The lockdown effect: A counterfactual for Sweden*

Benjamin Born, Alexander M. Dietrich, Gernot J. Müller

July 2020

Abstract

We quantify the lockdown effect: the extent to which a lockdown limits the spread of COVID-19 infections. For this purpose we focus on Sweden—one of the few countries without a lockdown—and use synthetic control techniques to develop a counterfactual lockdown scenario. Based on a donor pool of European countries, we construct a control unit that behaves just like Sweden in terms of infection dynamics before it is put under lockdown for 8.5 weeks. The outcome in the control unit approximates the counterfactual lockdown scenario for Sweden. Three findings stand out. First, at the end of the lockdown period, COVID-19 infections and deaths in Sweden would have been reduced by one half and one third, respectively. Second, the lockdown effect starts to materialize with a delay of 3 to 4 weeks only. Third, we analyze Google mobility reports and find that actual mobility adjustments in Sweden are similar to those in the control unit—in both cases there is evidence for social distancing, but more so in the control unit.

Keywords: COVID-19, lockdown, counterfactual, synthetic control unit, infections, voluntary social restraint, Google mobility reports

JEL-Code: C4, F6, I18

*Born: Frankfurt School of Finance & Management and CEPR, b.born@fs.de, Dietrich: University of Tübingen, alexander.dietrich@uni-tuebingen.de. Müller: University of Tübingen and CEPR, gernot.mueller@uni-tuebingen.de. This is a significantly revised and extended version of a pre-print published in issue 16 of COVID Economics: vetted and real-time papers under the title “Do lockdowns work? A counterfactual for Sweden”. We thank seminar participants in an online Global Health Economics Seminar of the Heidelberg Institute of Global Health and at the Swedish Riksbank, as well as Christian Bayer, Jop de Vrieze, Bruno Frey, Hartmut Egger, Zeno Enders, Philip Jung, Stefan Kohler, Willi Kohler, Kris Nimark, Dominik Papies, Julia Peters, Martin Schlegel, Harald Uhlig, and Adalbert Winkler for very useful comments and discussions. Lennart Fischer provided valuable research assistance. The usual disclaimer applies.

1 Introduction

As the COVID-19 pandemic spreads across the globe, policy makers face a trade-off. By imposing a lockdown, they may limit the spread of COVID-19 infections and deaths. But a lockdown also entails severe economic and social costs (Bonaccorsi et al., 2020; Frey and Osterloh, 2020). In order to inform policy makers, it is essential to quantify both the costs and the benefits of lockdowns. In this study we focus on the latter and quantify the extent to which a lockdown limits the spread of COVID-19 infections and deaths.

A major challenge in quantifying this lockdown effect is that infection dynamics are bound to change over time, even in the absence of a lockdown—at least for two reasons. First, the basic Susceptible-Infectious-Recovered (SIR) model assumes a constant infection probability, but predicts that infection growth slows over time as the pool of infected and susceptible people shrinks (Kermack and McKendrick, 1927). Second, people may adjust their behavior in the face of infection risk and restrain their social interactions voluntarily. This, too, reduces infection growth even if no lockdown is imposed. Several studies neglect these aspects and attribute any change in the growth rate of infections to lockdown measures or, more broadly speaking, to non-pharmaceutical interventions, NPI, for short (Flaxman et al., 2020; Zhanga et al., 2020).[1]

In this study, we account for changing infection dynamics as we quantify the lockdown effect on the basis of a purely empirical approach. Specifically, we focus on Sweden: it stands out from its European peers in that its government opted against a lockdown even though the exposure to COVID-19 in Sweden was not systematically different from the rest of Europe. Still, the Swedish authorities merely advised—rather than ordered—citizens to adjust their behavior in the face of the pandemic. For instance, people were told "to avoid unnecessary traveling and social events, to keep distance to others, and to stay at home" if they had any symptoms. In addition, those above age 70 were "advised to avoid social contact" and "visits

[1] The number of contributions that seek to measure the impact of COVID-19 in general and that of NPIs in particular is growing fast. Rather than providing a comprehensive overview, we discuss the contributions most relevant for our study in the main text below.

to retirement homes” were banned (Krisinformation, 2020).

In order to quantify the lockdown effect, we benchmark the actual developments in Sweden against a counterfactual lockdown scenario. To approximate this scenario, we construct a synthetic control unit on the basis of a donor pool of countries that actually imposed a lockdown (Abadie et al., 2010). In constructing the control unit, we make sure that it resembles Sweden in terms of infection dynamics *before* the lockdown. The difference between the control unit and Sweden *during* the lockdown provides us with a measure of the lockdown effect.

Our main result is that a lockdown during the period March 18 to May 17 would have reduced the number of COVID-19 infections in Sweden by one half and the number of deaths by one third. These numbers pertain to the end of the lockdown period (May 17) in the synthetic control unit. Second, the lockdown impacts infection dynamics with a delay of 3 to 4 weeks only. In order to rationalize this result, we analyze Google mobility reports and find a profound change of actual mobility patters in Sweden, even if no lockdown is imposed: people travel less and spend more time at home or in parks. The extent of this adjustment is strong, but not quite as strong as in the control unit. This finding suggests that voluntary social restraint limits infection growth to a considerable extent—in line with recent work that augments the basic SIR model to account for behavioral adjustments in the face of infection risk. (Eichenbaum et al., 2020; Farboodi et al., 2020; Krueger et al., 2020). It may also rationalize why the lockdown effect becomes manifest with a considerable delay only. At the same time, the size of the lockdown effect suggests that there is a non-trivial infection externality—people fail to internalize the costs they impose on others as they become infectious and, hence, voluntary social restraint fails to deliver sufficient social distancing.

The remainder of the paper is organized as follows. Section 2 introduces our approach and our data. The main results are presented and discussed in Section 3. A final sections summarizes our findings and highlights a number of caveats. The appendix provides details on our data sources and results for alternative specifications.

2 The approach

We develop a counterfactual scenario for Sweden in order to quantify the extent to which a lockdown would have helped to contain the spread of COVID-19 infections and deaths. For this purpose, we construct a synthetic control unit for Sweden. Ideally, it behaves just like Sweden before the lockdown such that any difference during the lockdown may be attributed to the lockdown in the control unit.

2.1 The donor pool

The *control unit* is a weighted average of the countries in the *donor pool*. To ensure a high degree of homogeneity between Sweden and the control unit, we restrict the donor pool to Norway and western EU countries with more than 1 million inhabitants. In total it includes 13 countries. Table 1 provides details on the countries in the donor pool, including the most important containment measures and the timing of the lockdown. The latter is determined on the basis of various sources, detailed in Appendix A.6. The first lockdown was imposed in Italy on March 9, the last in the Netherlands on March 24. These lockdowns typically involved the closing of non-essential shops as well as a ban on gatherings of more than two people. In some instances, the ban applies only to gatherings of 10 people or more. In Sweden only gatherings of more than 50 people were banned.

As a comprehensive measure, we also compute the average stringency of the lockdown based on data from the Coronavirus Government Response Tracker (Hale et al., 2020). As can be seen in the second-to-last column of Table 1, there is some variation in lockdown stringency across countries. These data show that there has also been a government response to the Coronavirus in Sweden, but—consistent with the premise of our analysis—the index number is considerably lower than for the other countries in the donor pool.

In our analysis, we use daily observations for the number of infections up to May 1, 2020. Our data source for infections (as well as for the number of COVID-19 deaths) is the Johns

Table 1: Donor pool countries and Sweden: lockdown measures and dates

Country	Lockdown Start	Lockdown End	Containment Measures	Day 1	Days to Lockdown	Days in Lockdown	Lockdown Stringency	% of control unit
Austria	16.03.	01.05.	non-essential shops closed, ban on gatherings of more than 5 people	29.02.	16	46	83	00.0
Belgium	18.03.	11.05.	non-essential shops closed, ban on gatherings of more than 2 people, stay-at-home order	03.03.	15	54	81	00.0
Denmark	18.03.	11.05.	non-essential shops closed, ban on gatherings of more than 50 people	03.03.	15	54	70	26.1
Finland	16.03.	01.06.	government agencies closed, ban on gatherings of more than 10 people, stay-at-home advice	01.03.	15	77	56	19.2
France	17.03.	11.05.	non-essential shops closed, ban on gatherings of more than 2 people, stay-at-home order	29.02.	17	55	90	00.0
Germany	23.03.	06.05.	non-essential shops closed, ban on gatherings of more than 2 people	01.03.	22	44	73	00.0
Greece	23.03.	11.05.	non-essential shops closed, ban on gatherings of more than 10 people, stay-at-home-order	05.03.	18	46	82	00.0
Ireland	28.03.	08.06.	non-essential shops closed, stay-at-home-order	04.03.	24	72	87	00.0
Italy	09.03.	18.05.	non-essential shops closed, stay-at-home-order	22.02.	16	70	85	00.0
Netherlands	24.03.	11.05.	non-essential shops closed, ban on gatherings	02.03.	22	48	79	39.1
Norway	13.03.	01.06.	restaurants, bars closed, ban on gatherings of more than 5 people (24.03)	28.02.	14	80	66	14.6
Portugal	19.03.	01.06.	no shops closed, government agencies closed, stay-at-home advice	06.03.	13	74	79	01.0
Spain	14.03.	08.06.	non-essential shops closed, stay-at-home-order	01.03.	13	86	78	00.0
Sweden	-	-	No Lockdown imposed, ban on gatherings of more than 50 people	29.02.	-	-	42	-

Notes: donor pool includes Western-EU countries with population size of at least one million and Norway. Day 1 is the day when the number of infections surpasses a threshold of one infection per one million inhabitants. Sources for lockdown dates and details are provided in the supplementary material appendix. Lockdown ends when shops were reopened, except for Finland and Norway (restaurant reopening). Lockdown stringency is average value during lockdown period, data source: Hale et al., 2020. For Sweden, average is for counterfactual lockdown period from March 18 to May 17. "% of control unit" pertains to baseline (specification A), rounded to first digit.

Hopkins University (Dong et al., 2020), see Appendix A.1 for details. In order to assess the impact of the lockdown, it is essential to ensure that infection dynamics are comparable across countries prior to the lockdown. As the virus arrived at possibly different dates in each country, one should not compare the infection dynamics on a calender-day basis. Therefore, we initialize observations for each country using a common reference point: day 1 is when the number of infections surpasses a threshold of one infection per one million inhabitants. Day 1 varies from country to country, see Table 1. For instance, in Sweden it is February 29, in Norway February 28, and in Denmark March 3. We find that it takes countries at least 13 days since day 1 to impose a lockdown (see again Table 1).

2.2 The control unit

For the construction of the control unit, we require it to track the infection dynamics in Sweden during the first 13 days as closely as possible, that is, before any country in the donor pool imposed a lockdown. Since the number of infections are initially very low, we target the log of infections rather than the level. In this way, we make sure that the early observations within the 13-day window play a non-negligible role for the construction of the control unit. In addition, we require the control unit to be comparable to Sweden in terms of population size and in terms of the urban population share because these factors may play an important role for infection dynamics. In total, we target 15 observations in order to construct the control unit: log infections at daily frequency within the 13 day window prior to the first lockdown, population size, and the urbanization rate.

Formally, we construct the control unit by selecting weights on the countries in the donor pool for which we obtain the best match between the control unit and Sweden for the 15 target observations. Formally, we let $\mathbf{x_1}$ denote the (15×1) vector of observations in Sweden and let $\mathbf{X_0}$ denote a (15×13) matrix with observations in the countries included in the donor pool. Finally, we let $\mathbf{w}$ denote a (13×1) vector of country weights w_j, $j = 1, \ldots, 13$. Then, the control unit is defined by $\mathbf{w}^*$ which minimizes the following mean squared error:

$$(\mathbf{x_1} - \mathbf{X_0}\mathbf{w})'\mathbf{V}(\mathbf{x_1} - \mathbf{X_0}\mathbf{w}) \ , \tag{1}$$

subject to $w_j >= 0$ for $j = 1, \ldots, 13$ and $\sum_{j=1}^{13} w_j = 1$. In this expression, $\mathbf{V}$ is a (15×15)

symmetric and positive semidefinite matrix. Here, $\mathbf{V}$ is a weighting matrix assigning different relevance to the characteristics in $\mathbf{x_1}$ and $\mathbf{X_0}$. Although the matching approach is valid for any choice of $\mathbf{V}$, it affects the weighted mean squared error of the estimator (Abadie et al., 2010). We choose a diagonal $\mathbf{V}$ matrix such that the mean squared prediction error of the outcome variable (and the covariates) is minimized for the pre-treatment period (Abadie et al., 2010; Abadie and Gardeazabal, 2003).

3 Results

In this section we first present our main result regarding the lockdown effect and show that is robust across alternative specifications. Afterwards we shed further light on our main result by comparing mobility patters in Sweden and the control unit.

3.1 The lockdown effect

We benchmark actual infection dynamics in Sweden against a counterfactual that is approximated by the outcome in the control group. As we minimize expression (1), we find that five countries from the donor pool receive a non-negligible weight in the control unit: the Netherlands (39.1%), Denmark (26.1%), Finland (19.2%), Norway (14.6%), and Portugal (1%). The weights of the other countries are smaller than 0.01%, see also the right column in Table 1. Note that the number of countries with non-negligible weight is not restricted by our procedure and may vary across specifications. The weighted average outcome across the countries in the control unit approximates the counterfactual outcome against which we benchmark actual developments in Sweden. On average, it takes 18 days after day 1 before the a 60-day lockdown is imposed. Its stringency is 70 and hence well above the actual value for Sweden. Since day 1 in Sweden is February 29, 2020, we assume that the counterfactual lockdown runs from March 18 to May 17.

The left panel of Figure 1 shows infection dynamics in Sweden (blue solid line) and the

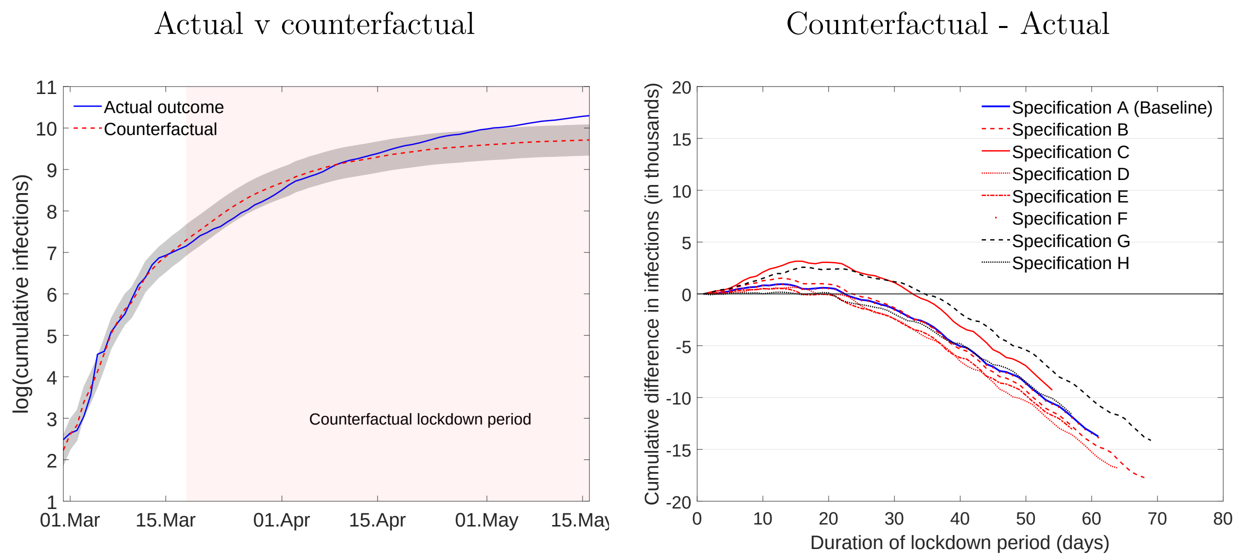

Figure 1: COVID-19 infections in Sweden. Left: actual outcome (blue solid line) v counterfactual (red dashed line) in logs. Right: counterfactual - actual outcome since start of the lockdown period in thousands. *Notes:* Counterfactual approximated by outcome of control unit, see Table 1 for country weights of specification A (baseline). Gray shaded area: two standard deviations of difference between infections in Sweden and control unit during the first 13 days. Pink shaded area: lockdown period (March 18–May 17). Specifications B-F: alternative control units based on restricted donor pools. Panel G: day 1 determined on the number of deaths, Panel H: control unit determined by matching number of deaths instead of log infections.

control unit that approximates the counterfactual outcome for Sweden (red dashed line). The vertical axis measures cumulative infections in logs, the horizontal axis represents calendar days. By construction, the control unit tracks infection dynamics closely during the matching period (day 1 to 13). We also obtain a good fit for population size (10.175 millions in Sweden and 10.187 in the control unit) and for the urbanization rate (0.874 and 0.877, respectively). Yet our procedure does not constrain infection dynamics during the lockdown period, indicated by the pink shaded area. The gray shaded area, in turn, represents two standard deviations of the difference between log infections in Sweden and the control unit during the matching period. It puts deviations between the actual outcome and the counterfactual into perspective, but is not a formal test of statistical significance. Importantly, for the end of the lockdown period, we observe that log infections in Sweden are considerably higher than in the lockdown counterfactual.

Table 2 reports the number of new infections and deaths (measured in absolute numbers rather than logs) during the lockdown period, contrasting the developments in Sweden and the

Table 2: The lockdown effect

	Infections		Deaths	
	Actual	Counterfactual	Actual	Counterfactual
A) Baseline	28,864	15,108	3,669	2,438
(lockdown ends May 17)		-48%		-34%
B) W/o Denmark in donor pool	32,180	14,457	3,988	2,479
(lockdown ends May 24)		-55%		-38%
C) W/o Finland in donor pool	24,731	15,453	3,213	1,724
(lockdown ends May 09)		-55%		-46%
D) W/o Netherlands in donor pool	29,274	12,464	3,692	2,329
(lockdown ends May 18)		-57%		-37%
E) W/o Norway in donor pool	26,630	13,583	3,450	2,269
(lockdown ends May 13)		-49%		-34%
F) W/o Portugal in donor pool	28,864	15,005	3,669	2,450
(lockdown ends May 13)		-48%		-33%
G) Day 1 determined using number of deaths	32,085	17,951	3,986	2,183
(lockdown ends May 23)		-44%		-45%
H) Matching number of deaths	27,143	15,511	3,518	2,330
(lockdown ends May 14)		-43%		-34%

Notes: new infections and deaths since start of lockdown. Counterfactual outcome and lockdown period depends on composition of the control unit. Panel A: baseline (see Figure 1). Panels B-F: alternative control units, for which each of the countries with non-negligible weight in baseline control unit, in turn, is excluded from donor pool. Panel G: day 1 determined on the number of deaths (see methods section), Panel H: control unit determined by matching number of deaths instead of log infections (see methods section).

counterfactual. Panel A shows numbers for the baseline specification: 28,864 new infections in Sweden and 15,108 in the counterfactual lockdown scenario since the start of lockdown. This is our main result: a lockdown would have reduced the number of infections in Sweden by 48%. At the same time, it would have reduced the number of deaths by 34% at the end of the lockdown period, from 3,669 to 2,438.

These results are based on a synthetic control unit defined by country weights that are plausible in the sense that Sweden's neighbours receive a lot of weight. Still, we verify that results do not depend on any specific country in the control group and report results for alternative specifications in the remaining panels of Table 2. Panels B to F show results once we construct alternative control units based on a restricted donor pool. Specifically, in this case, we exclude, in turn, each country that received a non-trivial weight in the baseline specification. We obtain five new control units that differ in terms of composition and also somewhat in the timing of the lockdown period and lockdown intensity. We report details in the appendix. Importantly, however, as Table 2 shows, we obtain similar results for the lockdown effect in each instance. Appendix A.5 reports details on each alternative specification, including the country weights.

Table 2 also reports results for specification G. In this case we consider an alternative criterion to specify day 1. Recall that we define as day 1 the day at which there is at least 1 infected person per 1 million inhabitants in the baseline specification. Since one may be worried that measurement error and/or mistakes in reporting have a large impact when the overall number of infections is still low, we consider an alternative approach to specify day 1. Specifically, for specification G we first determine the day at which the number of deaths exceeds one per 100,000 inhabitants. We then define day 1 as the day 30 days before that particular threshold. It turns out that results are not too different from the baseline. For instance, in specification G day 1 is February 26 in Sweden (rather than February 29 in the baseline). Given day 1, we again match log infections for the first 13 days (as well as country size and the urbanization rate). Because in some instances there are initially still zero infections, we add 1 infection to all observations before taking logs. For specification G, the control unit is composed of four countries: Denmark (3%), the Netherlands (31%), Norway (44%) and Portugal (22%).

In an additional robustness check, we match the total number of deaths instead of log infections (specification H). In this case we also start from the day at which a country suffers at least one death per 100,000 inhabitants. Next, we define day 1 as the day 10 days prior to that day. In this way we start matching observations later than in the other specifications because initially there are no deaths. We match the number of deaths for the first 13 days after day 1 as well as country size and the urbanization rate. As a result, our matching period overlaps somewhat with the lockdown period. This is inconsequential, however, because there is considerable delay before a lockdown may impact the number of deaths. In specification H, the control unit is composed of Denmark (45%), Finland (17%), and the Netherlands (38%).

A widely discussed shortcoming of the available data is that the number of reported infections is not independent of the number of tests, since infections may go undetected if symptoms are mild or even absent. Figure A.1 in the appendix shows that there where fewer tests conducted in Sweden relative to the control unit, but that the ratio of tests has been fairly stable during the lockdown period. Hence, it is unlikely that our results are significantly biased by different testing frequencies in Sweden and the control unit.

The right panel of Figure 1 shows the cumulative difference in infections between the counterfactual and the actual numbers for Sweden since the start of the lockdown. The blue solid line is the result for our baseline specification, the other lines pertain to alternative specifications B to H. A robust finding emerges across specifications: the lockdown effect emerges only with a considerable delay. In fact, initially Sweden is outperforming the control unit—the lockdown starts to slow down infections only after about 3 to 4 weeks. This is noteworthy because the incubation period for COVID-19 is on average only 5-6 days, and at most up to 14 days (WHO, 2020).

3.2 Social distancing

Our analysis shows, first, that a lockdown would have reduced the number of COVID-19 infections in Sweden by one half and the number of deaths by one third. These numbers pertain to the end of the lockdown period (May 17). Second, the lockdown impacts infection dynamics with a delay of 3 to 4 weeks only. In order to rationalize this result, two additional findings are helpful. We conjecture that it takes time for a lockdown effect to materialize because there is (voluntary) social distancing even in the absence of a lockdown. In order to explore this possibility, we compare actual mobility patterns in Sweden and the counterfactual on the basis of Google COVID-19 Community Mobility Reports (Google, 2020). Locations are classified according to six distinct categories: Work, Transit, Retail and Recreation, Grocery and Pharmacy, Parks, and Residential. The reports measure the change in the number and the length of stays at these locations relative to the median value of the same weekday between January 3 and February 6, 2020.

Figure 2 displays mobility dynamics for each category, contrasting once more actual data for Sweden (blue solid line) and for the counterfactual (red dashed line), approximated by the same control unit as above. We measure observations using five-day symmetric averages along the horizontal axis, and report the percentage change relative to the period in early 2020. As before, the pink shaded area indicates the lockdown period. Several findings stand out. First, we observe a pronounced decline of mobility in the top four panels. They provide a measure of activities associated with travel and work as well as shopping and dining in restaurants ("recreation"). At the same time, people spend more time in parks and at home (bottom panels). Second, the adjustment starts to take place about 10 days before the lockdown and, in line with our conjecture, it can be observed both in terms of actual developments in Sweden as well as for the counterfactual.[2] Last, we observe that while the adjustment of activities follows roughly the same pattern, it is more pronounced for the counterfactual.

A final statistic that is closely monitored as the pandemic unfolds is the reproduction number R. Intuitively, it measures the number of infections caused by one infected person. Values above (below) unity indicate that the number of infected people increases (decreases) over time. We compute R by a seven-day moving average of new infections divided by the seven-day moving average four days earlier (Robert-Koch-Institut, 2020). Figure 3 shows the results, both for Sweden as well as for the control unit. It is organized as the earlier figures and shows that R drops sharply well before the lockdown—both as far as actual developments in Sweden and the counterfactual are concerned. During the lockdown period, R fluctuates

[2]Consistent with these findings, it has been documented that COVID-19 has induced job seekers to reduce their search intensity and employers to reduce vacancy postings in the absence or prior to a lockdown (Hensvik et al., 2020; Kahn et al., 2020).

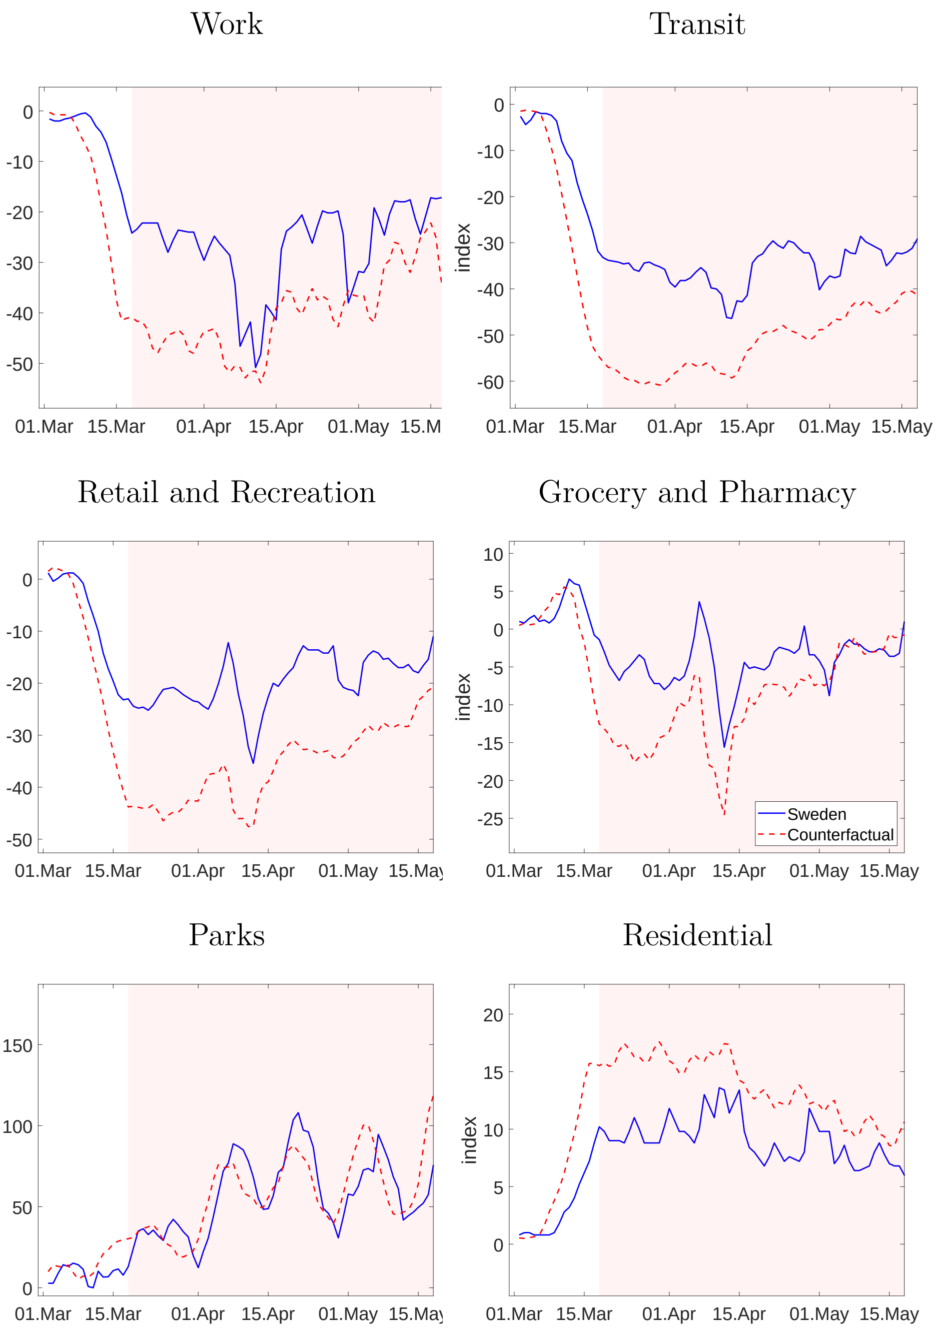

Figure 2: Mobility patterns in Sweden (blue solid line): actual outcome v counterfactual (red dashed line). Vertical axis: percentage change relative to median in early 2020. Horizontal axis: five-day symmetric moving average. Pink shaded area indicates lockdown period. Construction of counterfactual: see Figure 1 (Specification A, Baseline). Data source: Google mobility reports (Google, 2020).

around unity. However, there is an important difference across actual dynamics and the counterfactual: in the counterfactual, R drops somewhat below unity in the second half of

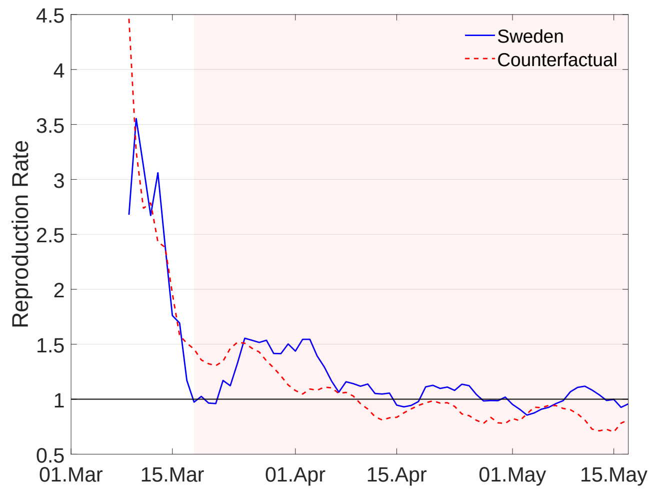

Figure 3: Actual reproduction R number in Sweden (blue solid line) v counterfactual (red dashed line) for specification A (baseline). Pink shaded area indicates lockdown period in control unit. R using actual infection data in Sweden as well as computed data for the counterfactual with lockdown. R rate constructed using the ratio of a 7-day average of newly reported infections relative to the 7-day average number of newly infected four days earlier. Construction of counterfactual: see Figure 1. Data source: Johns Hopkins University (Dong et al., 2020).

the lockdown period; instead, it fails do so in Sweden.

These results are consistent with the hypothesis that it takes time for the lockdown to make a difference because people adjust the behavior even in the absence of a lockdown: they reduce their social interactions in ways which limits the spread of the virus—there is, in other words, voluntary social restraint. This point emerges also clearly from a number of recent model-based contributions that have extended the basic SIR model to allow for peoples' (rational) adjustment of work, consumption and leisure activities in the face of infection risk. Model simulations illustrate that voluntary social restraint can have large effects, both in terms of limiting the maximum share of infections in the population and in terms of inducing an economic contraction (Eichenbaum et al., 2020). This behavioral response can also rationalize the observation that individuals across the United States reduced their social interactions well before lockdown restrictions were imposed (Farboodi et al., 2020). Importantly, the extent to which people adjust their behavior depends critically on the extent to which there are possibilities to substitute low-risk activities, such as working from home or take-away dining, for high-risk ones. There can be little doubt that actual economies offer a fairly rich array of substitutes for most goods and activities (Krueger et al., 2020).

That said, economic theory also identifies an infection externality: while people adjust their behavior to limit their exposure to the virus, they fail to internalize the costs they impose on others as they become infectious. This provides a rationale for government-imposed NPIs. In theory, the outcomes under laissez-faire and under the optimal policy can be quite different. In Sweden, instead, it takes a considerable amount of time for the difference to become visible—likely because lockdown policies in practise are fairly blunt and less well targeted than the optimal policy in model-based simulations. Still, over time, there is a lockdown effect—in line with the notion of a non-trivial infection externality (Chudik et al., 2020).

4 Conclusion

How effective are lockdowns in limiting the number of COVID-19 infections and deaths? Answering this question is difficult because most countries imposed a lockdown and we cannot directly observe the counterfactual outcome. Our study exploits the fact that Sweden was the only country among its European peers that did not impose a lockdown. Based on this peer group, we construct a synthetic control unit that approximates what would have happened, had Sweden imposed a lockdown. The counterfactual lockdown period runs from March 18 to May 17: at the end of this period we observe 48% fewer infections and 34% fewer deaths compared to what actually happened in Sweden. It takes about 3-4 weeks before the lockdown effect starts to materialize.

In concluding, we would like to point out a number of limitations of our study. First, we use data on COVID-19 infections and deaths even though there are serious issues related to measurement, not least the fact that the number of reported infections depends on the number of tests. Still, our analysis is based on the same data that informs public discussions and actual policy design. Second, we assume a macro perspective throughout and study the effect of “a” lockdown (with stringency 70). Clearly, specific lockdown measures may differ strongly in terms of effectiveness. For instance, it has been shown on the basis of an approach similar to ours that making face masks mandatory is fairly effective in limiting the spread of infections (Mitze et al., 2020).

Third, we focus on the short-term effects of the lockdown, that is, on the total effect up to the end of the counterfactual lockdown period. Clearly, some of the effects of the lockdown may unfold only after it has been lifted, that is, in the medium and long run. In particular, since a lockdown would have lowered the number of infections, it may lead to persistently lower infection growth. In the longer run, however, there may be the opposite effect because, in the absence of a lockdown, a larger fraction of the population will have gained immunity, thus limiting the impact of a “second wave” (Giesecke, 2020). We remain deliberately agnostic about these issues and focus exclusively on the short-run impact of the lockdown.

Fourth, there is the issue of external validity: we have developed a counterfactual for Sweden and cannot be sure that results carry over to other contexts and countries. In particular, we cannot rule out that the behavioral adjustment in Sweden was influenced by the fact that other countries in Europe imposed a lockdown. We note, however, that our results are similar to those obtained for California, the first state in the US to issue a shelter-in-place order, on the basis of an approach similar to ours (Friedson et al., 2020).

Last, we stress that our study quantifies benefits of a lockdown in terms of limiting COVID-19 infections and deaths. We find that these benefits are not trivial. Yet, since our analysis is altogether silent on the costs of a lockdown, the final verdict on lockdowns as a policy tool is still out. We would hope that our results inform a broad-based debate on the best policy response to the COVID-19 pandemic.

References

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

Abadie, Alberto and Javier Gardeazabal (2003). "The economic costs of conflict: a case study of the Basque country". *American Economic Review 93* (1), 113–132.

Bonaccorsi, Giovanni et al. (2020). "Economic and social consequences of human mobility restrictions under COVID-19". *Proceedings of the National Academy of Sciences 117* (27), 15530–15535.

Chudik, Alexander, M. Hashem Pesaran, and Alessandro Rebucci (2020). "Voluntary and mandatory social distancing: evidence on COVID-19 exposure rates from chinese provinces and selected countries". NBER Working Papers 27039.

Cori, Anne, Neil M. Ferguson, Christophe Fraser, and Simon Cauchemez (2013). "A New Framework and Software to Estimate Time-Varying Reproduction Numbers During Epidemics". *American Journal of Epidemiology 178* (9), 1505–1512.

Dong, Ensheng, Hongru Du, and Lauren Gardner (2020). "An interactive web-based dashboard to track covid-19 in real time". *Lancet Infect Dis 20* (5), 533–534.

Eichenbaum, Martin, Sergio Rebelo, and Mathias Trabandt (2020). "The macroeconomics of epidemics". NBER Working Paper 26882.

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer (2020). "Internal and external effects of social distancing in a pandemic". NBER Working Papers 27059.

Flaxman, Seth et al. (2020). "Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe". *Nature*.

Frey, Bruno S. and Margit Osterloh (2020). *Was sagt die Glücksforschung zu Corona?* Die Volkswirtschaft 6/2020.

Friedson, Andrew I, Drew McNichols, Joseph J Sabia, and Dhaval Dave (2020). "Did California's shelter-in-place order work? early coronavirus-related public health effects". NBER Working Paper 26992.

Giesecke, Johan (2020). "The invisible pandemic". *The Lancet 395* (10238).

Google (2020). *Google COVID-19 Community Mobility Reports*. `https://www.google.com/covid19/mobility/`. Accessed: 01.07.2020.

Hale, Thomas et al. (2020). "Variation in government responses to COVID-19". Working Paper BSG-WP-2020/032. Blavatnik School.

Hensvik, Lena, Thomas Le Barbanchon, and Roland Rathelot (2020). *Job search during the COVID-19 crisis*. CEPR Discussion paper 14748.

Kahn, Lisa B, Fabian Lange, and David G Wiczer (2020). “Labor demand in the time of COVID-19: evidence from vacancy postings and ui claims”. NBER Working Paper 27061.

Kermack, William Ogilvy and A. G. McKendrick (1927). “A contribution to the mathematical theory of epidemics”. *Procceedings of the Royal Society 115* (772), 700–721.

Krisinformation (2020). “Https://www.krisinformation.se/en/hazards-and-risks/disasters-and-incidents/2020/officialinformation- on-the-new-coronavirus”. *Online*.

Krueger, Dirk, Harald Uhlig, and Taojun Xie (2020). “Macroeconomic dynamics and reallocation in an epidemic”. Working Paper 20-015. PIER.

Mitze, Timo, Reinhold Kosfeld, Johannes Rode, and Klaus Wälde (2020). “Face masks considerably reduce COVID-19 cases in Germany: a synthetic control method approach”. *IZA Discussion Paper*.

Robert-Koch-Institut (2020). “Erläuterung der Schätzung der zeitlich variierenden Reproduktionszahl R”. Technical Note.

Roser, Max, Hannah Ritchie, Esteban Ortiz-Ospina, and Joe Hasell (2020). “Coronavirus pandemic (covid-19)”. *Our World in Data*. https://ourworldindata.org/coronavirus.

WHO (2020). *Coronavirus disease 2019 (COVID-19) situation report – 73*. `https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200402-sitrep-73-covid-19.pdf?sfvrsn=5ae25bc7_4`.

Worldbank (2019 Revision[a]). *Population, total, sp.pop.totl*.

——— (2019 Revision[b]). *Urban population (% of total population), sp.urb.totl.in.zs*.

Zhanga, Renyi, Yixin Li, Annie L. Zhangc, Yuan Wang, and Mario J Molina (2020). “Identifying airborne transmission as the dominant route for the spread of COVID-19”. *PNAS 26* (117), 14857–14863.

A Appendix

A.1 Data sources

In our analysis we rely on several data sources. First, we use daily observations for the number of COVID-19 infections and deaths up to June 15, 2020. Here our data source is the Johns Hopkins University (Dong et al., 2020). Observations are available at daily frequency. They are assembled using national (e.g. ministry of health, government) as well as international official sources (e.g. WHO, European Centre for Disease Prevention and Control). We show the raw time series in the appendix in Figure A.2. Second, we also use data for the frequency of testing (Roser et al., 2020). These data are not always available at daily frequency such that we have to resort to interpolation. In particular, we verify that the relative frequency of testing in Sweden and in our control unit did not change much during our sample period. Figure A.1 in the appendix shows the results. Third, our data for population size and urbanization rate is provided by the World Bank (Worldbank, 2019 Revision[a],[b]). Fourth, we rely on official sources in order to specify and date the lockdown measures in each country in the donor pool. We provide details in the appendix. We also compute the average stringency of the lockdown measures in each country using the Coronavirus Government Response Tracker (Hale et al., 2020). Last, we use Google COVID-19 Community Mobility Reports to measure mobility changes due to the pandemic (Google, 2020). They are available for each country in our donor pool and provide a measure for how long and how frequently certain types of locations are visited. Google collects location data in various ways using mobile phone positions (via mobile networks or GPS data), a user's IP address, search queries, or navigation requests. Google uses this information only if users actively agree to share their "Location History".

A.2 Reproduction Number

We compute the reproduction rate for Sweden and the counterfactual displayed in Figure 3 in the main text as described by Cori et al. (2013):

$$R_t = \frac{\bar{E}_t^7}{\bar{E}_{t-4}^7} = \frac{\sum_{k=t-6}^{t} E_k}{\sum_{k=t-6}^{t} E_{k-4}} \, . \tag{2}$$

Here, E_t denotes the number of newly infected on day t. $\bar{E}_t^j$ denotes a j-day moving average of infection numbers, from day $t-j+1$ to day t, such that $\bar{E}_t^j = \frac{1}{j}\sum_{k=t-6}^{t} E_k$. Following the method proposed by the Robert-Koch-Institut, 2020, the reproduction rate in our case is then given by a seven-day moving average (i.e., $j=7$) of new infections divided by the seven-day moving average four days earlier, in order to allow for an incubation period of four days. The reproduction rate R_t, as in (2) thus measures the average number of people infected between $t-7$ and t by a person who had itself been infected with COVID-19 between $t-11$ and $t-4$.

A.3 Ratio of Tests

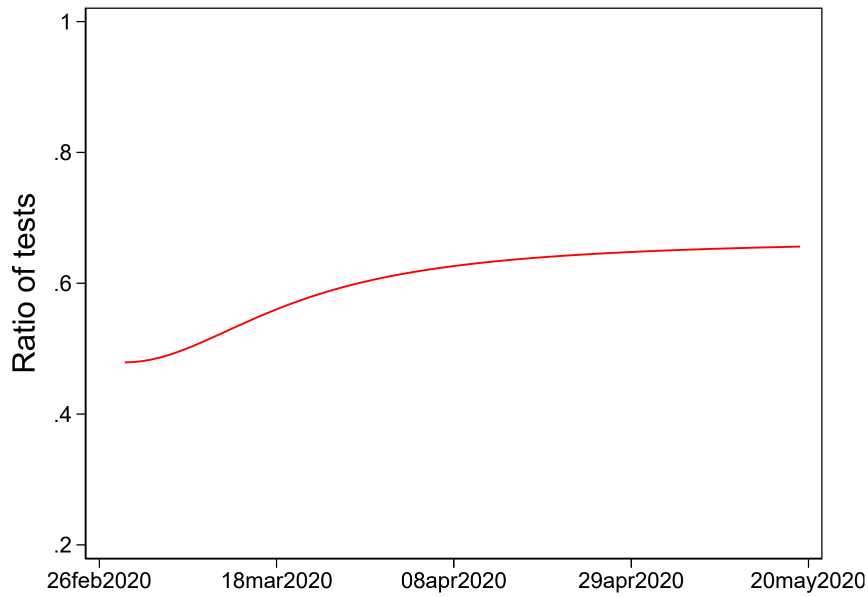

Figure A.1: Ratio of tests between Sweden and the counterfactual for the baseline specification.

In Figure A.1 we display the ratio of tests between Sweden and the control group. We use data provided by Roser et al., 2020. Since for both Sweden and the countries in the control group, data on tests conducted is only available infrequently, we estimate a quadratic trend for both the number of tests conducted in Sweden and the counterfactual. The ratio displayed is then based on the predicted trend.

A.4 Raw Data of Donor Pool Countries: Infections and Deaths

This section presents figures on the raw data series for deaths and infections of the countries in the donor pool, as made available by the Johns Hopkins University Dong et al., 2020. Data displayed runs from February 22 to June 15.

Infections

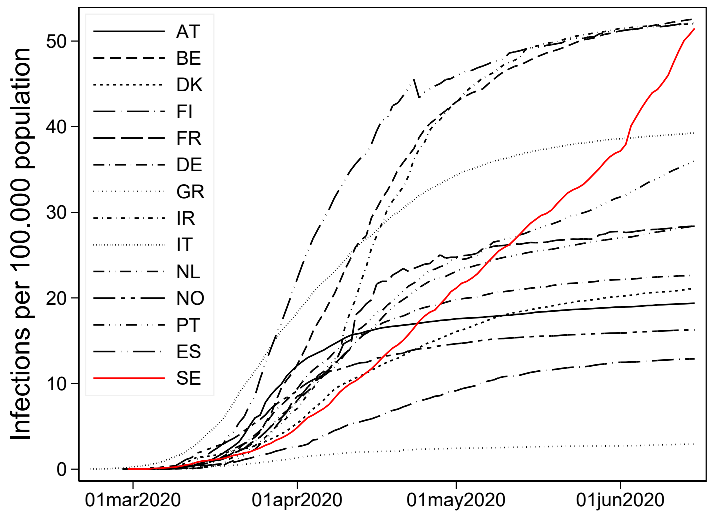

Deaths

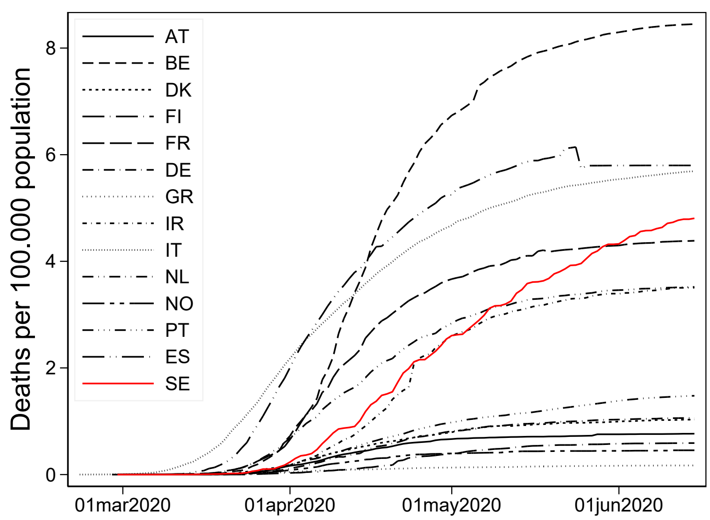

Figure A.2: Upper panel: Cumulative Infections per 100.000 population in countries of donor pool. Lower panel: Cumulative COVID-19 deaths per 100.000 population in countries of donor pool. Red line represents data for Sweden. Data source: Dong et al., 2020

A.5 Robustness

This section provides additional information and results on the specifications outlined in Table 1 in the main text. For each specification, weights of the countries in the donor pool for the counterfactual group as well as figures for log infections and cumulative deaths in levels are displayed.

Weights of Counterfactual Scenario

Table A.1: Country weights

Specification	A (Baseline)	B	C	D	E	F	G	H
Austria	00.0	00.0	38.3	00.0	02.6	00.0	00.0	00.0
Belgium	00.0	00.0	00.0	18.5	00.0	00.0	00.0	00.0
Denmark	26.1	NA	25.0	31.2	34.5	26.2	02.8	45.2
Finland	19.2	13.1	NA	14.0	21.2	19.4	00.0	17.0
France	00.0	00.0	00.0	00.0	00.0	00.0	00.0	00.0
Germany	00.0	00.0	00.7	04.3	00.0	00.0	00.0	00.0
Greece	00.0	00.0	00.0	04.5	05.5	00.0	00.0	00.0
Ireland	00.0	0.01	00.0	00.0	00.0	00.0	00.0	00.0
Italy	00.0	00.0	00.0	00.0	00.0	00.0	00.0	00.0
Netherlands	39.1	40.5	23.8	NA	36.1	39.5	31.2	37.7
Norway	14.6	46.4	12.2	27.4	NA	14.9	43.9	00.0
Portugal	01.0	00.0	00.0	00.0	00.0	NA	22.1	00.0
Spain	00.0	00.0	00.0	00.0	00.0	00.0	00.0	00.0

Table gives the weights in percent for each country in the optimized control group for each specification (A to H), as outlined in Table 1 in the main part. If table entry is NA country is excluded from the donor pool for the respective specification.

Specification A

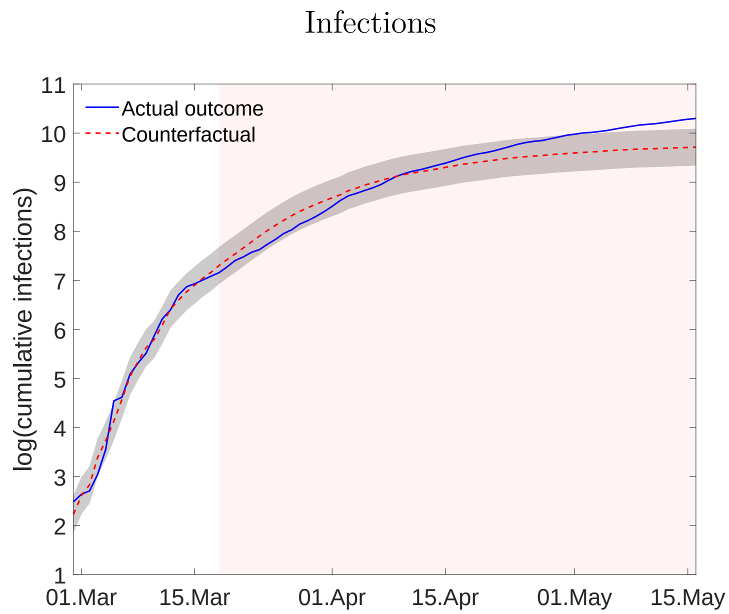

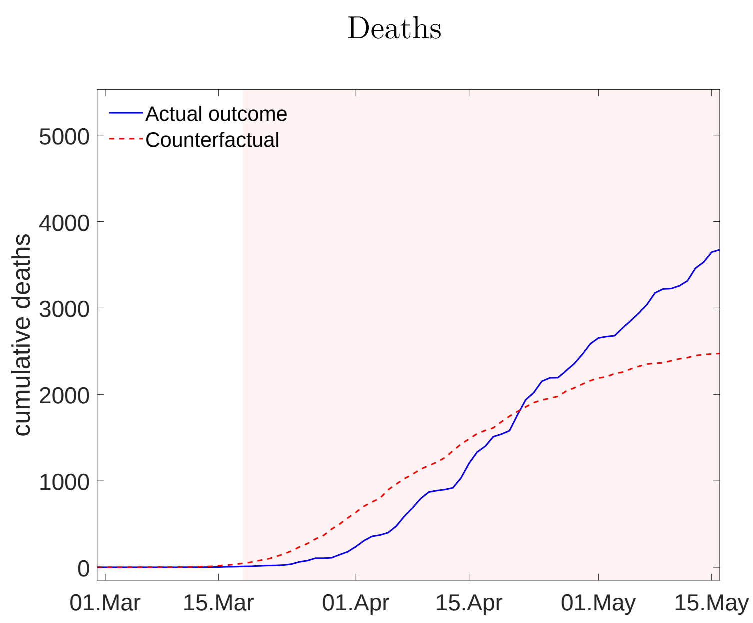

Figure A.3: Specification A: Baseline: Infections and Deaths relative to Sweden. Shaded area indicates Lockdown period in counterfactual.

Specification B

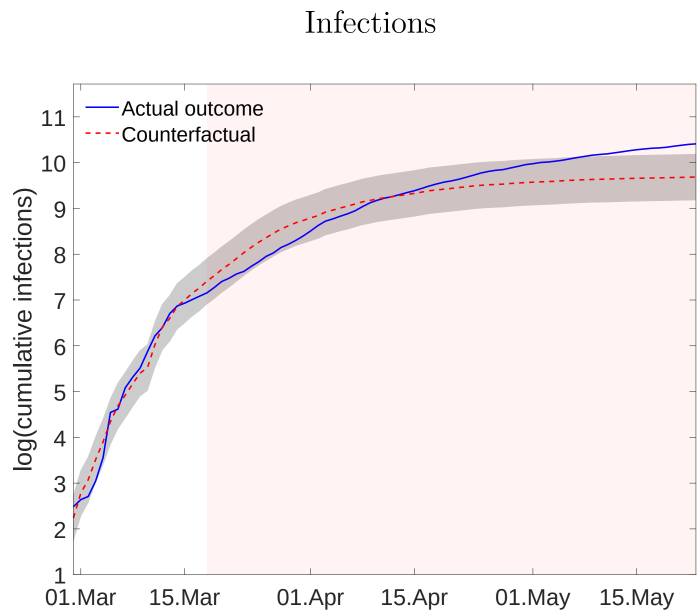

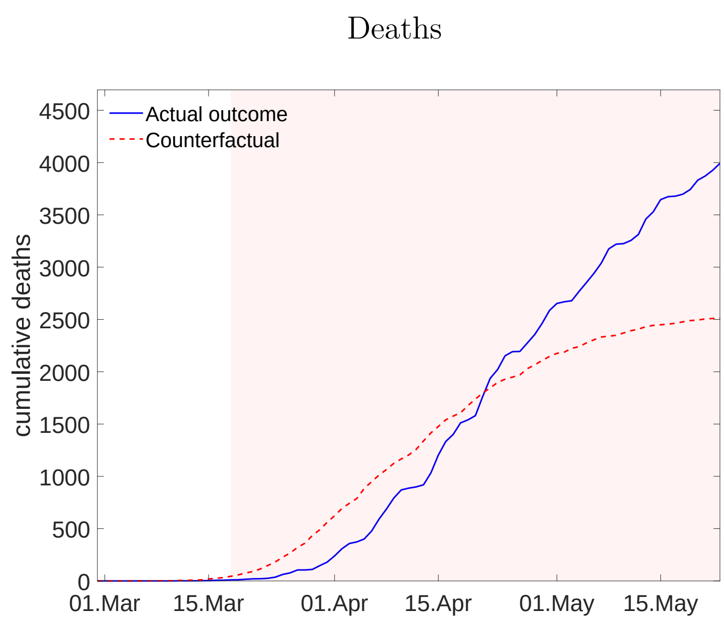

Figure A.4: Specification B: w/o Denmark: Infections and Deaths relative to Sweden. Shaded area indicates Lockdown period in counterfactual.

Specification C

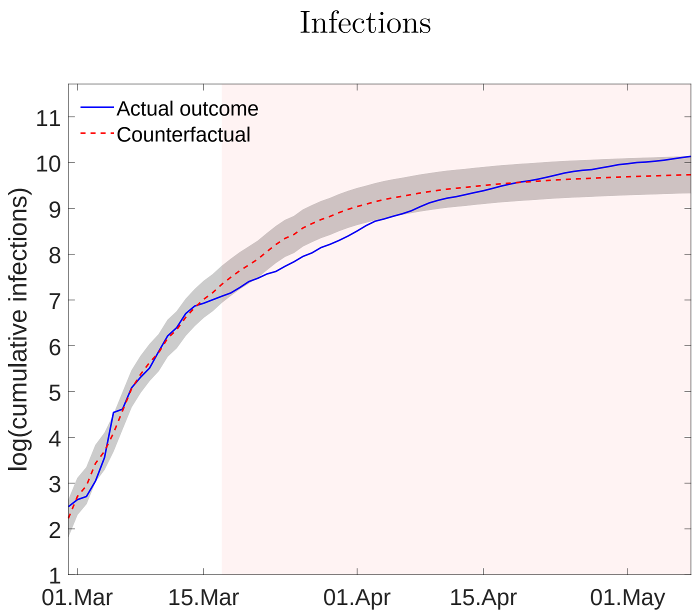

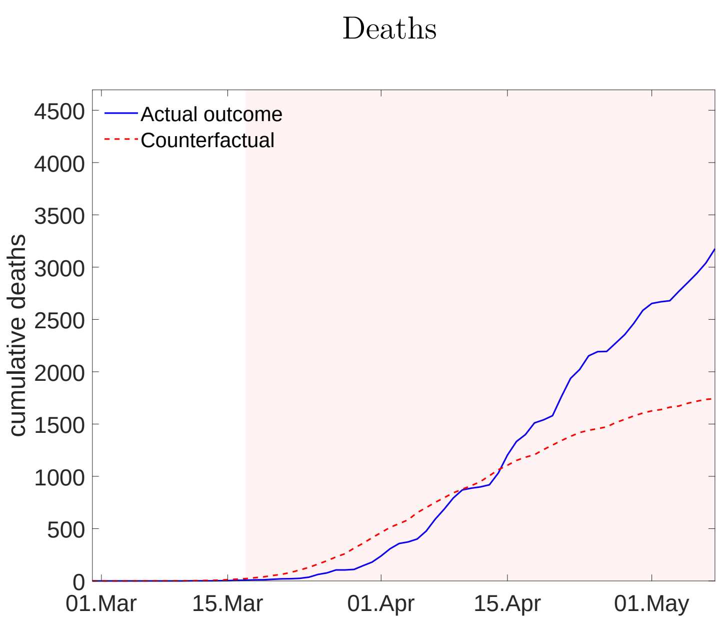

Figure A.5: Specification C: w/o Finland: Infections and Deaths relative to Sweden. Shaded area indicates Lockdown period in counterfactual.

Specification D

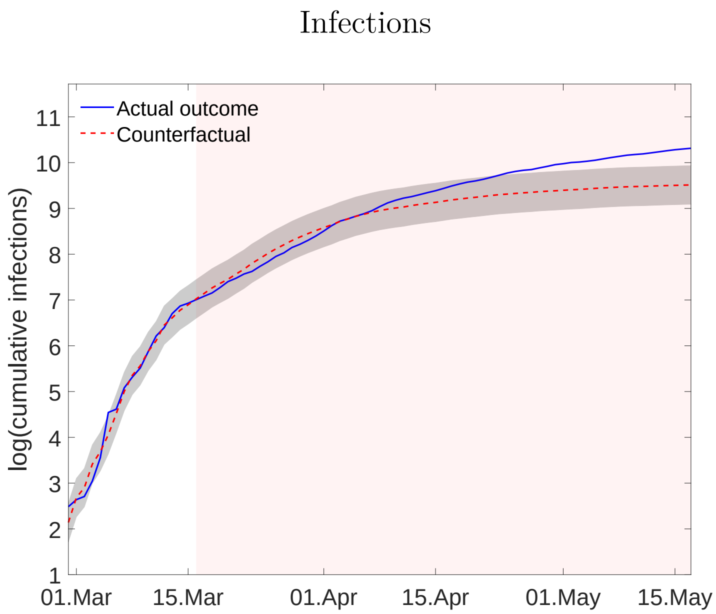

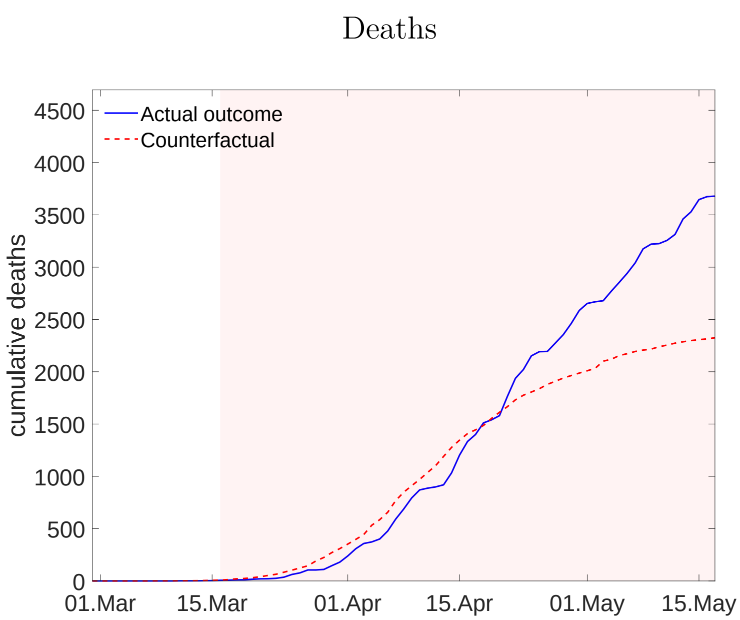

Figure A.6: Specification D: w/o Netherlands: Infections and Deaths relative to Sweden. Shaded area indicates Lockdown period in counterfactual.

Specification E

Infections

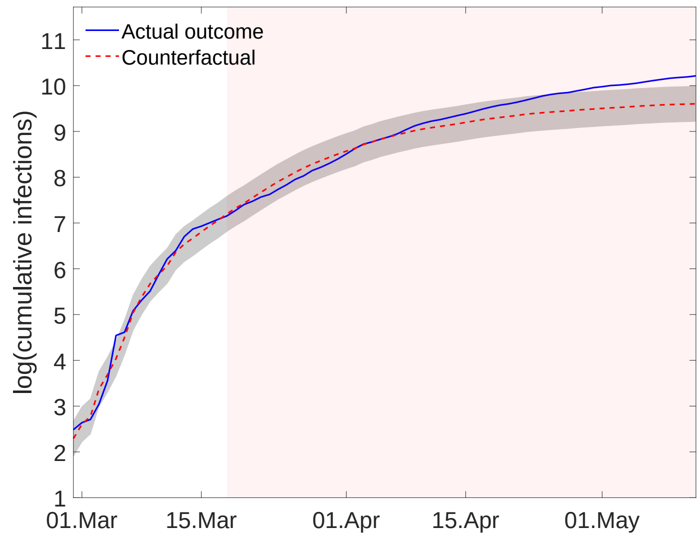

Deaths

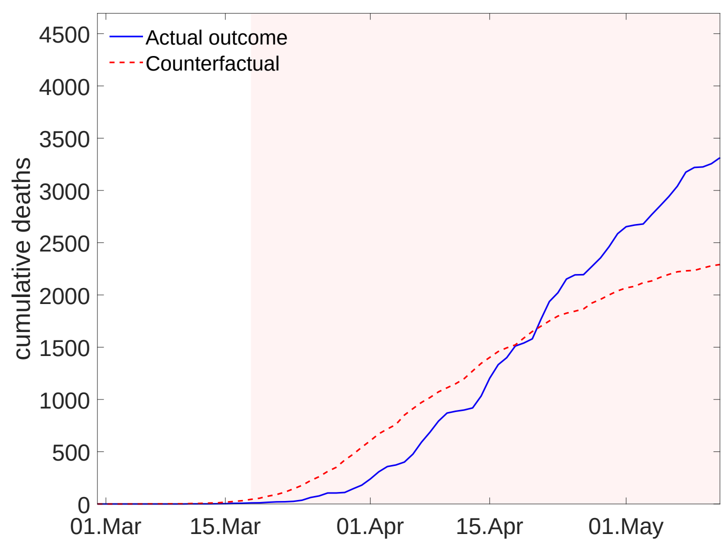

Figure A.7: Specification E: w/o Norway: Infections and Deaths relative to Sweden. Shaded area indicates Lockdown period in counterfactual.

Specification F

Infections

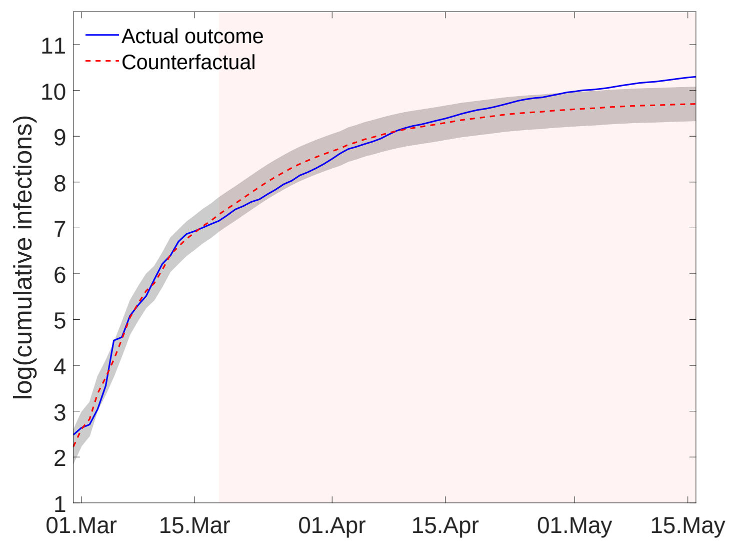

Deaths

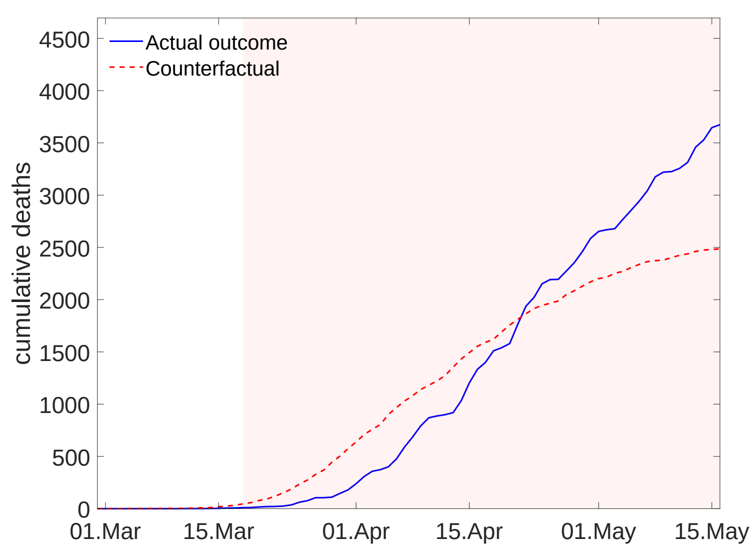

Figure A.8: Specification F: w/o Portugal: Infections and Deaths relative to Sweden. Shaded area indicates Lockdown period in counterfactual.

Specification G

Infections

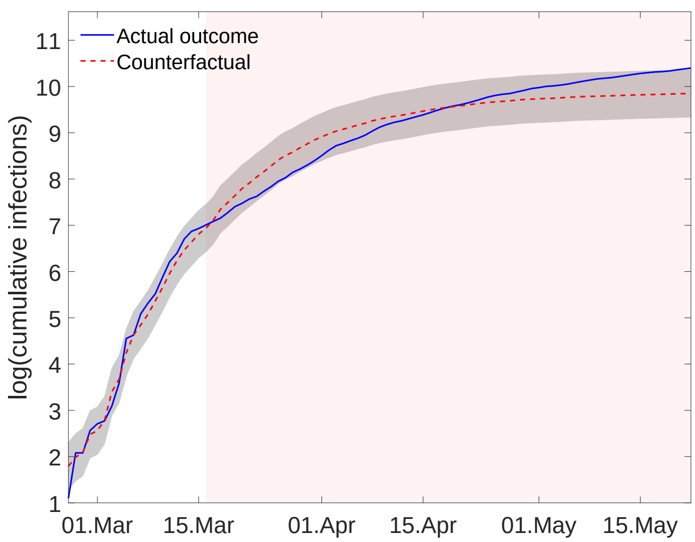

Deaths

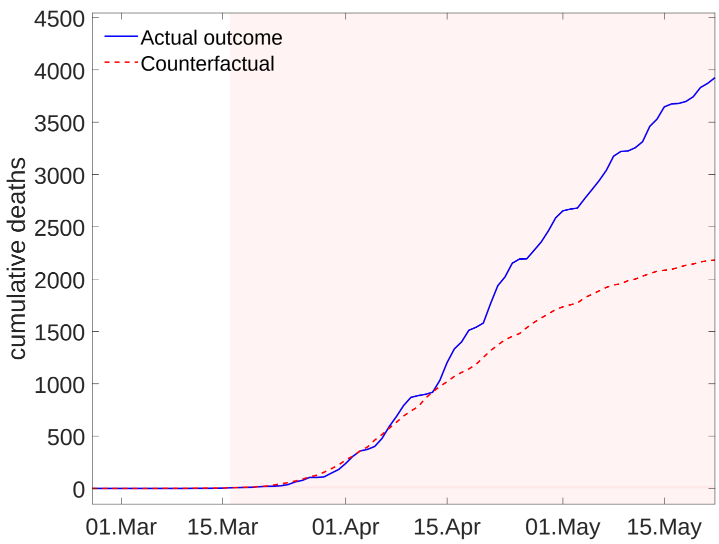

Figure A.9: Specification G. Shaded area indicates Lockdown period in counterfactual.

Specification H

Infections

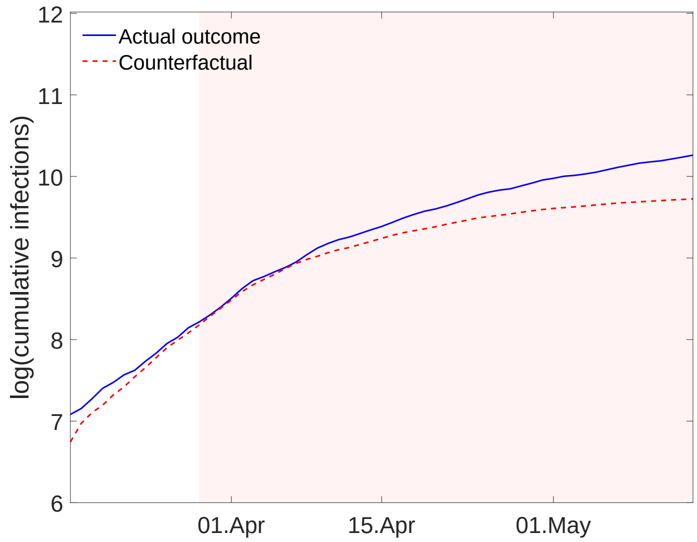

Deaths

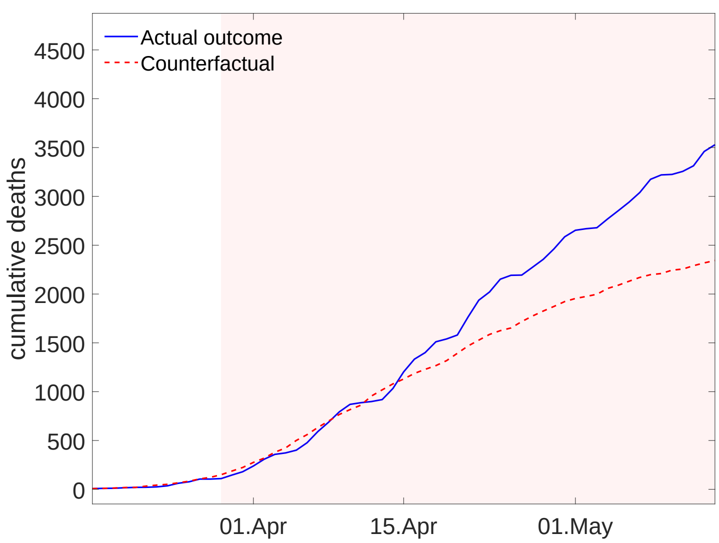

Figure A.10: Specification H. Shaded area indicates period from the end of matching to Lockdown end.

A.6 Data Sources for Lockdown Start and End Dates

Table A.2: Lockdown measures in Europe (start date): sources

Country	Source
Austria	*www.parlament.gv.at*
Belgium	*www.belgium.be*
Denmark	*politi.dk*
Finland	*valtioneuvosto.fi*
France	*www.diplomatie.gouv.fr*
Germany	*www.bundesregierung.de*
Greece	*gr.usembassy.gov/covid-19-information*
Ireland	*www.gov.ie*
Italy	*www.trovanorme.salute.gov.it*
Netherlands	*www.government.nl*
Norway	*www.helsedirektoratet.no*
Portugal	*www.acm.gov.pt*
Spain	*www.gov.uk/foreign-travel-advice/spain/coronavirus*

Table A.3: Lockdown measures in Europe (end date): sources

Country	Source
Austria	*www.wko.at*
Belgium	*www.belgium.be/en*
Denmark	*politi.dk*
Finland	*valtioneuvosto.fi*
France	*www.tagesschau.de*
Germany	*www.bundesregierung.de*
Greece	*www.visitgreece.gr*
Ireland	*www.gov.ie*
Italy	*www.salute.gov.it*
Netherlands	*www.government.nl*
Norway	*www.regjeringen.no*
Portugal	*www.acm.gov.pt*
Spain	*www.bbc.com*