} and z_t to weight the objective function (10). To adjust for elapsed time since the treatment implementation, we compute $\text{diag}(\mathbf{z})\hat{\mathbf{W}}$, where \mathbf{z} concatenates two sequence of values drawn from a logistic model, the first sequence descending to T_0 and the second sequence ascending from T_0 . This computation places more weight on the ends of $\hat{\mathbf{W}}$ for the later-treated units. The weights used to adjust the loss for later-treated units in Eq. (10) is then $I_{i \in LT} \frac{1 - \text{diag}(z_t)\hat{w}_{it}}{\text{diag}(z_t)\hat{w}_{it}}$.

Table 1. Staggered treatments.

Period	Eastern cluster	Swiss cluster	Always-Treated
2005Q1-2007Q2	-	(Mild FoM)	Schengen + FoM
2007Q3-2007Q4	-	FoM	Schengen + FoM
2008Q1-2008Q4	Schengen	FoM	Schengen + FoM
2009Q1-2010Q4	Schengen	Schengen + FoM	Schengen + FoM
2011Q1-2019Q4	Schengen + FoM	Schengen + FoM	Schengen + FoM

5. Empirical Application

In this section, we describe the survey data used for the application; we then estimate the effects of Schengen and FoM on the share of residents working in a neighbouring country, conduct placebo tests to assess if the effects are truly attributable to the treatment, and assess the accuracy of the matrix completion estimator.

5.1. Data

We rely on the harmonized quarterly ELFS, which covers a sample of resident households in the different regions of the EU from the 1st quarter of 2005 to the 4th quarter of 2019. The ELFS is the largest household survey in Europe with about 1.7 million interviews per quarter and an effective sampling rate of 0.29% in 2017. Individuals living in the household are asked questions on their labour status, working time, education, sex, age, nationality and, importantly for our analysis, the region of residence and the country of work. Information on the location of residence is usually provided at regional level though some countries provide this information only at an aggregated level. Most countries have a rotation scheme in place, though the anonymised ELFS microdata do not contain the information which would allow tracking people across waves. We focus on the working age population and retain all the individuals aged between 20 and 59. As we rely on panel data methods in the econometric analysis, we then collapse the data at the level of the region of residence and follow the evolution of the aggregated outcomes over time. To enhance harmonization, we start following the regions from the first quarter of 2005, when all countries have transited from a yearly survey (usually in Spring) to a continuous quarterly survey, covering all weeks of the year. We exclude all regions not sharing an inland border with another EU or Schengen country, countries for which information on the country of work is missing in any year (Portugal, Ireland, Norway, Sweden and Switzerland) and the treated regions that tend to host the cross-border flows (i.e. regions in Austria, Germany and Italy sharing a border with East of Europe).

The outcome of interest is the share of residents working in another country, which shares the border with the region of residence, conditional or unconditional on employment. As the survey covers only residents, this outcome should include cross-border commuters or temporary migrants who come back to their home country regularly. The treatment statuses and periods, mentioned in Section 2.4, are summarized in Figure 1 and Table 1.

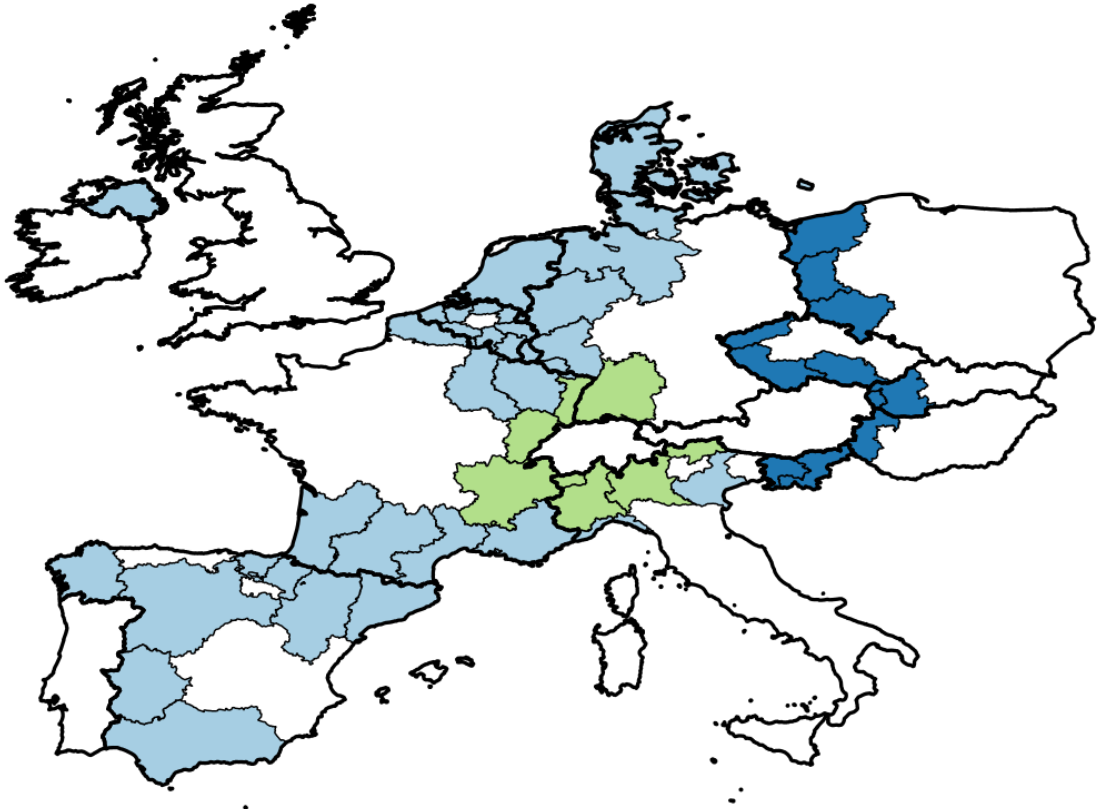


Fig. 1. Treated and control regions. Key: ■, controls (always-treated); ■, Schengen in 2008 and FoM in 2011; ■, FoM in 2007 and Schengen in 2009.

5.2. Results

We estimate the effect of the European integration on the share of residents working in a neighbouring country. Since there are two treatments in the application — Schengen and FoM — we set T_0 at the time when both treatments are in place. In the Eastern cluster, the Schengen policy and the FoM are adopted in the first quarters of 2008 and in 2011, respectively, so we set $T_0^{(\text{Eastern})} = 2011$. In the Swiss cluster, the FoM is adopted in the second quarter of 2007 and the Schengen policy is adopted in the first quarter of 2009, so we set $T_0^{(\text{Swiss})} = 2009$.

The top panel of Figure 2 shows the evolution of the (conditional) cross-border observed outcomes (solid lines) for the two clusters of later-treated regions and the always-treated border regions (dashed line). The descriptive evidence shows an increase in the outcome of both later-treated clusters of regions after the treatments. At the same time, the outcome of the always-treated regions remains quite stable and increases only marginally throughout the period of analysis.

We then check whether this suggestive descriptive evidence on a positive effect of the European institutions is confirmed by our matrix completion method. As described, we specify a propensity-score and elapsed-time weighted loss function (8)-(10). We then predict the counterfactual outcomes of the later-treated regions, $\hat{\mathbf{Y}}(1)_{\text{LT, pre}}$, using information from the observed outcomes of the later-treated and always-treated. In the bottom panel we plot the estimated treatment effects on later-treated units, $\hat{\tau}_t^{\text{LT, pre}}$, which are calculated by differencing the predicted outcomes of the later-treated units (dotted lines) and their observed outcomes. The estimated treatment effects are surrounded by percentile bootstrap confidence intervals constructed by block resampling the columns (i.e, time dimension) of the observed outcomes, in order to preserve temporal dependence structure of the original data (Davison and Hinkley, 1997; Politis and White, 2004). To infer the overall effect of treatment, we estimate percentile bootstrap confidence intervals for the counterfactual trajectory, $\hat{\tau}^{\text{LT, pre}}$, by resampling the trajectories without the time component.

The decomposition of the low-rank matrix $\hat{\mathbf{L}}_{N \times T} = \hat{\mathbf{S}}_{N \times N} \hat{\mathbf{\Sigma}}_{N \times T} \hat{\mathbf{R}}_{T \times T}^\top = \hat{\mathbf{U}}_{N \times R} \hat{\mathbf{V}}_{T \times R}^\top$ helps to explain the imputed outcome trends. The dominant (first) detected factor in $\hat{\mathbf{V}}_{R \times T}$, is mainly loaded by the later-treated (by means of the loading factors in the sub-matrix $\hat{\mathbf{U}}_{\text{LT} \times R}$), and visually resembles the imputed outcome trends of the later-treated in Figure 2. The rest of the factors are more stationary and are mainly loaded by the always-treated. The graphs for the detected first and second factors for both the Eastern and Swiss blocks are reported in Figure SM-1 in the Supplementary Material (SM).

Our results are in line with the descriptive evidence. The counterfactual trajectory estimates, reported in Table 2, show that had the Schengen and FoM both been implemented in the Eastern cluster prior to 2011, the share of cross-border workers out of the total working population would have been 0.8 percentage points [0.2, 1.6] higher. Relative to the observed outcome in the same period, this represents an increase of 86%. We find a positive effect on the share of cross-border workers in Swiss region of about 0.6 percentage points [-0.2, 2.3], which correspond to a relative increase of about 32%, although this is not statistically significant at the 5% level. The effects are marginally smaller if we do not condition the outcome on employment.

We compare the matrix completion estimates with difference-in-differences and syn-



Fig. 2. Matrix completion estimates of the effect of combined treatment (Schengen + FoM) on share of residents working in another country, conditional on employment, for the later-treated regions ($N \times T = 53 \times 60$). *Note:* The dashed vertical lines represent the initial treatment times of the combined treatment, $T_0^{(\text{Eastern})} = 2011$ and $T_0^{(\text{Swiss})} = 2009$. The shaded regions are 95% bootstrap confidence intervals of the effects constructed with 999 block bootstrap samples. *Key:* —, observed outcomes of Eastern regions; —, observed outcomes of Swiss regions; - - -, observed outcomes of always-treated regions; ○, counterfactual Eastern; △, counterfactual Swiss; ·····, $\hat{\tau}_t^{\text{Eastern}}$; ·····, $\hat{\tau}_t^{\text{Swiss}}$.

Table 2. Counterfactual trajectory estimates.

Estimator	Eastern (2005Q1-2010Q4)		Swiss (2005Q1-2008Q4)	
	Uncond.	Cond.	Uncond.	Cond.
Matrix com.	0.006 [0.002, 0.012]	0.008 [0.002, 0.016]	0.004 [-0.002, 0.016]	0.006 [-0.002, 0.023]
Diff.-in-diff.	0.008 [0.002, 0.015]	0.010 [0.003, 0.018]	0.005 [-0.003, 0.020]	0.006 [-0.003, 0.025]
Synth. con.	0.007 [0.002, 0.013]	0.009 [0.003, 0.016]	-0.001 [-0.007, 0.005]	0.001 [-0.008, 0.008]

Notes: Outcome is % working in border region, unconditional or conditional on employment. Bracketed values represent 95% percentile bootstrap confidence intervals for $\hat{\tau}^{LT.pre}$.

thetic control estimates, also reported in Table 2. The difference-in-differences estimator is a regression of outcomes on treatment and unit and time fixed effects (Athey and Imbens, 2018). The synthetic control method is a regression of the pre-treatment outcomes of each treated unit on the control unit outcomes during the same periods, with the restrictions of the original synthetic control method of Abadie et al. (2010), namely zero intercept and non-negative regression weights that sum to one (Doudchenko and Imbens, 2016; Athey et al., 2017). These restrictions are equivalent to imposing a restriction of linear dependence between the loading factors in the context of matrix completion or interactive fixed effects models. We provide the exact form of these estimators in Section SM-1.

The difference-in-differences and synthetic control method point estimates are within the confidence intervals of the matrix completion estimates. Note that we express the difference-in-differences as a two-way fixed effects as in Athey and Imbens (2018) to allow a comparability of results over the entire period 2005Q1-2019Q4. This collapses the temporal effects into a single parameter and thus loses the interpretation in terms of inter-temporal dynamics compared to the matrix completion method. The small difference in the estimated effect between difference-in-differences and matrix completion is due to the fact that the difference in the averaged outcomes in the post-treatment period between the later-treated and always-treated (for both the Eastern and the Swiss block) is similar to the total estimated effect from matrix completion. This occurrence is not indicative of whether the parallel trends assumption holds or not.

We also investigate the role of the two treatments, Schengen or the FoM, in driving the estimated positive effect. To do that, we divide the pretreatment period into two institutional periods for each regional cluster in which either no treatment or only one of the two treatments was implemented. As the estimated counterfactual represents the outcome under both treatments, the effect on the former provides evidence of the joint impact of the two European institutions, while the effect on the latter represents the impact of the treatment that was not yet implemented, conditional on the other treatment. As shown in Table SM-1, the treatment effect for the Eastern cluster fully exhibits during the period when Schengen was already operative, from 2008Q1 to 2010Q4. This finding suggests that the main driver of the estimated effect is the FoM, an institution that has removed the work permit requirement across European countries.

Similarly, the Swiss cluster shows a very small treatment effect when the FoM was operative (+0.3 percentage points during the period 2007Q3-2008Q4), which is more than

double when no treatment was implemented (+0.7 percentage points before 2007Q2), though it remains statistically insignificant at the 5% level. Removing border controls does not seem sufficient in boosting cross-border employment, which is in contrast with the results of Parenti and Tealdi (2019) for the Swiss cluster. This study however selected only individuals commuting between regions and ignored that the control group is composed of always-treated.

Finally, when interpreting the results for the FoM in the Swiss cluster we should keep in mind that our estimates only captures the “top-up” effects of the full liberalization of 2007. Indeed, as explained in Section 2.2, from 2004 cross-border workers were already allowed working in the border regions of Switzerland without a work permit, which according to the previous literature had increased the relative incidence of non-resident employment in the border regions of Switzerland (e.g., Beerli et al., 2020). The 2007 reform extended the FoM to the whole Swiss territory and included also migrants. Our estimates provide evidence that the this extension did not have a significant additional impact.

5.3. *Placebo tests and comparison with alternatives*

To assess whether the estimated effects are attributable to Schengen and FoM rather than other policy changes during the same period, we conduct placebo tests by conducting a “no treatment” evaluation; that is, we re-run the analysis on the post-treatment data, when no treatment effect is expected. In the first set of placebo tests, we fit the matrix completion estimator on the post-treatment years and randomize the initial treatment period for the actual treated units in both simultaneous or staggered treatment adoption settings. Letting $\hat{\tau}_{\pi}^{LT,pre}$ denote the average causal effects estimated for each block bootstrap sample $\pi \in \Pi$, we estimate nonparametric bootstrap p -values under the null hypothesis, $H_0 : |\hat{\tau}_{\pi}^{LT,pre}| = 0$:

$$\hat{p} = 1 + \frac{1}{\Pi + 1} \sum_{\pi \in \Pi} I \{ |\hat{\tau}_{\pi}^{LT,pre}| > |\hat{\tau}^{LT,pre}| \}. \quad (14)$$

The estimated p -values, reported in Table SM-2, exceed the standard level of statistical significance except for the Swiss cluster when the initial placebo treatment period is two time periods away from the actual initial treatment period. These results provide indirect evidence that the estimated effects are truly attributable to the Schengen and FoM policies instead of other policy changes or spurious errors during the same period.

We also leverage the post-treatment period—when the true “treatment” effect is null—to compare the matrix completion estimator with difference-in-differences and synthetic control method. Specifically, we randomly select $N/2$ treated units and impute their missing values following a given initial placebo treatment period. We evaluate the matrix completion estimator in terms of the average root mean squared error (RMSE), comparing actual versus imputed values, and specify a model without covariates to facilitate comparability with the alternative estimators, which do not rely on covariates.

Figure 3 reports the estimators’ performance under simultaneous or staggered treatment adoption settings. The x axis is the ratio of the initial placebo treatment time to the number of periods in the placebo data, so higher values represent more training data.




Fig. 3. RMSE for retrospective imputation for placebo treated units on post-treatment data, under staggered treatment adoption (left) or simultaneous treatment adoption (right). Outcome: % working in border region, conditional on employment. Error bars are calculated by taking average RMSE ± 1.96 times the standard error of the average RMSE across 100 runs. The x axis is the ratio of the placebo initial treatment time to the number of periods in the placebo data. Estimates are jittered across the x axis to avoid overlap. *Method:* \triangle , Matrix com.; \circ , Diff.-in-diff.; $+$, Synth. con.

In most settings, the matrix completion estimator outperforms difference-in-differences, both of which significantly outperform the synthetic control method. This bolsters the usage of the matrix completion methods in our application.

6. Conclusion

Motivated by a setting with later-treated and always-treated units, we formulate a novel identification strategy of retrospective counterfactual imputation to predict potential outcomes under treatment for the later-treated in the pre-treatment period. We treat these potential outcomes as missing and impute them using a matrix completion estimator.

Our methodological contributions are three-fold: first, we introduce a framework for causal inference in the retrospective setting, comparing always-treated and later-treated units; second, we extend the matrix completion estimator by propensity- and elapsed-time weighting the objective function to correct for differences in the observed covariate and unobserved fixed effects distributions, and the elapsed time since treatment between groups. Third, we propose a procedure for imputing endogenous covariate values when reconstructing potential outcomes under treatment.

We apply this framework to evaluate the effect of two key European institutions, namely the Schengen Agreement and the Freedom of Movement, on cross-border employment in Europe. While previous literature has studied the labour market effects of

these institutions on natives in the host regions (Naguib, 2019; Basten and Siegenthaler, 2019; Aepli and Kuhn, 2019; Åslund and Engdahl, 2019; Beerli et al., 2020), our analysis focuses on sending regions.

By relying on ELFS data aggregated at the regional level for the period 2005-2019, we provide evidence that these institutions have almost doubled the probability of working beyond the border for individuals living in the border regions of Eastern Europe. Interestingly, the effect appears to be driven mostly by the Freedom of Movement, which abolished the requirement of work permits, while the abolition of border controls is not the main driver. Our results are consistent with the results of previous studies which showed that opening the border increases the share of non-resident workers in host countries (e.g., Dustmann et al., 2016; Beerli et al., 2020). The findings are robust to whether difference-in-differences or the synthetic control method is used in place of the matrix completion estimator. Placebo tests provide indirect evidence that the key identification assumption of exogeneity is not violated and that the estimator is consistent.

Acknowledgements

We thank Eurostat for granting access to the European Labour Force Survey (project RPP 9/2017-LFS). Supported by the Luxembourg National Research Fund (C17/SC/11700060). Poulos and Li's research is partially supported by the National Science Foundation under Grant DMS-1638521 to the Statistical and Applied Mathematical Sciences Institute. The views expressed herein are those of the authors and do not necessarily reflect those of the institutions they represent. Code and microdata access instructions are available at: <https://github.com/jvpoulos/schengen>.

References

- Abadie, A. (2005) Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, **72**, 1–19.
- 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**, 493–505.
- (2015) Comparative politics and the synthetic control method. *American Journal of Political Science*, **59**, 495–510.
- Abadie, A. and Gardeazabal, J. (2003) The economic costs of conflict: A case study of the Basque Country. *The American Economic Review*, **93**, 113–132.
- Aeppli, M. and Kuhn, A. (2019) Open labor markets and firms’ substitution between training apprentices and hiring workers. *IZA discussion papers*, Institute of Labor Economics (IZA).
- Ahrens, A., FitzGerald, J. and Lyons, S. (2020) Commuting across the Irish border. *Economics Letters*, **190**, 109060.
- Amjad, M., Shah, D. and Shen, D. (2018) Robust synthetic control. *The Journal of Machine Learning Research*, **19**, 802–852.
- Angrist, J. D. and Pischke, J.-S. (2009) *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton NJ: Princeton University Press.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W. and Wager, S. (2019) Synthetic difference in differences. *Tech. rep.*, National Bureau of Economic Research.
- Ashenfelter, O. (1978) Estimating the effect of training programs on earnings. *The Review of Economics and Statistics*, **60**, 47–57.
- Ashenfelter, O. and Card, D. (1985) Using the longitudinal structure of earnings to estimate the effect of training programs. *The Review of Economics and Statistics*, **67**, 648–660.
- Athey, S., Bayati, M., Doudchenko, N., Imbens, G. and Khosravi, K. (2017) Matrix completion methods for causal panel data models. *arXiv:1710.10251*.
- Athey, S. and Imbens, G. (2018) Design-based analysis in difference-in-differences settings with staggered adoption. *arXiv:1808.05293*.
- Bai, J. (2003) Inferential theory for factor models of large dimensions. *Econometrica*, **71**, 135–171.
- (2009) Panel data models with interactive fixed effects. *Econometrica*, **77**, 1229–1279.
- Bartz, K. and Fuchs-Schündeln, N. (2012) The role of borders, languages, and currencies as obstacles to labor market integration. *European Economic Review*, **56**, 1148–1163.

- Basten, C. and Siegenthaler, M. (2019) Do immigrants take or create residents' jobs? Evidence from free movement of workers in Switzerland. *The Scandinavian Journal of Economics*, **121**, 994–1019.
- Beerli, A., Ruffner, J., Siegenthaler, M. and Peri, G. (2020) The abolition of immigration restrictions and the performance of firms and workers: Evidence from Switzerland. *Working Paper 25302*, National Bureau of Economic Research.
- Beine, M., Bourgeon, P. and Bricongne, J.-C. (2019) Aggregate fluctuations and international migration. *The Scandinavian Journal of Economics*, **121**, 117–152.
- Bello, P. (2019) Exchange rate effects on cross-border commuting: Evidence from the Swiss–Italian border. *Journal of Economic Geography*.
- Belloni, A., Chernozhukov, V., Fernández-Val, I. and Hansen, C. (2017) Program evaluation and causal inference with high-dimensional data. *Econometrica*, **85**, 233–298.
- Ben-Michael, E., Feller, A. and Rothstein, J. (2018) The augmented synthetic control method. *arXiv:1811.04170*.
- Borjas, G. J. (2001) Does immigration grease the wheels of the labor market? *Brookings Papers on Economic Activity*, **2001**, 69–133.
- Callaway, B. and Sant'Anna, P. (2018) Difference-in-differences with multiple time periods and an application on the minimum wage and employment. *arXiv:1803.09015v2*, 1–47.
- Candès, E. J. and Recht, B. (2009) Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, **9**, 717.
- Card, D. and Krueger, A. B. (1994) Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*, **82**, 772–793.
- Carvalho, C., Masini, R. and Medeiros, M. C. (2018) ArCo: An artificial counterfactual approach for high-dimensional panel time-series data. *Journal of Econometrics*, **207**, 352–380.
- Chen, N. and Novy, D. (2011) Gravity, trade integration, and heterogeneity across industries. *Journal of International Economics*, **85**, 206 – 221.
- Davis, D. and Gift, T. (2014) The positive effects of the Schengen agreement on European trade. *The World Economy*, **37**, 1541–1557.
- Davison, A. C. and Hinkley, D. V. (1997) *Bootstrap Methods and their Application*, vol. 1. Cambridge: Cambridge University Press.
- Ding, P. and Li, F. (2019) A bracketing relationship between difference-in-differences and lagged-dependent-variable adjustment. *Political Analysis*, **27**, 605–615.
- Doudchenko, N. and Imbens, G. W. (2016) Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. *arXiv:1610.07748*.

- Dustmann, C., Schönberg, U. and Stuhler, J. (2016) Labor supply shocks, native wages, and the adjustment of local employment. *The Quarterly Journal of Economics*, **132**, 435–483.
- Egger, P. and Lassmann, A. (2015) The causal impact of common native language on international trade: Evidence from a spatial regression discontinuity design. *Economic Journal*, 699–745.
- Elsner, B. (2013) Does emigration benefit the stayers? Evidence from EU enlargement. *Journal of Population Economics*, **26**, 531–553.
- Felbermayr, G., Gröschl, J. and Steinwachs, T. (2018) The trade effects of border controls: Evidence from the European Schengen agreement. *Journal of Common Market Studies*, **56**, 335–351.
- Gobillon, L. and Magnac, T. (2016) Regional policy evaluation: Interactive fixed effects and synthetic controls. *Review of Economics and Statistics*, **98**, 535–551.
- Gobillon, L., Selod, H. and Zenou, Y. (2007) The mechanisms of spatial mismatch. *Urban Studies*, **44**, 2401–2427.
- Hazlett, C. and Xu, Y. (2018) Trajectory balancing: A general reweighting approach to causal inference with time-series cross-sectional data. Available at: <https://ssrn.com/abstract=3214231>.
- Heckman, J. J., Ichimura, H. and Todd, P. E. (1997) Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The Review of Economic Studies*, **64**, 605–654.
- Hirano, K. and Imbens, G. W. (2001) Estimation of causal effects using propensity score weighting: An application to data on right heart catheterization. *Health Services and Outcomes Research Methodology*, **2**, 259–278.
- Kain, J. (1968) Housing segregation, negro employment, and metropolitan decentralization. *The Quarterly Journal of Economics*, **82**, 175–197.
- Kim, K. I. I. S. and Wang, E. (2018) Matching methods for causal inference with time-series cross-sectional data. Available at: <https://ssrn.com/abstract=3214231>.
- Lechner, M. (2010) The estimation of causal effects by difference-in-difference methods. *Foundations and Trends in Econometrics*, **4**, 165–224.
- Li, F., Morgan, K. L. and Zaslavsky, A. M. (2018) Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, **113**, 390–400.
- Mathä, T. and Wintr, L. (2009) Commuting flows across bordering regions: A note. *Applied Economics Letters*, **16**, 735–738.
- Mazumder, R., Hastie, T. and Tibshirani, R. (2010) Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, **11**, 2287–2322.

- Mundell, R. A. (1961) A theory of optimum currency areas. *The American Economic Review*, **51**, 657–665.
- Naguib, C. (2019) Estimating the heterogeneous impact of the free movement of persons on relative wage mobility. Available at: <https://ssrn.com/abstract=3344222>.
- Neyman, J. (1923) On the application of probability theory to agricultural experiments. *Annals of Agricultural Sciences*, **51**. Reprinted in Splawa-Neyman et al. (1990).
- Niebuhr, A., Granato, N., Haas, A. and Hamann, S. (2012) Does labour mobility reduce disparities between regional labour markets in Germany? *Regional Studies*, **46**, 841–858.
- Parenti, A. and Tealdi, C. (2019) Does the implementation of the Schengen agreement boost cross-border commuting? Evidence from Switzerland. *IZA Discussion Papers 12754*, Institute of Labor Economics (IZA).
- Persyn, D. and Torfs, W. (2015) A gravity equation for commuting with an application to estimating regional border effects in Belgium. *Journal of Economic Geography*, **16**, 155–175.
- Politis, D. N. and White, H. (2004) Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, **23**, 53–70.
- Recht, B. (2011) A simpler approach to matrix completion. *Journal of Machine Learning Research*, **12**, 3413–3430.
- Rubin, D. B. (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, **66**, 688.
- Sandner, M. and Wassmann, P. (2018) The effect of changes in border regimes on border regions crime rates: Evidence from the Schengen treaty. *Kyklos*, **71**, 482–506.
- Splawa-Neyman, J., Dabrowska, D. M., Speed, T. P. et al. (1990) On the application of probability theory to agricultural experiments. *Statistical Science*, **5**, 465–472.
- Strezhnev, A. (2018) Semiparametric weighting estimators for multi-period difference-in-differences designs. Available at: https://www.antonstrezhnev.com/s/generalized_did.pdf.
- Xu, Y. (2017) Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, **25**, 57–76.
- Åslund, O. and Engdahl, M. (2019) Open borders, transport links, and local labor markets. *International Migration Review*, **53**, 706–735.

