DISCUSSION PAPER SERIES

DP15178

COVID-19 ACROSS EUROPEAN REGIONS: THE ROLE OF BORDER CONTROLS

Matthias Eckardt, Kalle Kappner and Nikolaus Wolf

INTERNATIONAL TRADE AND REGIONAL ECONOMICS

ISSN 0265-8003

COVID-19 ACROSS EUROPEAN REGIONS: THE ROLE OF BORDER CONTROLS

Matthias Eckardt, Kalle Kappner and Nikolaus Wolf

Discussion Paper DP15178
Published 17 August 2020
Submitted 12 August 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Trade and Regional Economics

COVID-19 ACROSS EUROPEAN REGIONS: THE ROLE OF BORDER CONTROLS

Abstract

Attempts to constrain the spread of Covid-19 included the temporal reintroduction of travel restrictions and border controls within the Schengen area. While such restrictions clearly involve costs, their benefits have been disputed. We use a new set of daily regional data of confirmed Covid-19 cases from the respective statistical agencies of 18 Western European countries. Our data starts with calendar week 10 (starting 2nd March 2020) and extends to calendar week 17 (ending 26th April 2020), which allows us to test for treatment effects of border controls. We use Poisson models with fixed effects and controls for the stringency of national measures, as well as a Bayesian spatio-temporal specification using an integrated nested Laplace approximation (INLA) to take unobserved spatio-temporal heterogeneity into account. Both approaches suggest that border controls had a significant effect to limit the pandemic.

JEL Classification: C33, I18, R23

Keywords: COVID-19, Border effects, INLA

Matthias Eckardt - m.eckardt@hu-berlin.de
Humboldt University Berlin

Kalle Kappner - kalle.kappner@hu-berlin.de
Humboldt University Berlin

Nikolaus Wolf - nikolaus.wolf@wiwi.hu-berlin.de
Humboldt University Berlin and CEPR

Covid-19 across European Regions: the Role of Border Controls

Matthias Eckardt*, Kalle Kappner†, Nikolaus Wolf‡

August 2020

Abstract

Attempts to constrain the spread of Covid-19 included the temporal reintroduction of travel restrictions and border controls within the Schengen area. While such restrictions clearly involve costs, their benefits have been disputed. We use a new set of daily regional data of confirmed Covid-19 cases from the respective statistical agencies of 18 Western European countries. Our data starts with calendar week 10 (starting 2nd March 2020) and extends to calendar week 17 (ending 26th April 2020), which allows us to test for treatment effects of border controls. We use Poisson models with fixed effects and controls for the stringency of national measures, as well as a Bayesian spatio-temporal specification using an integrated nested Laplace approximation (INLA) to take unobserved spatio-temporal heterogeneity into account. Both approaches suggest that border controls had a significant effect to limit the pandemic.

*HU Berlin. E-mail: m.eckardt@hu-berlin.de
†HU Berlin. E-mail: kalle.kappner@hu-berlin.de
‡HU Berlin and CEPR. E-mail: nikolaus.wolf@wiwi.hu-berlin.de

1 Introduction

The outbreak of the Covid-19-pandemic led to a massive return of the nation state. National governments around the world took far-reaching measures to control the spread of the disease, from decisions to close shops, restaurants and schools to a full-blown lock-down of public life. In Europe, the crisis is a fundamental challenge to the principles of the European Union, notably solidarity, policy coordination and free movement across national borders.

In this paper we focus on the temporal reintroduction of national border controls within the Schengen area. Border controls are obviously costly, even though it is hard to estimate such costs. For example, Felbermayr, Gröschl, and Steinwachs (2016) suggest that the controls imposed in the wake of the refugee crisis in 2015 amounted to a reduction of EU28 real GDP by over 12 billion Euro (or 0.10%) per year. For the current crisis, Meninno and Wolff (2020) argue that these costs must be substantially larger, given the increase in cross-border commuting since 2015. The major question is whether the (temporary) closure of borders had benefits that could justify such costs. According to Nicolas Schmit, Jobs and Social Rights Commissioner the closure of borders, such as the border between Germany and Luxembourg was *just a reflex, which doesn't add anything to health security.* (cf. The New York Times, 17$^{\text{th}}$ April 2020). But maybe border controls did help to contain Covid-19?

Attempts to conduct an encompassing cost-benefit analysis for policy measures to contain Covid-19 are difficult and contentious (see for example Gros (2020), Broughel and Kotrous (2020)). Instead we focus on one crucial aspect: to what extent did the reintroduction of border controls reduce the number of infections? Arguably, if we would not find any systematic evidence for the effectiveness of controls on limiting the spread of the disease it would be hard to justify them. Our approach is to collect daily data at the level of European regions within

nation states. Our data starts roughly one week before the introduction of border controls, which allows us to test for treatment effects. Based on two quite different approaches, we find that border controls reduced the number of Covid-19 cases significantly, by about 6% to 25%, depending on the specification. Our paper is related to several attempts to test for the effects of control measures on the spread of Covid-19, including Hartl, Wälde, and Weber (2020) who use a simple linear trend model to test for the effect of the public shutdown on the spread of Covid-19 in Germany, and Mitze et al. (2020), who use a synthetic control method to test for the effect of introducing face masks.

The rest of our paper is organized as follows: we first provide a short survey of the spread of Covid-19 across European regions and the introduction of border controls. Next, we describe our data and our main estimation strategy using a PPML estimator. We then discuss the robustness of our findings using a Bayesian count specification implemented through the integrated nested Laplace approximation introduced by Rue, Martino, and Chopin (2009) to capture unobserved heterogeneity in the spatial structure of our data, and conclude.

2 The spread of Covid-19 and border controls

According to the WHO, the pandemic reached Europe on 25th January 2020 with first cases reported in France, followed by Germany on 28th January and Italy on 30th January 2020 (WHO Situation Reports, 5, 8, 11, 2020). By 1st March 2020, there were 1,457 confirmed cases (with 31 deaths) in the European region (WHO definition), spreading rapidly. One month later, by 1st April 2020 there were 463,677 cases and 30,085 deaths, most of which occurred within the European Union (WHO SR 41 and 72, 2020). Italy introduced the first large-scale measures on 21st February 2020 with a lock-down of initially 11 municipalities, next 4 provinces and on 8 March for the whole country (Maurice

et al. 2020). At the European level, the Commission mobilized additional funds for research on 1st February, extended on 24th February, set-up a "response team" on 2nd March 2020 and suggested to relax the fiscal rules of the Stability and Growth Pact. Most importantly, the European Central Bank announced on 18th March 2020 measures of unprecedented scale to support economic activity in the Euro-Zone. While this was a very quick response to the pandemic, notably if compared to the financial crisis, it still left the impression that the European Union was caught off guard and unprepared.

The main reason for this perception is that only national and regional governments could take immediate action against the disease, because health care falls within the competence of the member states, according to the Treaty of Lisbon (Brehon 2020). Among the first actions taken by national governments were the reintroduction of border controls. On 11th March 2020 Austria introduced controls on the land border with Italy, followed by Hungary on the border with Austria and Slovenia on 12nd March, Switzerland on the continental borders with Italy on 13th March. Within a few weeks most countries in the Schengen area have reintroduced border controls, with few exceptions such as the border between the Netherlands and Germany.

3 Data, method and main results

While we still lack the data to fully understand the dynamics of the pandemic, we can approximate the spread of the disease by looking at confirmed Covid-19 cases across regions within nation-states and over time. To this end, we collected daily regional data of confirmed new Covid-19 cases from the respective statistical agencies of 18 Western European countries[1] from calendar week 10

[1]Those are Andorra, Austria, Belgium, Denmark, France, Finland, Germany, Ireland, Italy, Liechtenstein, Luxembourg, the Netherlands, Norway, Portugal, Spain, Switzerland, Sweden and UK – in other words: All of Western Europe except for the isolated island of Iceland.

(starting 2nd March 2020) to calendar week 17 (ending 26th April 2020). For France, we approximate daily new cases by the number of hospitalized per day and region and rescale them to the number of confirmed cases with national data. We aggregate the data to the level of 213 roughly equally sized sub-national regions that closely follow the definition of European NUTS2 regions. Finally, we rescale all regional daily case counts to match the national totals reported by Johns Hopkins University's Coronavirus Resource Center.

Figure 1 shows the spread of Covid-19 across European regions during this period in terms of new confirmed cases per two-week period. As we can see, the spread of the disease shows a very strong regional pattern, while the effects of national borders are not obvious. During the first two weeks of our sample (panel 1a), incidence was concentrated in Norther Italy and parts of Spain. Calendar weeks 12 and 13 (panel 1b) saw a quick spread, in many cases across national borders (with interesting exceptions – see France-Spain). During this period, border controls were enacted. Weeks 14 and 15 (panel 1c) saw the apex of new cases, with incidence all over the map. Calendar weeks 16 and 17 (panel 1d) already saw a reduction in new cases as most countries surpassed the height of their incidence curve during the first wave of 2020.

The spatial patterns in the raw daily case data are hard to interpret due to differences in national testing and reporting schemes, differences in data quality, and possible confounders. For the remainder of the paper, we thus condition our data on region fixed effects, and country-specific time fixed effects.[2]

How did borders and border controls matter? Figure 2 illustrates the extent of border controls and two definitions of our treatment group. We also list the date of border control enactments for each country pair. Note that the

[2]Since we include region fixed effects, we do not need to control for population or GDP and absolute case numbers are just as informative as case rates. Also, country-specific time trends should absorb major differences in testing and reporting behavior, as well as changes in containment policies such as social distancing.

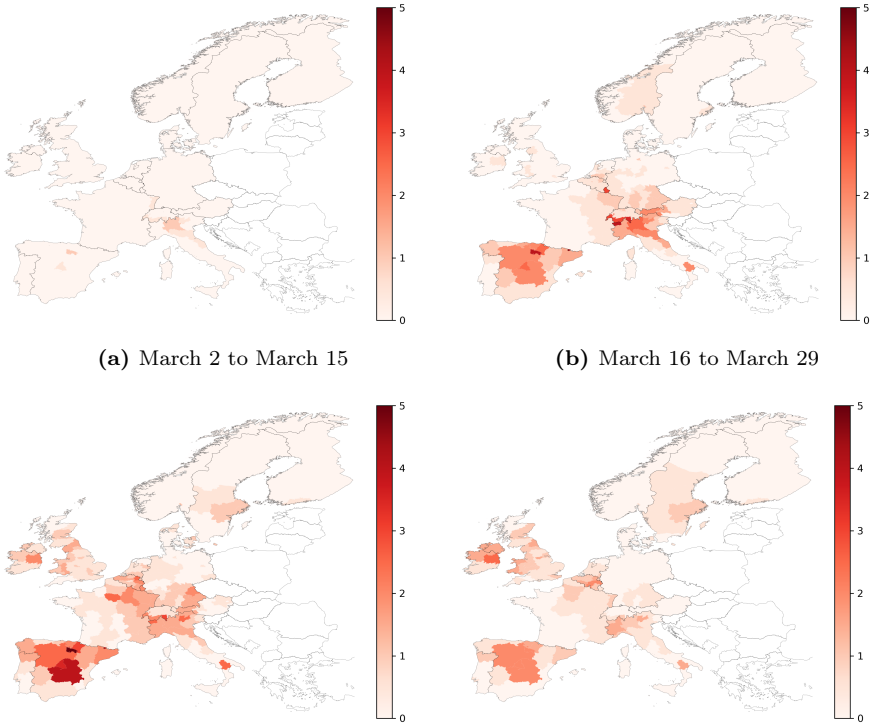

(a) March 2 to March 15

(b) March 16 to March 29

(c) March 30 to April 12

(d) April 13 to April 26

Figure 1: New confirmed Covid-19 cases per 1,000 inhabitants in specified calendar weeks. See text for details on the data sources.

we always assume a symmetric impact of border controls: If France controls its border with Germany, both French and German border regions are treated, even if Germany technically introduces border controls only later (or not at all).

To test for the role of national borders for the spread of the disease we estimate a series of difference-in-differences regressions of the form

$$I_{r,n,d} = exp(\alpha_i + \beta_{d,(n)} + \gamma D_{r,d} + \epsilon_{r,d}) \tag{1}$$

where $I_{r,n,d}$ are new cases in region r, in country n on day d, α_i and $\beta_{d,(n)}$ are region and time fixed effects (which in some specifications are allowed to be country-specific), $D_{r,d}$ is a dummy for regions affected by border controls. Since our sample extends well before the onset of border controls, this dummy is time-varying. γ is our coefficient of interest, capturing the causal effect of border controls on daily cases.[3]

We present results for two different definitions of the treatment: a broad definition and a narrow definition. In our first set of results, we distinguish regions located at controlled borders from those not located at controlled borders. In this specification γ will pick up variation between the two groups over time that is not explained by average (or country-specific) time effects, depending on the specification. However, borders might have mattered for some regions much more than for others in the first place. Intuitively, the introduction of travel restrictions should have mattered for regions that experienced intense cross-border commuting beforehand, such as regions on the border between Belgium and Germany, but much less (or not at all) for border regions with little cross-border commuting.

Therefore, in an alternative specification we consider only those border

[3]We note that due to the staggered treatment timing, the estimated coefficients present weighted averages of the underlying group-time average treatment effects that are likely to underestimate the actual average treatment effect (Callaway and Sant'Anna 2019).

regions as treated that experienced intense cross-border commuting before the introduction of border controls: in this case, the treated regions are all border regions with an above-mean share of their workforce (> 0.9 %) commuting to a workplace across a national border in 2019. For example, 30 % of the workforce of the Belgian region Luxembourg and 11.3 % of the workforce of bordering French Lorraine were cross-border commuters in 2019, hence both regions belong to our intensity-based treatment definition. In contrast, Spanish Aragon and bordering French Midi-Pyrénées both had no significant cross-border commuting in 2019, hence they are excluded from the intensity-based treatment definition.

Moreover, we account for the possibility of a time lag in the effect of border controls. In one specification we assume that controls can have an immediate effect on the spread of the disease. In a second specification we assume that controls have an effect only with a time lag of at least one week, following Lauer et al. (2020).

Before we consider the findings of our difference-in-differences regressions, we need to discuss whether regions in both, treatment and control groups followed similar trends before the treatment. If not, our results might be spurious as they would pick up differences in trends rather than effects from some treatment (Bertrand, Duflo, and Mullainathan 2004). A major challenge in our setting is the staggered introduction of border controls across European regions, together with the relatively limited number of pre-treatment observations. Hence, we lack the data to formally test for common trends. Instead, we rely on a graphical analysis as shown in figure 3. Note that we show the narrow definition of the treatment group here, based on border regions with above average commuting before the treatment and country-specific time effects.

The key takeaway from figure 3a is that treated regions showed a somewhat higher level of (conditional) confirmed Covid-19 cases compared to control regions

before the (staggered) introduction of controls, but that trends in both groups were similar. After the introduction of border controls, the levels in treated regions converge to those in control regions, i.e. there is evidence for a trend break sometime after the treatment. In figure 3b we plot the exponentiated γ-coefficients over time, where an exponentiated coefficient significantly lower than 1 indicates a reduction in cases. The effect of being a treated border region becomes consistently (and significantly) negative only after the introduction of border controls, whereas we see no clear pattern before the treatment(s).

Table 1 shows our first set of results, using a PPML estimator. This method consistently estimates average marginal effects even if the data is not (conditionally) Poisson-distributed. The first three models allow for immediate effects of border controls. The last three models "shift" the onset of border controls by 7 days to take the incubation time and reporting delay into account.

The point estimates for γ in all models suggest that border controls led to a reduction in the number of reported Covid-19 cases. In table 1 we transform them to report the percentage change in cases relative to the control group together with the p-values. We see that the size of the effect is much larger, once we use the narrow, intensity-based treatment definition (compare columns 1, 2 and 4, 5). Intuitively, the introduction of border controls mattered much more for regions with a substantial number of cross-border commuters beforehand, compared to border regions with little or no commuting. The effects become statistically significant once we use an intensity-based definition of the treatment together with country-specific time effects.[4] We report heteroscedasticity-robust standard errors clustered on the region level.

How to read these coefficients? Our preferred specification is shown in column 6, where we control for region effects and country-specific time effects,

[4]The number of observations decreases because three countries composed of a single region (Andorra, Liechtenstein, Luxembourg) drop out of the sample. Neither the drop of these three regions nor the introduction of country-specific time effects itself are driving our results.

	Instant impact			7-day lagged impact		
	(1) Location	(2) Intensity	(3) Intensity	(4) Location	(5) Intensity	(6) Intensity
Border control	-7.04	-40.83	-42.93***	-18.31	-30.92	-25.08**
	(0.75)	(0.21)	(0.00)	(0.22)	(0.29)	(0.02)
Fixed effects	region day	region day	region day*country	region day	region day	region day*country
Observations	11928	11928	11462	11928	11928	11462
Regions	213	213	210	213	213	210
Pseudo R2	0.81	0.81	0.91	0.81	0.81	0.91

Table 1: The reported coefficients are the percentage changes in cases relative to the control group due to border controls, i.e. $(e^{\gamma} - 1) * 100$. Numbers in parentheses are p-values.

use a narrow definition of the treatment group and test for lagged effects. A value of -25.08 (column 6) means that daily cases are reduced by 25.08 % due to border controls. Using the mean number of daily new cases across our whole sample (92.387), this amounts to a reduction by $92.387 * 0.251 = 23.189$, about 23 cases per day less for an "average" region.

4 Robustness

A major challenge in our setting is the spatial nature of our data and the fact that we cannot control for temporal variation of local containment policies that differed from national policies. As indicated by figure 1, the spread of Covid-19 followed a particular spatial pattern, which is not well captured by our PPML model. In figure 4 below we provide a measure of spatial correlation in our dependent variable, conditional on region and country-time effects. To this end, we first compute from table 1, col. 6 the spatial lag of the residuals for each region. Next, we group the (standardized) residuals into 100 equally sized bin and plot each bin's mean against the average spatial lag in that bin. Given the strong spatial patterns seen in figure 1 this suggests that our PPML method

helps to reduce spatial correlation in the residuals, but does not eliminate it.

Moreover, it is likely that our PPML estimation overstates the true treatment effect from border controls, because the higher incidence in some border regions before the treatment (see figure 3a) would have led to the implementation of local containment policies before they were introduced at the national level. Our country-specific time effects would thus not control for such local measures, which should bias our estimated treatment effects upwards.

To control for both the spatial structure and temporal dynamics of the data and also account for potentially unobserved spatio-temporal heterogeneity, we specify a Bayesian spatial-temporal count data model which we implement using the INLA formalism for Bayesian inference in latent Gaussian models. This is provided by the *R-INLA* project (`www.r-inla.org`) using the capacities of the R environment (R Core Team 2020). Bayesian methods have become widespread in applied epidemiology and public heath research, notably due to the development of Markov Chain Monte Carlo methods (MCMC) and, more recently, the development of more computationally efficient alternatives including INLA and variational Bayes approaches, see Blangiardo et al. (2013), Bakka et al. (2018). To this end, we construct a first-order spatial lag structure over the regional entities defined through a contiguity-based spatial weighting matrix which we use to set up a conditional autoregressive specification of the spatial effect (Besag 1972; Besag 1974). Further, to allow for potentially unobserved spatial heterogeneity, we additionally included a spatial random effect assuming an iid Gaussian distribution. By this, we allow for both structured and unstructured spatial effects such that the model also absorbs unobserved spatial heterogeneity (Fahrmeir, Kneib, and Lang 2004). Again, treating the number of new confirmed Covid-19 cases as outcome, the spatio-temporal count model includes time effects, the distance from a continental border, the share of

	Global effect	Std. Dev.	Implied percentage change
Border control	-0.06	0.01	-6.20***
Constant	-4.18	0.88	

Table 2: Main coefficients from the INLA model. See text for details on the specification.

commuters in the workforce, a time-varying dummy for border controls and an offset.[5] Table 2 shows the main parametric results, figure 5 shows the distribution of the structured and unstructured component of spatial effects.

The main finding from this exercise is that even if we allow for a very flexible form of unobserved spatio-temporal effects, we still find that border controls reduced the number of confirmed Covid-19 cases significantly. According to the INLA approach, the introduction of border controls reduced the number of daily new cases by roughly 6 %, compared to 25 % suggested by the PPML estimator.

5 Conclusion

The temporal reintroduction of border controls within the Schengen area helped to contain the spread of Covid-19. While such restrictions clearly involve costs, their benefits have been disputed. In this paper we used a new set of daily regional data of confirmed Covid-19 cases from the respective statistical agencies of 18 Western European countries, running from calendar week 10 (starting 2 March 2020) to calendar week 17 (ending 26 April 2020). This allowed us to test for treatment effects of border controls. Based on a PPML estimator with region fixed effects and country-specific time effects, we show that border controls were associated with a 25% reduction in daily cases. Importantly, we show that border controls mattered only for regions with a substantial number of cross-border commuters prior to the crisis. As a robustness check, we use a Bayesian

[5] We model both the temporal dependence and the duration of the border controls by a second-order random walk specification of the effect.

INLA approach to take unobserved spatio-temporal heterogeneity into account, for example due to local containment policies that might have differed from nation-wide measures. With this we find smaller, but still significant effects in the area of 6 %. We conclude that the temporal introduction of border controls was certainly costly, but made an important contribution to contain the spread of Covid-19. At the same time it is likely that better policy coordination at the European level could have generated these benefits at lower economic (and political) costs, for example if based on a closer monitoring of cross-border commuting flows. We leave this question for further research.

References

Bakka, Haakon et al. (2018). "Spatial modeling with R-INLA: A review". In: *WIREs Computational Statistics* 10:e1443.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004). "How Much Should We Trust Differences-In-Differences Estimates?" In: *Quarterly Journal of Economics* 119.1, pp. 249–275.

Besag, Julian (1972). "Nearest-neighbor systems and the auto-logistic model for binary data". In: *Journal of the Royal Statistical Society, Series B, Statistical Methodology* 34.1, pp. 75–83.

— (1974). "Spatial interaction and the statistical analysis of lattice systems". In: *Journal of the Royal Statistical Society, Series B, Statistical Methodology* 36, pp. 192–236.

Blangiardo, Marta et al. (2013). "Spatial and spatio-temporal models with R-INLA". In: *Spatial and Spatio-temporal Epidemiology* 4, pp. 33–49.

Brehon, Nicholas-Jean (2020). "The European Union and the Coronavirus". In: *Fondation Robert Schuman Policy Paper* 553.7 April.

Broughel, James and Michael Kotrous (2020). "The Benefits of Coronavirus Suppression. A Cost-Benefit Analysis of the Response to the First Wave of COVID-19". In: *Mercatus Center, Mercatus Working Paper*.

Callaway, Brantly and Pedro H. C. Sant'Anna (2019). "Difference-in-Differences with Multiple Time Periods". In: *Working Paper*. URL: `http://conference.iza.org/conference_files/EVAL_2019/28549.pdf`.

Fahrmeir, Ludwig, Thomas Kneib, and Stefan Lang (2004). "Penalized Structured Additive Regression for Space-Time Data: A Bayesian Perspective". In: *Statistica Sinica* 14.3, pp. 731–761.

Felbermayr, Gabriel, Jasmin Gröschl, and Thomas Steinwachs (2016). "Trade costs of border controls in the Schengen area". In: *VOXEU* 27 April 2016.

Gros, Daniel (2020). "The great lockdown: was it worth it?" In: *CEPS Policy Insights* No 2020–11, May 2020.

Hartl, Tobias, Klaus Wälde, and Enzo Weber (2020). "Measuring the impact of the German public shutdown on the spread of Covid-19". In: *Covid Economics* 1.3 April 2020, pp. 25–32.

Lauer, Stephen A. et al. (2020). "The Incubation Period of Coronavirus Disease 2019 (Covid-19) From Publicly Reported Confirmed Cases: Estimation and Application". In: *Annals of Internal Medicine* 172.9, pp. 577–582.

Maurice, Eric et al. (2020). "Covid-19: European Responses, a complete picture". In: *Fondation Robert Schuman Policy Paper* 28 April 2020.

Meninno, Raffaella and Guntram Wolff (2020). "As the Coronavirus spreads, can the EU afford to close its borders?" In: *VOXEU* 28 February 2020.

Mitze, Timo et al. (2020). "Face masks considerably reduce COVID-19 cases in Germany: A synthetic control method approach". In: *COVID Economics: Vetted and Real-Time Papers* 27, pp. 74–103.

R Core Team (2020). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: `https://www.R-project.org/`.

Rue, Håvard, Sara Martino, and Nicolas Chopin (2009). "Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71.2, pp. 319–392.

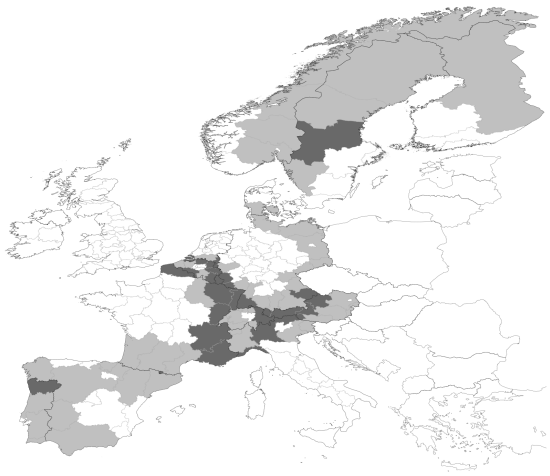

country pair	enactment
AND-ESP	March 17
AND-FRA	March 17
AUT-CHE	March 11
AUT-DEU	March 11
AUT-ITA	March 11
AUT-LIE	March 11
AUT-HUN	March 12
AUT-SVK	March 12
AUT-CZE	March 12
BEL-DEU	March 16
BEL-FRA	March 20
BEL-LUX	March 20
BEL-NLD	March 20
CHE-DEU	March 13
CHE-FRA	March 13
CHE-ITA	March 13
DEU-CZE	March 14
DEU-DNK	March 12
DEU-FRA	March 16
DEU-LUX	March 16
DEU-POL	March 15
ESP-PRT	March 16
ESP-FRA	March 17
FIN-NOR	March 16
FIN-RUS	March 16
FIN-SWE	March 19
FRA-ITA	March 14
FRA-LUX	March 14
ITA-SVN	March 14
NOR-SWE	March 16

Figure 2: Border controls in European regions. The map on the left shows our treatment and control groups. Light-gray regions (■) are located at controlled borders. Dark-gray regions (■) are a subset of the former with high levels of cross-border commuting in 2019. The table on the right lists the dates of border control enactment for all country pairs, including East European countries not included in our sample.

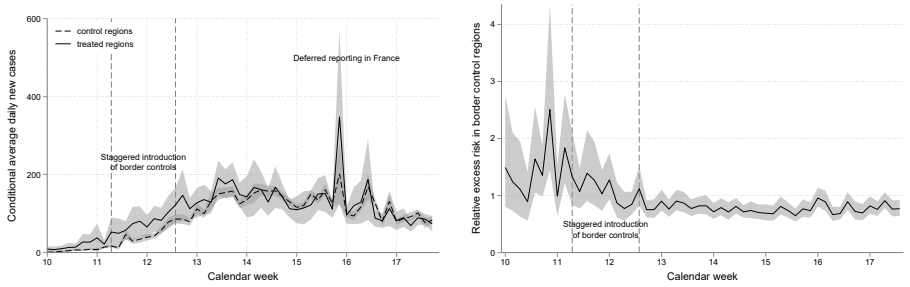

(a) Groupwise conditional means **(b)** Periodwise treatment effects

Figure 3: Visual checks for parallel trends. Panel 3a plots average daily new cases in the treatment and control groups, conditional on day and region fixed effects. Panel 3b shows (exponentiated) coefficients of the treatment group dummy for each day, conditional on country-day and region fixed effects In both panels, gray areas show the 10 % confidence interval for robust standard errors clustered at the region level. Also note that the "France spike" seen in panel 3a does not show up in panel 3b, because it is absorbed by the France-specific time effects.

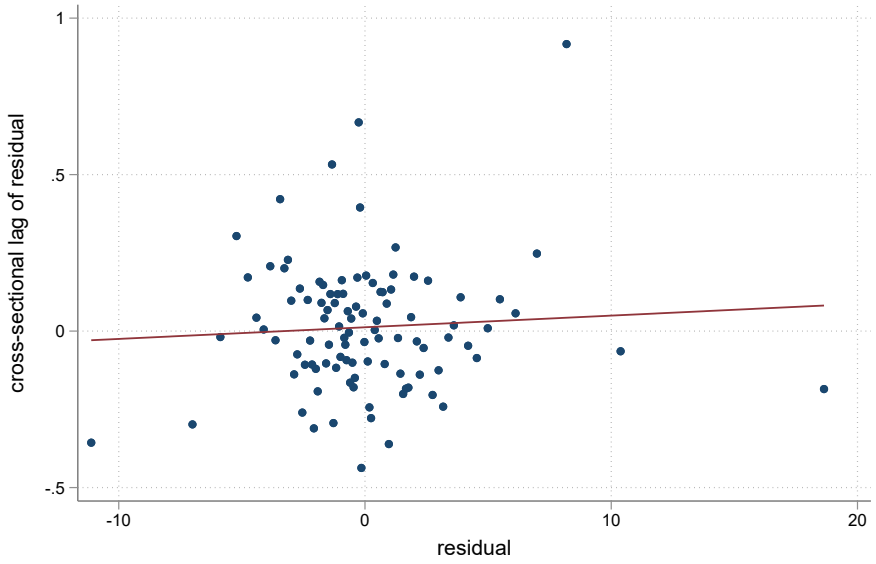

Figure 4: Binscatter of residuals against their cross-sectional lag. We first compute the cross-sectional spatial lag of the residuals from table 1, col. 6, counting as neighbors all regions with a shared border. We then group the (standardized) residuals into 100 equally sized bins and plot each bin's mean against the average spatial lag in that bin.

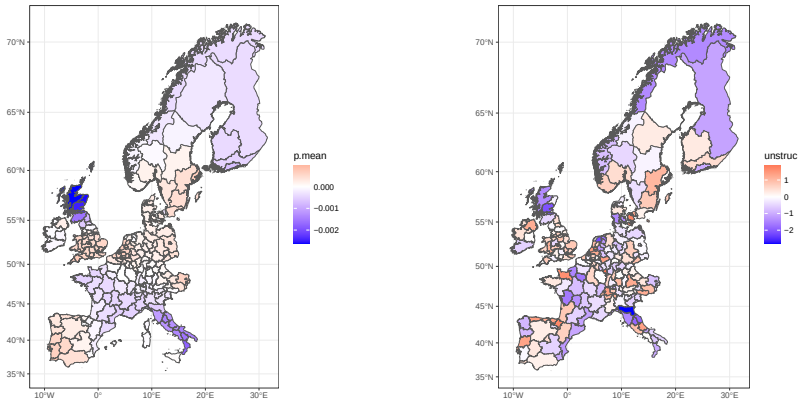

(a) Structured heterogeneity **(b)** Unstructured heterogeneity

Figure 5: Spatial heterogeneity