Do lockdowns work?
A counterfactual for Sweden[*]

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

May 8, 2020

Abstract

Is a lockdown an effective means to limit the spread of the COVID-19 pandemic? We study the case of Sweden—one of the few countries without a lockdown—and use synthetic control techniques to develop a counterfactual lockdown scenario. First, we use a "donor pool" of European countries to construct a doppelganger that behaves just like Sweden in terms of infections *before the lockdown*. Second, we find that infection dynamics in the doppelganger *since the lockdown* do not systematically differ from the actual dynamics in Sweden. Third, we study Google mobility data and find that Swedes adjusted their activities in similar ways as in the doppelganger, although to a somewhat lesser extent.

Keywords:	COVID-19, lockdown, counterfactual, synthetic control, doppelganger, infections, social restraint, Google mobility reports
JEL-Code:	E60

"And when countries do all the wrong things and still end up seemingly not as battered by the virus as one would expect, go figure." (New York Times, May 3, 2020)

1 Introduction

As the COVID-19 pandemic spread around the globe in early 2020, many countries shut down their economies in order to limit social interactions. In China, the Hubei province was put under lockdown for more than two months, starting in late January 2020. In Europe, Italy imposed a lockdown on March 9, followed by the large majority of countries in Europe and elsewhere. A widely discussed exception is Sweden, which opted against a lockdown in its approach towards the COVID-19 pandemic. And while the Swedish authorities advised citizens to adjust their behavior in the face of the pandemic, one may ask if the spread of the pandemic would have been more limited, if the government had imposed a proper lockdown instead?[1]

From a theoretical point of view, the answer seems obvious because, as it constrains social interactions, a lockdown is bound to limit the transmission of the disease. It is important to recognize, however, that even in the absence of a lockdown, (economic) activities will slow down because people reduce social interactions voluntarily out of sheer self interest. Eichenbaum et al. (2020) put forward and formalize this notion. They also show that because people fail to internalize the costs they impose on others as they become infectious—in the presence of an infection externality—government restrictions on social interaction will generally be optimal (see also Alvarez et al., 2020; Farboodi et al., 2020). Still, Krueger et al. (2020) stress that actual economies offer a lot of possibilities to shift (economic) activities towards sectors where the risk of infection is small(er). As a result, voluntary adjustment is

---

[1]For instance, "everyone" was "advised to avoid unnecessary traveling and social events, to keep distance to others, and to stay at home" if they had any symptoms. In addition, as "age is a risk factor, those above 70" were "advised to avoid social contact and visits to retirement homes are banned". Quoted from website: https://www.krisinformation.se/en/hazards-and-risks/disasters-and-incidents/2020/official-information-on-the-new-coronavirus.

relatively easy and goes a long way towards containing the spread of the virus.

Against this background we take up the issue empirically and ask whether the spread of COVID-19 infections in Sweden would have been more limited, if a lockdown had been imposed. To answer this question, we develop a counterfactual scenario based on a variant of the synthetic control method (Abadie et al., 2010, 2015; Abadie and Gardeazabal, 2003). Ideally, in order to assess the effectiveness of a lockdown, we would like to compare the developments in a country that imposed a lockdown to a synthetic control group of countries where no lockdown is imposed. And indeed, Friedson et al. (2020) benchmark infection dynamics after California's Shelter-in-place order against those in other US states. However, in this case—like elsewhere—the comparison is complicated by the fact that the control unit is composed of states that also imposed a lockdown, if only somewhat later than California.

In our analysis of the Swedish case, we turn things upside down: we compare actual infection dynamics in the absence of a lockdown to a counterfactual outcome for which we assume that a lockdown is imposed. In order to develop this counterfactual, we also rely on the synthetic control approach and construct a "doppelganger" for Sweden (Born et al., 2019). We then compare the outcome in Sweden and its doppelganger while subjecting the doppelganger to a lockdown "treatment" and not Sweden. To the extent the doppelganger is behaving like Sweden prior to treatment, any difference in outcomes after the lockdown treatment may be attributed to it. Hence, this counterfactual experiment should inform us about what would have happened in Sweden had a lockdown been imposed.

We construct the doppelganger as a weighted average of the countries in a donor pool of European countries. Because infections do not start spreading simultaneously in all countries, we initialize observations for each country: day 1 is when the number of infections surpasses a threshold of one infection per one million inhabitants. After day 1 it took all countries in our donor pool at least 13 days to impose a lockdown. We select the weights for the doppelganger by matching infection dynamics within the 13 days window between day 1 and the first lockdown in our donor pool. At the same time, as we select the weights, we make

sure that the doppelganger equals Sweden in terms of population size and urbanization rate.

We find that the (almost) perfect doppelganger of Sweden is a weighted average of the Netherlands (weight of 0.39), Denmark (0.26), Finland (0.19), Norway (0.15), and Portugal (0.01). By construction the doppelganger behaves like Sweden during the 13 day window prior to the lockdown in terms of infection dynamics. Yet, and this is our main result, it also behaves not systematically different from Sweden after the lockdown is imposed in the doppelganger (which happens to be at day 18). This result suggests that the infection dynamics in Sweden would not have been different in case it had imposed a lockdown. Importantly, our analysis is not speaking to the question if the lockdown was effective in the countries that make up the doppelganger.

But we can verify that our result for Sweden does not depend on any individual country that enters the doppelganger. In a robustness exercise we exclude, in turn, each of the five countries from the donor pool and, as we construct new doppelgangers on the basis of the restricted donor pool, obtain very similar results as in the baseline, even though the composition of the doppelganger is necessarily somewhat different in each case. Moreover, we also compare the number of tests for COVID-19 and find that Sweden conducted fewer tests than the doppelganger. However, this is unlikely to be the (only) cause for our result, because we find that the number of COVID-19 deaths in Sweden has not been higher than in the doppelganger.

In sum, we find that a lockdown would not have helped much in Sweden. In order to account for this apparently puzzling result, we analyze Google mobility reports. We find that—even in the absence of a lockdown—Swedes adjusted their activities considerably, and in similar ways as in the doppelganger. Yet, by and large, the adjustment has been somewhat weaker. This suggests that voluntary social restraint goes some way in resolving the "lockdown puzzle".

One may object that this is merely about semantics. If all that matters is that people adjust their activities, why do we care if this is happening on a voluntary basis, because of a

government recommendation, or because of a full-fledged lockdown? The answer is twofold: First, from a theoretical point of view, it is interesting to learn whether people adjust their behavior voluntarily. Second, from a practical angle, it seems preferable if the necessary adjustment comes about voluntarily and does not need to be enforced. More importantly, still, it is conceivable that the benefits of social restraint (voluntary or due to a lockdown) are declining in the scope of the measures that are taken and, hence, a moderate restraint as observed in Sweden may already be sufficient to limit the spread of infections considerably.

The remainder of this note is organized as follows. Section 2 details our approach, notably the construction of the donor pool and the doppelganger. Section 3 presents results for the baseline. It also discusses our robustness exercises. A final section concludes. The appendix provides further details.

2   The approach

We develop a counterfactual scenario for Sweden in order to assess whether a lockdown would have helped to contain the spread of COVID-19 infections. For this purpose, we construct a "doppelganger" for Sweden based on synthetic control methods. Ideally, the doppelganger behaves just like Sweden before the lockdown such that any difference after the lockdown may be attributed to the "treatment" to which the doppelganger is subjected while Sweden is not.

2.1   The donor pool

The doppelganger is a weighted average of the countries in the "donor pool". To ensure a high degree of homogeneity between Sweden and its doppelganger, 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 when and what type of lockdown measures where imposed in the countries of the donor pool as well as in Sweden. The first lockdown was imposed in Italy on March 9, the last in the Netherlands on March 24. These lockdowns

Table 1: Donor pool vs. Sweden

Country	Lockdown Start	End	Note on containment measures	Day 1	Days to Lockdown
Austria	16.03.	-	non-essential shops closed, ban on gatherings >5	29.02.	16
Belgium	18.03.	-	non-essential shops closed, ban on gatherings >2	03.03.	15
Denmark	18.03.	-	non-essential shops closed	03.03.	15
Finnland	16.03.	-	governm. agencies closed, ban on gatherings >10	01.03.	15
France	17.03.	-	non-essential shops closed, ban on gatherings >2	29.02.	17
Germany	23.03.	-	non-essential shops closed, ban on gatherings >2	01.03.	22
Greece	23.03.	-	non-essential shops closed, stay-at-home-order	05.03.	18
Ireland	28.03.	-	non-essential shops closed, stay-at-home-order	04.03.	24
Italy	09.03.	-	non-essential shops closed, stay-at-home-order	22.02.	16
Netherlands	24.03.	-	non-essential shops closed, ban on gatherings	02.03.	22
Norway	13.03.	-	restaurants, bars closed, ban on gatherings >5 (24.03)	28.02.	14
Portugal	19.03.	-	no shops closed, governm. agencies closed, stay-at-home advice	06.03.	13
Spain	14.03.	-	non-essential shops closed, stay-at-home-order	01.03.	13
Sweden	-	-	No Lockdown imposed, ban on gatherings >50	29.02.	-

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 appendix.

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.

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

Hopkins University (Dong et al., 2020).[2] 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.

In the rightmost column, Table 1 also reports the number of days between that date and the day the lockdown was imposed.[3] Note that no country in the donor pool imposed a lockdown within the first 13 days after day 1. For the construction of the doppelganger, 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 doppelganger.

We also require the doppelganger 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 doppelganger: log infections at daily frequency within the 13 day window prior to the first lockdown, population size, and the urbanization rate.

2.2 Constructing the doppelganger

We construct the doppelganger by selecting weights on the countries in the donor pool for which we obtain the best match between the doppelganger and Sweden for the 15 target

---

[2]Observations are available at daily frequency. They are assembled using national (e.g. ministry of health, government) as well as international (e.g. WHO, European Centre for Disease Prevention and Control) official sources. For further Information see the website: `https://coronavirus.jhu.edu/us-map-faq`.

[3]Figure A.1 in the appendix displays the number of infections in each country per calendar date.

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, $j = 2, \ldots, 14$. Then, the doppelganger 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 $j = 2, \ldots, 14$ and $\sum_{j=2}^{14} w_j = 1$. In this expression, $\mathbf{V}$ is a (15×15) symmetric and positive semidefinite matrix.[4]

3   Results

We now turn to the results. First, we compare the infection dynamics in Sweden and the doppelganger. Afterwards, we suggest a tentative explanation of our findings based on data from Google's mobility reports.

3.1   Infection dynamics

Figure 1 shows the main result. The vertical axis measures cumulative infections in logs, the horizontal axis measures time in days since the number of infections exceeds one per one million inhabitants (day 1). The solid (blue) line represents data for Sweden. The dashed (red) line represents the doppelganger. It is a weighted average of the following countries: Netherlands (weight of 0.39), Denmark (0.26), Finland (0.19), Norway (0.15), and Portugal (0.01). By construction, that is by minimizing expression (1), the number of infections in the doppelganger tracks the dynamics in Sweden up to day 13 very closely. Still, the

---

[4]$\mathbf{V}$ 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 (see the discussion in Abadie et al. (2010), p. 496). Following Abadie and Gardeazabal (2003) and 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. Including the covariates in the optimization differs from Kaul et al. (2018) who have raised concerns about including all pre-intervention outcomes together with covariates when using the SCM.
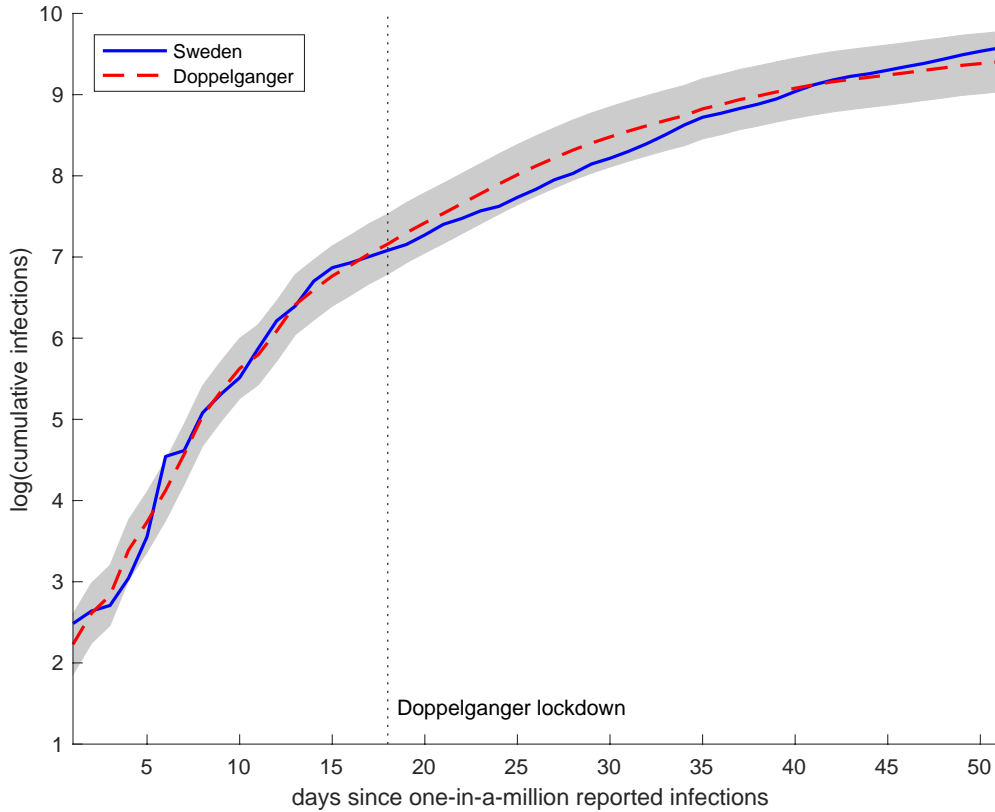
Figure 1: Cumulative infections in Sweden (blue solid line) and in its doppelganger (red dashed line) since day 1 (in logs). *Notes:* Doppelganger is weighted average of the Netherlands (0.39), Denmark (0.26), Finland (0.19), Norway (0.15), and Portugal (0.01). Targeted observations are log infections during the first 13 days as well as population size and urbanization rate. Shaded areas represent two standard deviations of the difference between infections in Sweden and its doppelganger during the first 13 days. Dotted line indicates the time of the lockdown in the doppelganger. Data source: Johns Hopkins University (Dong et al., 2020).

match is not perfect and the shaded area indicates two standard deviations of the difference between infection growth in Sweden and the doppelganger in the matching period. Recall that we also match the population size and the urbanization rate: it is 10.2 million and 0.87, respectively—both in Sweden and its doppelganger.[5]

Given that the doppelganger behaves like Sweden in terms of infection dynamics during the first 13 days, is of equal size and has the same urbanization rate, it provides a meaningful counterfactual for Sweden. The average number of days until a lockdown was imposed in the countries that make up the doppelganger is 18 days. It is indicated by the dotted

---

[5]The precise numbers for the urbanization rate are 0.87431 for Sweden and 0.87756 for the Doppelganger. For population they are 10.175 and 10.187 million, respectively.
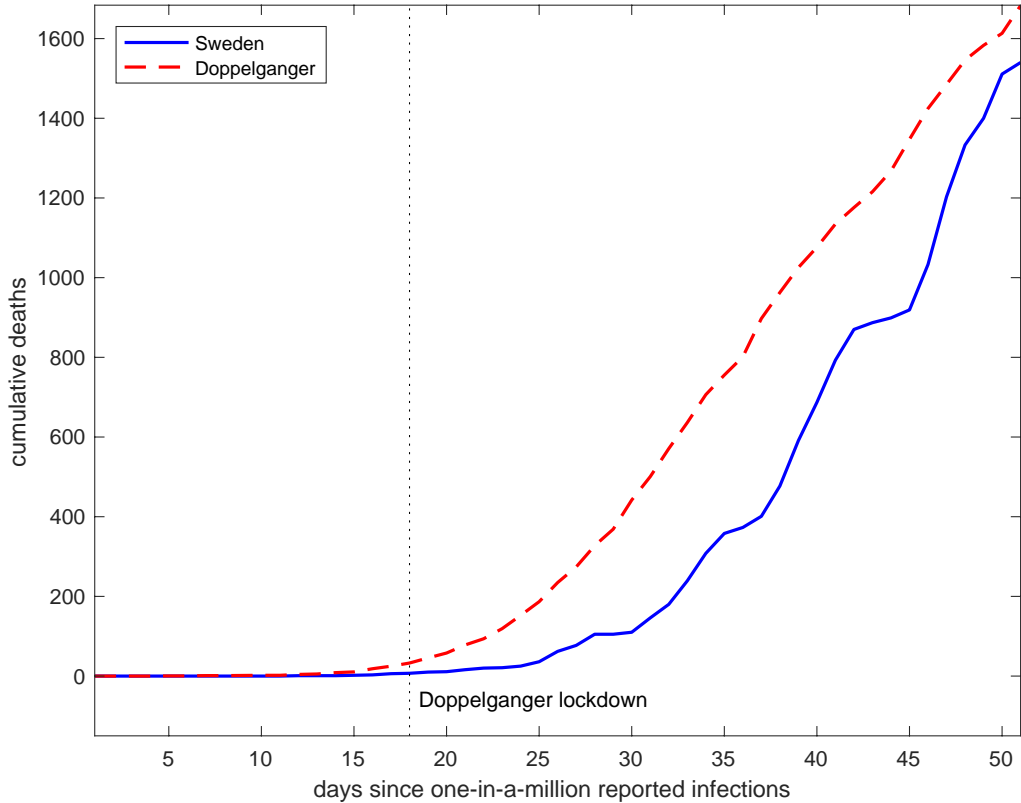
Figure 2: Cumulative COVID-19 deaths in Sweden (blue solid line) and in its doppelganger (red dashed line) since day 1. *Notes:* Based on the country weights used in Figure 1. See the notes to Figure 1 for additional details.

vertical line in Figure 1. We interpret the infection dynamics in the doppelganger afterwards as the outcome that would have been observed in Sweden, had a lockdown been imposed. Our identification assumption is that given pre-lockdown trends as well as other country characteristics, the probability of a lockdown was the same for Sweden and the doppelganger. The "treatment" was, in other words, random.

Under these assumptions, Figure 1 leaves little room for doubt: it shows that a lockdown would not have altered infection dynamics in Sweden in a meaningful way, at least not during the first four weeks after the lockdown. Given the available data at the time of writing we can not analyze the effect of the lockdown in the longer run. We expect this will be the subject of further analyses in the time to come.

However, for now we can make sure that our results are not driven by any individual country that contributes to the doppelganger. To do so, we construct doppelgangers again

9

on the basis of objective (1), but exclude from the donor pool one country at a time that has a non-zero weight in the baseline doppelganger. Table A.2 in the appendix reports the results. We obtain new weights in each instance, but the countries that enter the baseline are generally important (except, of course, if they are excluded from the pool). Figure A.4 shows the infection dynamics for each of the five alternative doppelgangers. They are close to the baseline doppelganger and, hence, close to the actual dynamics observed in Sweden.

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.3 in the appendix shows that there where fewer tests conducted in Sweden than in the doppelganger and increasingly so in the second half of the sample. Still the change is moderate and it is hard to say to what extent this matters for reported infections, notably since testing strategies differ across countries.

Against this background we also compare the number of COVID-19 deaths in Sweden and the doppelganger because the statistics are arguably less distorted in this case. Figure 2 displays the time series for the cumulative number of deaths. The figure is organized in the same way as Figure 1 above. The solid (blue) line represents data for Sweden. The dashed (red) line represents the outcome for the doppelganger. Note that we do not construct a new doppelganger but use exactly the same country weights as above. In the figure we display absolute numbers rather than logs, because initially there were no deaths reported, neither in Sweden nor in the doppelganger.

We observe that Sweden did not suffer more COVID-19 deaths than the doppelganger, even though the doppelganger and Sweden are of equal population size. Still, during the last two weeks in the period under consideration, we observe faster growth of COVID-19 deaths in Sweden than in the doppelganger. It remains to be seen whether Sweden will eventually suffer a larger number of deaths. Still, we note that a similar pattern emerges for the five alternative doppelgangers that we obtain for the restricted donor pool. Results are shown in Figure A.5 in the appendix.

3.2 Voluntary social restraint

On the basis of the available data, we find that a lockdown in Sweden would not have limited the number of infections or the number of COVID-19 deaths. Theory suggests that this may be the result of people maintaining a larger social distance even in the absence of a lockdown—there could be, in other words, voluntary social restraint. Krueger et al. (2020), in particular, show this in the context of a formal model and suggest that this may be the relevant case for Sweden.

In order to assess this hypothesis, we rely on the Google COVID-19 Community Mobility Reports (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.[6] Locations are classified according to six distinct categories: Grocery and pharmacy, Parks, Residential, Retail and recreation, Transit stations, and Workplaces. 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 3 displays mobility dynamics for each category, contrasting once more data for Sweden (blue solid line) and the doppelganger (red dashed line). As before, the horizonal axis measures time since day 1. The dotted line indicates the lockdown in the doppelganger. A number of findings stand out. First, we observe a pronounced decline for the locations retail and recreation, transit, and work. At the same time, there is an increase in domestic activities (residential) as well as for parks. Second, the adjustment starts to take place about 10 days before the lockdown and, importantly, both in the Sweden and the doppelganger. Last, we observe that while the adjustment of activities follows roughly the same pattern in Sweden and the doppelganger, it is somewhat more pronounced in the doppelganger. Hence, social restraint does seem to take place voluntarily but not to the same extent as in case of a lockdown. Still, since our counterfactual lockdown scenario for Sweden shows no improvement

---

[6]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".
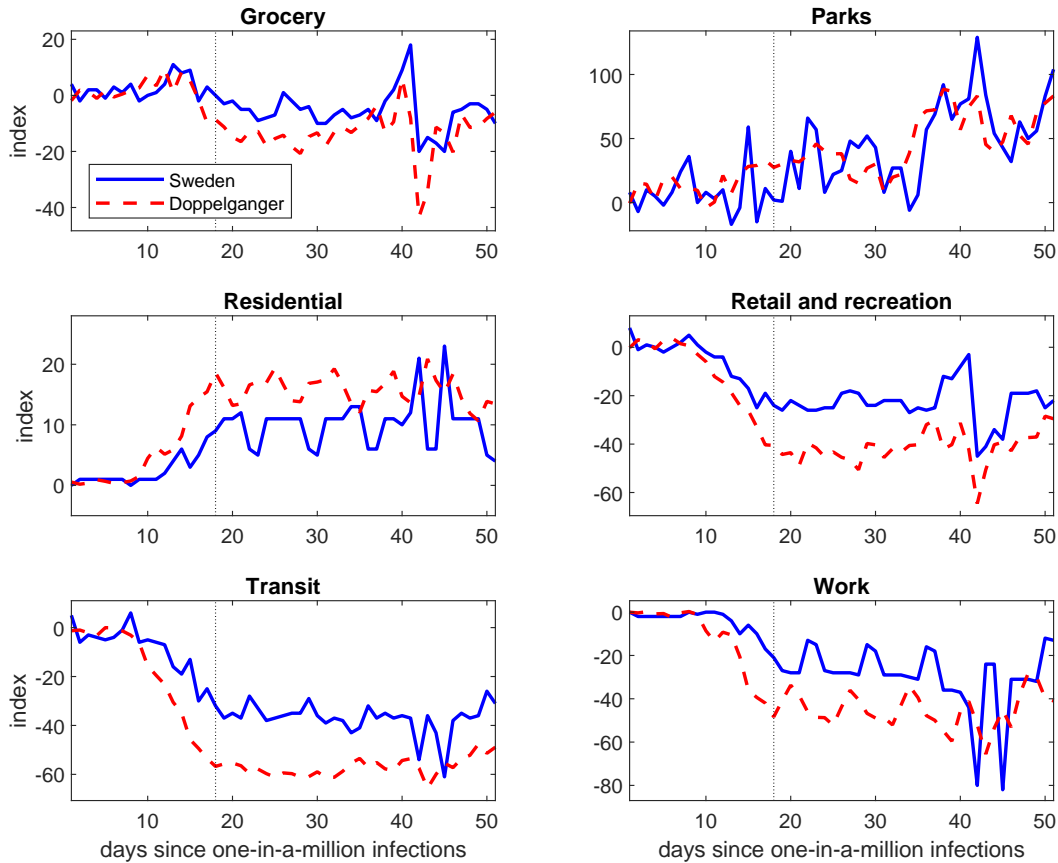
Figure 3: Mobility in Sweden (blue solid line) vs. doppelganger (red dashed line) since day 1. Dotted line: lockdown in doppelganger, see Figure 1 for details. Data source: Google (2020).

in terms of infection outcomes, it seems conceivable that the benefits of social restraint decline as its scope increases.

4   Conclusion

In this paper, we take up the question of whether lockdowns work, that is, whether they limit the spread of COVID-19. We develop a counterfactual for Sweden, one of the few countries where no lockdown was imposed. The counterfactual is based on data from other European countries, selected in such a way as to mimic key aspects of Sweden before any of those countries imposed a lockdown.

We find that a lockdown would not have helped in terms of limiting COVID-19 infections

or deaths in Sweden. As we analyze Google mobility reports, we find that Swedes have adjusted their activities in the absence of a lockdown in ways similar to what we observe for the counterfactual lockdown scenario. This suggests that voluntary social restraint plays an important role when it comes to accounting for our result. Against this background it is natural to ask whether voluntary social restraint implies significantly lower economic costs than a full-blown lockdown. We leave this question for future research.

Instead, in concluding, we stress some caveats. 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, our results are necessarily preliminary because it is based on data up to May 1, 2020 only. We cannot know whether our results change as more data comes in.

Last, 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—just because we find a lockdown would not have made a big difference for Sweden does not mean that it didn't work elsewhere. For it seems that for the COVID-19 pandemic to impact a country strongly, several conditions are required to be met at the same time (New York Times, 2020). In our view, the present study is but one of many that can advance our understanding of whether and how lockdowns work.

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.

——— (2015). "Comparative politics and the synthetic control method". *American Journal of Political Science 59* (2), 495–510.

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.

Alvarez, Fernando E, David Argente, and Francesco Lippi (2020). "A simple planning problem for COVID-19 lockdown". NBER Working Paper 26981.

Born, Benjamin, Gernot Müller, Moritz Schularick, and Petr Sedláček (2019). "The costs of economic nationalism: evidence from the Brexit experiment". *The Economic Journal 129* (623), 2722–2744.

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

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

Farboodi, Maryam, Gregor Jarosch, and Robert Shimer (2020). "Internal and external effects of social distancing in a pandemic". Working Paper 2020-47. Becker Friedman Institute.

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.

Google (2020). *Google COVID-19 Community Mobility Reports.* `https://www.google.com/covid19/mobility/`. Accessed: 01.05.2020.

Hasell, Joe et al. (2020). *To understand the global pandemic, we need global testing – the our world in data covid-19 testing dataset.* `OurWorldInData.org`.

Kaul, Ashok, Stefan Klößner, Gregor Pfeifer, and Manuel Schieler (2018). "Synthetic control methods: never use all pre-intervention outcomes together with covariates". Mimeo, University of Hohenheim.

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

New York Times (2020). *The covid-19 riddle: why does the virus wallop some places and spare others?* `https://www.nytimes.com/2020/05/03/world/asia/coronavirus-spread-where-why.html?referringSource=articleShare`.

A   Appendix

Table A.1: Lockdown measures in Europe: sources

Country	Source
Austria	*www.parlament.gv.at*
Belgium	*www.belgium.be*
Denmark	*politi.dk*
Finnland	*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.2: Doppelganger weights: restricted donor pool

	Baseline	I	II	III	IV	V
Austria	< 0.01	< 0.01	0.43	< 0.01	0.10	< 0.01
Belgium	< 0.01	< 0.01	< 0.01	0.19	< 0.01	< 0.01
Denmark	0.26	NA	0.24	0.31	0.33	0.26
Finland	0.19	0.13	NA	0.14	0.18	0.19
France	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
Germany	< 0.01	< 0.01	< 0.01	0.04	< 0.01	< 0.01
Greece	< 0.01	< 0.01	< 0.01	0.05	0.04	< 0.01
Ireland	< 0.01	0.01	0.01	< 0.01	< 0.01	< 0.01
Italy	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
Netherlands	0.39	0.41	0.27	NA	0.35	0.40
Norway	0.15	0.46	0.05	0.27	NA	0.15
Portugal	0.01	< 0.01	< 0.01	< 0.01	< 0.01	NA
Spain	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01

Notes: doppelganger weights in baseline and for restricted donor pool. In column I to V we exclude, in turn, one of the five countries that receive a positive weight in our doppelganger of the baseline specification.
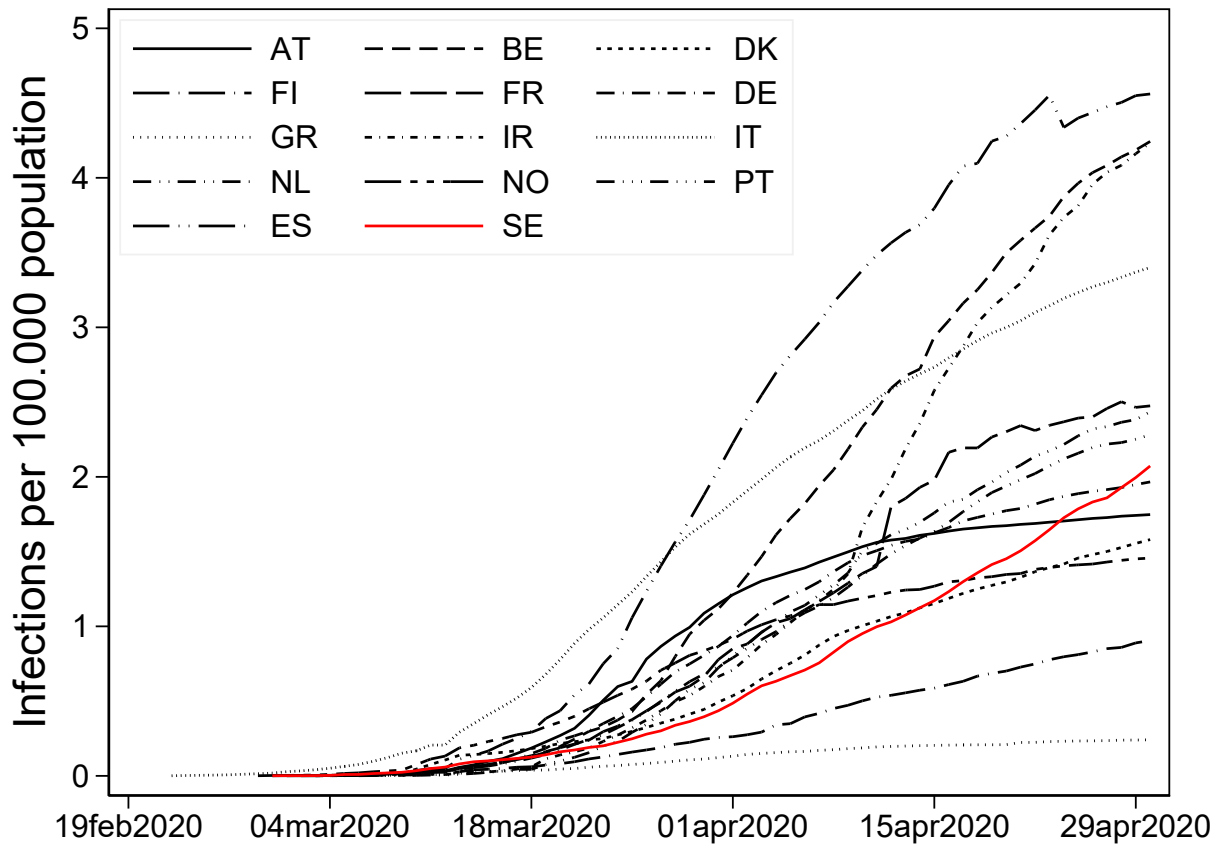
Figure A.1: Infections per 100.000 population for european countries vs. Sweden (red line). Countries grouped by region.

COVID-19 related Deaths per 100.000 by date for donor pool countries
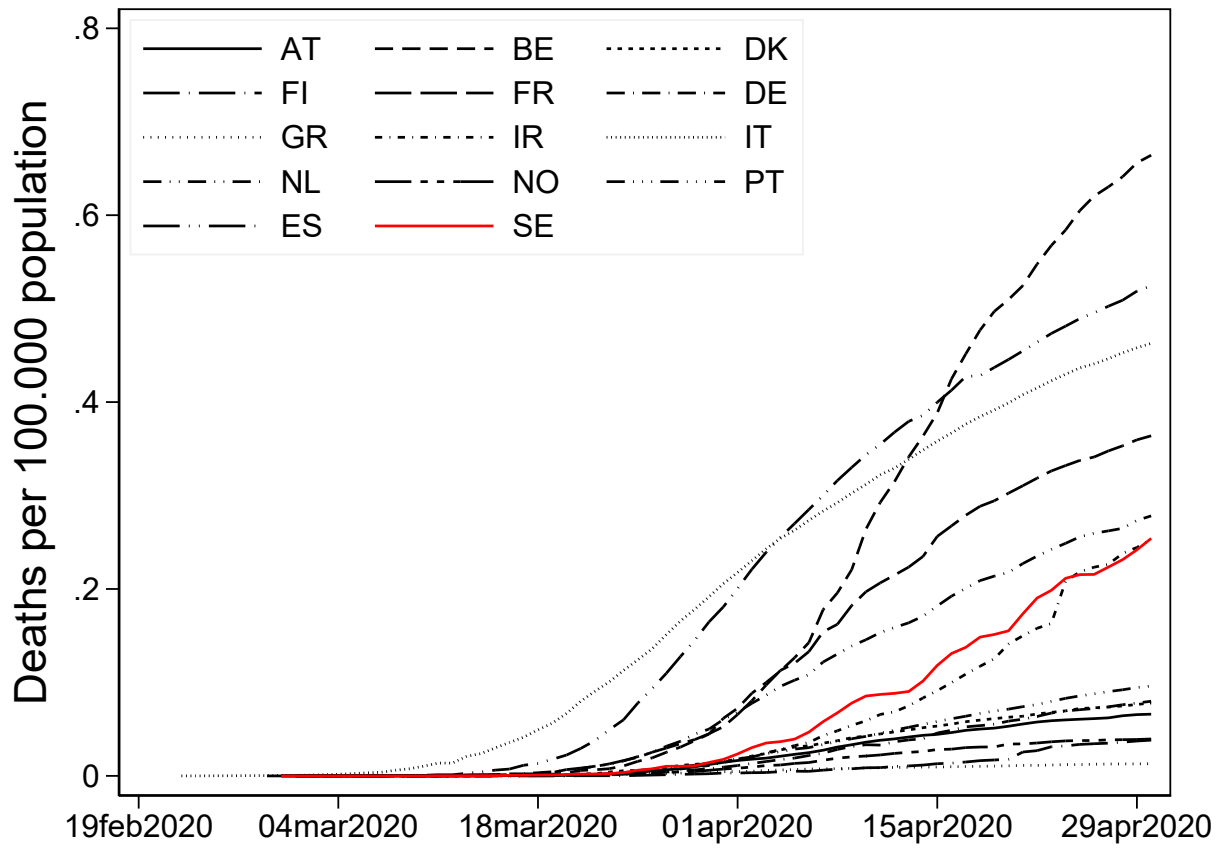
Figure A.2: Deaths per 100.000 population for european countries vs. Sweden (red line). Countries grouped by region.
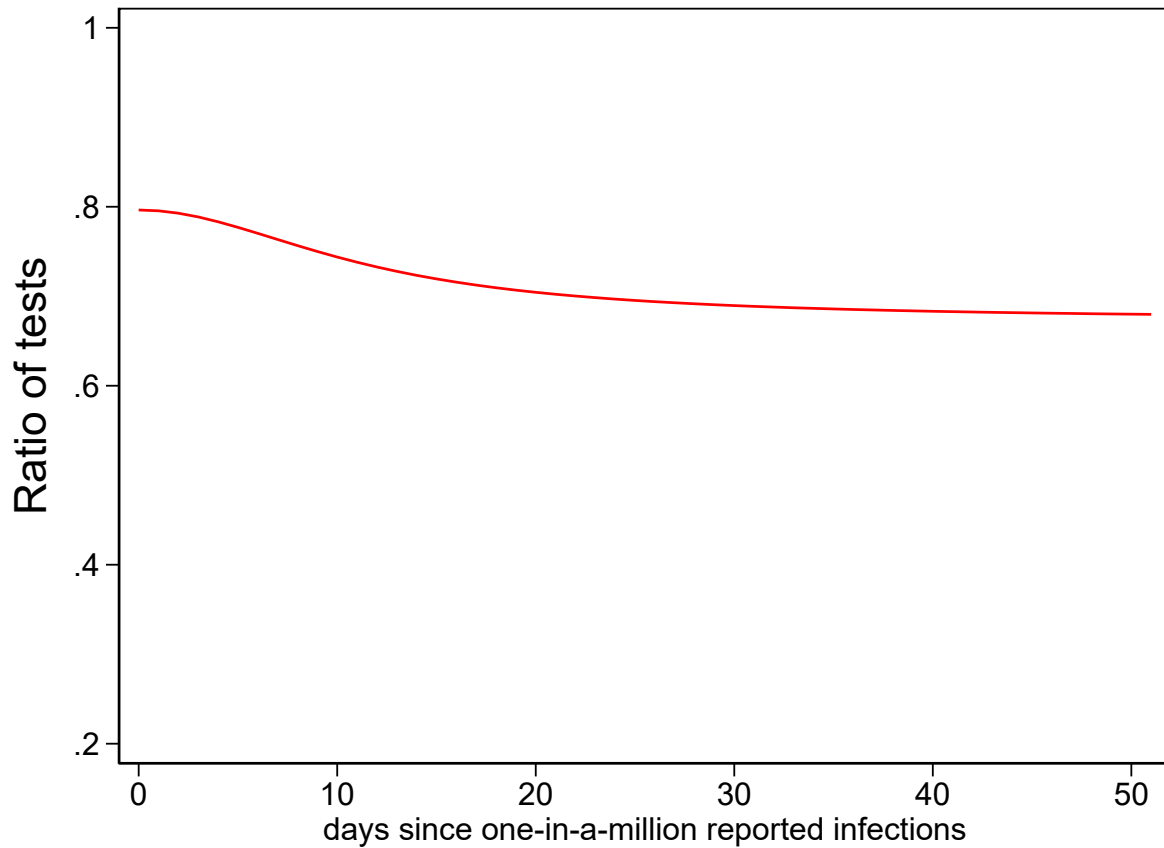
COVID-19 tests in Sweden relative to the Doppelganger

Figure A.3: Number of COVID-19 tests in Sweden relative to the Doppelganger. Source: Hasell et al. (2020). Observations on the number of tests are not reported at daily frequency both for Sweden and the countries that make up the doppelganger. We compute the ratio based on interpolated data after fitting a quadratic trend on the available observations.

COVID-19 Infections in Sweden vs. Doppelganger: Robustness
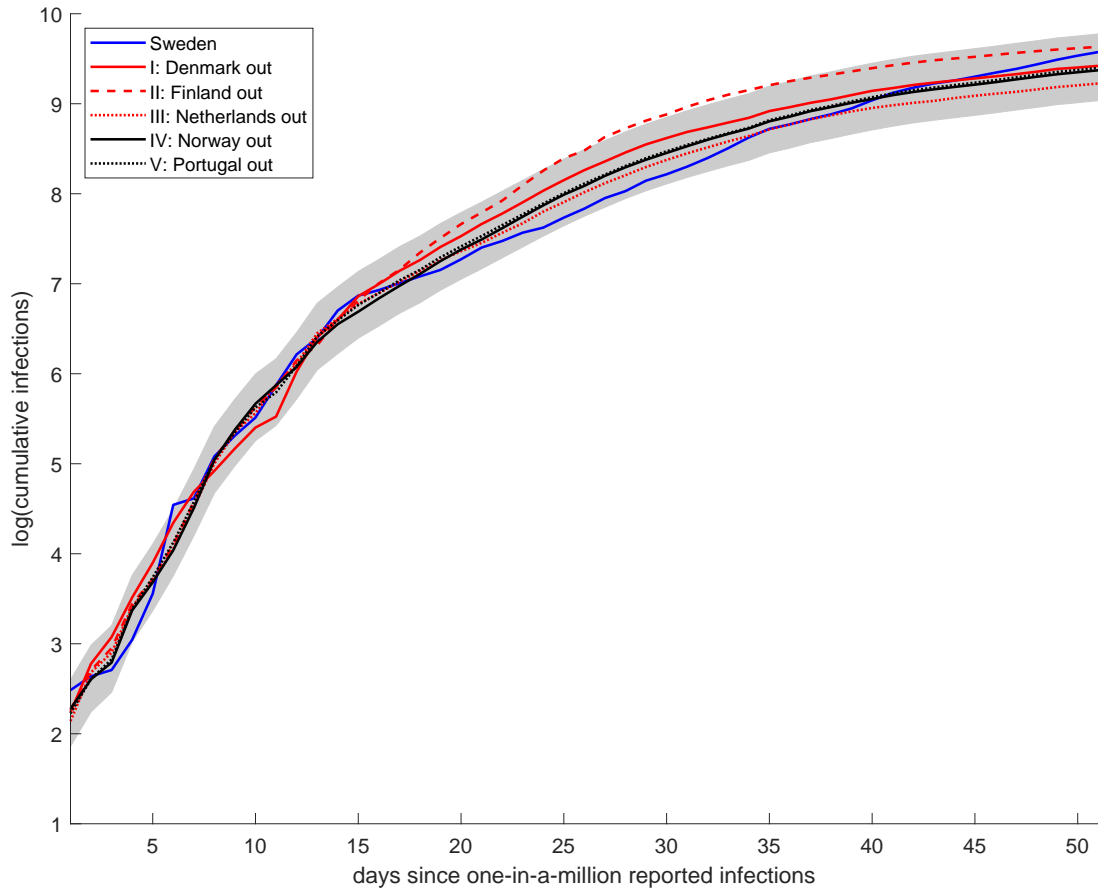
Figure A.4: Cumulative infections in Sweden (blue solid line) and in five doppelgangers since day 1 (in logs). *Notes:* Doppelganger are constructed as each country that contributes to the baseline doppelganger is, in turn, excluded from the donor pool. Table A.2 reports country weights. Shaded areas represent two standard deviations of the difference between infections in Sweden and the baseline doppelganger during the first 13 days. Data source: Johns Hopkins University (Dong et al., 2020).

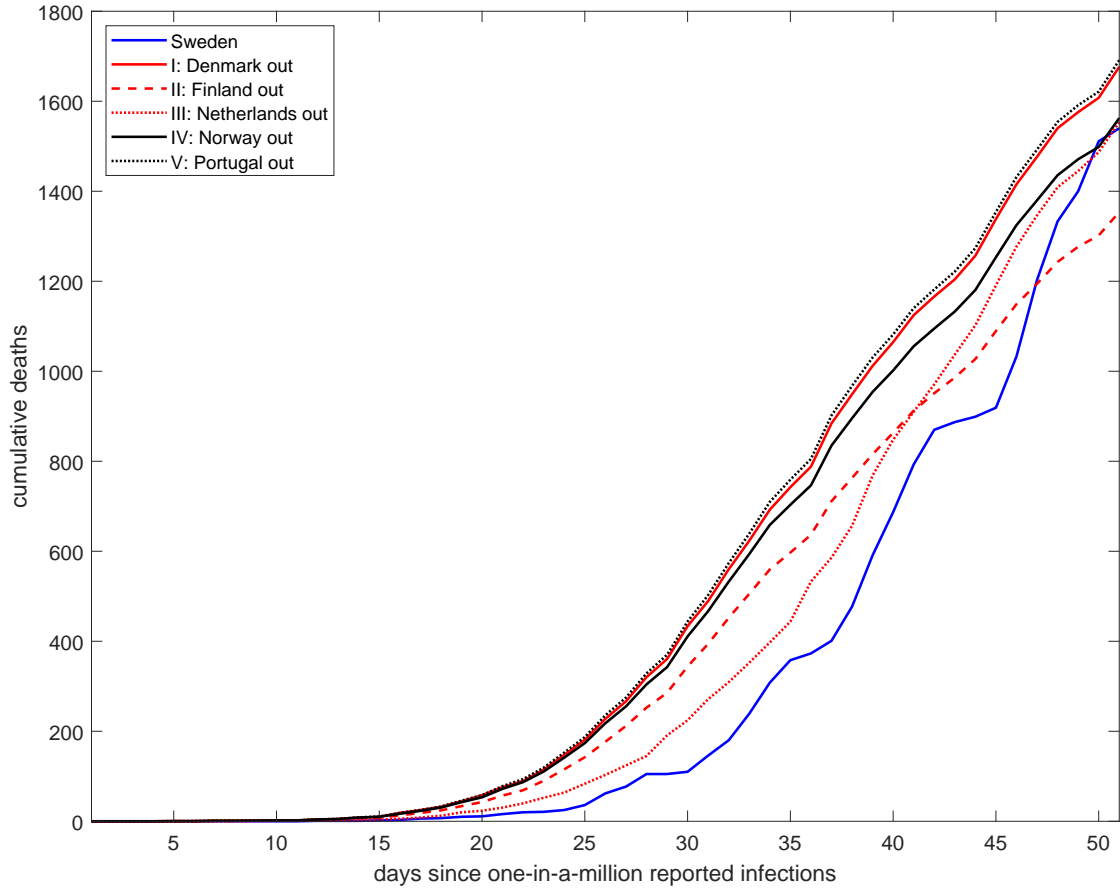COVID-19 Deaths in Sweden vs. Doppelganger: Robustness

Figure A.5: Cumulative Deaths in Sweden (blue solid line) and in five doppelgangers since day 1 (in logs). *Notes:* Doppelganger are constructed as each country that contributes to the baseline doppelganger is, in turn, excluded from the donor pool. Table A.2 reports country weights. Data source: Johns Hopkins University (Dong et al., 2020).