

WZB

Berlin Social Science Center



Ruud Koopmans

A Virus That Knows No Borders? Exposure to and Restrictions of International Travel and the Global Diffusion of COVID-19

Discussion Paper

SP VI 2020-103

October 2020

WZB Berlin Social Science Center

Research Area

Migration and Diversity

Research Unit

Migration, Integration, Transnationalization

WZB Berlin Social Science Center
Reichpietschufer 50
10785 Berlin
Germany
www.wzb.eu

Copyright remains with the author.

Discussion papers of the WZB serve to disseminate the research results of work in progress prior to publication to encourage the exchange of ideas and academic debate. Inclusion of a paper in the discussion paper series does not constitute publication and should not limit publication in any other venue. The discussion papers published by the WZB represent the views of the respective author(s) and not of the institute as a whole.

Ruud Koopmans

A Virus That Knows No Borders? Exposure to and Restrictions of International Travel and the Global Diffusion of COVID-19

Discussion Paper SP VI 2020-103
Wissenschaftszentrum Berlin für Sozialforschung (2020)

Affiliation of the author

Ruud Koopmans

Director, Migration and Diversity Research Area
WZB Berlin Social Science Center
Email: ruud.koopmans@wzb.eu

Abstract

A Virus That Knows No Borders? Exposure to and Restrictions of International Travel and the Global Diffusion of COVID-19

by Ruud Koopmans

„Closing borders is naive, the virus will come regardless" – this was the policy assumption that was repeatedly stated until mid-March by the WHO, the EU, as well as responsible authorities in Germany and other countries. Meanwhile, other states had started closing their borders to travellers from high-risk countries or to introduce mandatory quarantines. On 17 March, the EU did what it had previously argued against, and closed its borders to travellers from outside the EU and the Schengen Area. Germany, too, changed its line, and closed its borders to France, Switzerland, and Austria and on 18 March also to travellers from Italy. Who was right? Those who initially rejected travel restrictions as useless or those countries that decided to introduce them early on? Results from a global analysis of travel restrictions and cross-national differences in mortality rates as a result of the COVID-19 pandemic suggest that the belief that the spread of the virus could not be significantly slowed down by entry restrictions was fatally mistaken. The paper also shows that exposure of a country to international travel, as indicated by centrality in air travel networks and tourist numbers is strongly associated with higher COVID-19 mortality rates. By contrast, island states, which have lower exposure to international travel because of their lack of land borders, have much lower mortality. The results are robust across a wide variety of model specifications and controls, including domestic COVID-19 containment measures. The findings have important policy implications and suggest that in containing upcoming waves of the COVID-19 pandemic as well as similar pandemics in the future, the risks of exposure to international travel and the advantages of early travel restrictions should be given much greater weight. Among various types of travel restrictions, the findings suggest prioritizing targeted restrictions over global ones, and mandatory quarantines for travellers over entry bans.

Keywords: Covid-19 pandemic, diffusion, social networks, international travel, World Health Organization (WHO)

JEL classification: I18, D85, L93

Zusammenfassung

Ein Virus, das keine Grenzen kennt? Ausmaß und Einschränkungen des internationalen Reiseverkehrs und die globale Verbreitung von COVID-19

von Ruud Koopmans

"Grenzen zu schließen ist naiv, das Virus wird trotzdem kommen" – das war die politische Annahme, die bis Mitte März wiederholt von der WHO, der EU sowie den zuständigen Behörden in Deutschland und anderen Ländern vertreten wurde. In der Zwischenzeit hatten andere Staaten begonnen, ihre Grenzen für Reisende aus Hochrisikoländern zu schließen oder obligatorische Quarantänen einzuführen. Am 17. März tat die EU das, wogegen sie zuvor argumentiert hatte, und schloss ihre Grenzen für Reisende von außerhalb der EU und des Schengen-Raums. Auch Deutschland änderte seine Linie und schloss seine Grenzen zu Frankreich, der Schweiz und Österreich und am 18. März auch für Reisende aus Italien. Wer hatte Recht? Diejenigen, die die Reisebeschränkungen zunächst als nutzlos ablehnten, oder die Länder, die sich frühzeitig zu ihrer Einführung entschlossen hatten? Die Ergebnisse einer globalen Analyse von Reisebeschränkungen und der internationalen Unterschiede in der Sterblichkeitsrate infolge der COVID-19-Pandemie legen den Schluss nahe, dass die Annahme, die Ausbreitung des Virus könne durch Einreisebeschränkungen nicht wesentlich verlangsamt werden, ein fataler Irrtum war. Das Papier zeigt auch, dass die Exposition eines Landes gegenüber dem internationalen Reiseverkehr, wie sie sich aus der Zentralität in Flugverkehrsnetzwerken und den Touristenzahlen ergibt, stark mit höheren COVID-19-Mortalitätsraten verbunden ist. Im Gegensatz dazu haben Inselstaaten, die aufgrund fehlender Landgrenzen dem internationalen Reiseverkehr weniger ausgesetzt sind, eine viel geringere Sterblichkeit. Die Ergebnisse sind für eine Vielzahl von Modellspezifikationen und Kontrollen, einschließlich inländischer COVID-19-Eindämmungsmaßnahmen, robust. Die Ergebnisse haben wichtige politische Implikationen und deuten darauf hin, dass bei der Eindämmung bevorstehender Wellen der COVID-19-Pandemie sowie ähnlicher Pandemien in der Zukunft den Risiken der Exposition gegenüber dem internationalen Reiseverkehr und den Vorteilen frühzeitiger Reisebeschränkungen viel größeres Gewicht beigemessen werden sollte. Unter den verschiedenen Arten von Reisebeschränkungen deuten die Ergebnisse darauf hin, dass gezielte Beschränkungen für Hochrisikoländer über globale Einschränkungen zu bevorzugen sind und dass obligatorische Quarantänen für Einreisende effizienter sind als Einreisebeschränkungen.

Schlüsselwörter: Covid-19-Pandemie, Diffusion, soziale Netzwerke, internationale Reisen, Weltgesundheitsorganisation (WHO)

JEL-Klassifikation: I18, D85, L93

Executive summary

1. *This study investigates the roles of exposure to, and restrictions of international travel in explaining differences in COVID-19 mortality across the global population of independent states. It does so using publicly available data on COVID-19 deaths and exposure to international travel, as well as newly gathered data on the exact timing of the introduction of six types of travel restrictions: entry bans and mandatory quarantines, respectively targeting China, Italy, or all foreign countries.*
2. *Higher exposure of countries to international travel (as measured by the air travel exposure index AEF or the yearly number of tourist arrivals) is strongly and consistently associated with higher COVID-19 death tolls. Island states, by contrast, have much lower mortality rates.*
3. *Early introduction of policies that restrict international travel (entry bans and quarantines) strongly and consistently reduces the COVID-19 death toll. “Early” here means both early in absolute time, and relative to the local timing of the pandemic. Travel restrictions were especially powerful when countries introduced them before the local pandemic had passed a certain threshold, which seems to lie around the time of the 10th domestic death.*
4. *Among different types of travel restriction policies, mandatory quarantines were more effective than entry bans. The reason likely is that entry bans in most cases contain exceptions for returning citizens and permanent residents and therefore exclude an important part of traveller inflow. Quarantines, by contrast, usually apply to all incoming travellers, regardless of nationality or residence status.*
5. *Targeted travel restrictions (here measured through entry bans and quarantines for travellers from China and Italy) turn out to be more efficient than general restrictions that target all foreign countries. While general restrictions are effective to the extent that they encompass restrictions on high-risk countries, they have no measurable added value beyond what targeted travel restrictions can achieve.*
6. *The results for travel restrictions hold across a wide range of model specifications and robustness checks (including additional controls for domestic containment policies such as school closures and bans on public gatherings), as well as in a quasi-experimental design that compares treatment and control groups that differ only regarding the timing of travel restrictions.*
7. *The effect sizes of travel restrictions are substantial. Comparing in the quasi-experimental design the early adopter group of countries that were among the one third of the sample that introduced travel restrictions the earliest, to the latecomer group consisting of the one third that introduced travel restrictions the latest (or not at all), we find that early adopters have an estimated 62 percent lower COVID-19 mortality. Regression results indicate a 0.8 percent reduction in cumulative mortality per day that travel restrictions were introduced earlier.*

8. *Beyond the effects of exposure to international travel and of travel restrictions, the study provides some evidence that more affluent countries and democracies have higher death tolls. The reason probably is not that these countries have higher actual numbers of deaths, but that they report more deaths because of their more developed health systems and greater willingness to admit the true extent of the pandemic. These effects are however relatively small compared to the effects of travel exposure and restrictions.*

9. *The study has important policy implications. Countries that are highly exposed to international travel because of their centrality in airline networks and high tourist flows should be aware that they run a much-increased risk of early and multiple seeding from pandemic source regions. Contrary to the common wisdom during the early phases of the COVID-19 pandemic and the recommendations of the World Health Organization at the time, restrictions of international travel are an efficient means of pandemic containment, especially if they are implemented when domestic case and fatality numbers are still low.*

Acknowledgements

Earlier versions of this study were presented in the colloquium “Coronasozio­logie” organized by Jutta Allmendinger and Armin Nassehi; the series WZB Talks organized by Macartan Humphreys and Steffen Huck; and twice in the colloquium of the Research Area Migration & Diversity at the WZB Berlin Social Science Center. I have profited enormously from comments and suggestions received during and in the aftermath of these presentations. In particular, I would like to thank Daniel Auer, Macartan Humphreys, Ines Michalowski, Daniel Meierrieks, Max Schaub and Jonas Wiedner for their generous help and suggestions. I also thank Bao-An Nguyen, Aaron Lauterbach, Gizem Ünsal and Jegor Wolowikow, and especially Jasper Jansen for their meticulous assistance in gathering the data and preparing the online appendix. Elisabeth von Bressendorf and Melinda Biolchini brought the text into the format for the WZB paper series. All remaining errors are of course mine. While all care has been taken to collect as accurate as possible information on the introduction dates of various types of travel restrictions, I would be grateful for any notification of errors in the data, should there be any. I am also grateful for any other comments and suggestions.

Contents

1. Introduction: “Viruses know no borders”	1
1.1 The role of the World Health Organization	2
1.2. A short history of travel restrictions during the COVID-19 pandemic	5
2. Theoretical framework.....	10
2.1. The strength of cutting weak ties: what sociology teaches us about diffusion in social networks	10
2.2. International seeds and domestic growth.....	17
3. Data and variables	29
3.1. Country sample.....	29
3.2. Dependent variables.....	30
3.3. Exposure to international travel.....	34
3.4. Policies restricting international travel	36
3.5. Control variables.....	41
3.6. Modelling strategy	43
4. Main results	45
4.1. Exposure to international travel.....	45
4.2. International travel restrictions	48
5. Robustness checks	58
5.1. Reverse causality?	58
5.2 Considering other COVID-19 containment policies	61
5.3. Different observation points of cumulative mortality.....	64
5.4. Further robustness checks.....	65
6. A quasi-experimental approach	66
7. What about China and Italy?	69
8. Conclusions	72
9. References	76
10. Appendix: Additional robustness checks	79

1. Introduction: “Viruses know no borders”

At a federal press conference on the corona pandemic on 11 March, 2020, German Chancellor Angela Merkel stated: “Our opinion in Germany is that border closures are not an adequate answer to the challenges” („*Wir sind in Deutschland der Meinung, dass Grenzschießungen keine adäquate Antwort auf die Herausforderungen sind*“). She was seconded at the event by Lothar Wieler, President of the Robert-Koch-Institute, Germany’s leading epidemiological institute, who put it succinctly and without a shimmer of doubt: “Isolation is naïve. The virus will come regardless” („*Abriegelung ist naiv. Das Virus wird trotzdem kommen*“).

The assumption that entry restrictions are an ineffective means to stop or significantly slow down the spread of the virus was widespread at the time. At the EU Health Council of 13 February, Health Commissioner Stella Kyriakides began her press statement with an emphatic “viruses know no borders,”¹ and she kept repeating that phrase on numerous later occasions, calling instead for “cooperation” and “working together” as the only way to effectively combat the virus.² Germany’s Health Minister Jens Spahn also frequently repeated the phrase³ and categorically rejected border closures, for instance on 6 March on the occasion of an extraordinary meeting of the 27 EU health ministers in Brussels: “I continue to find any measure that restricts travel across borders inappropriate, considering what we know today about the virus” („*Ich fände jede Maßnahme, die zur Einschränkung des Reiseverkehrs über die Grenze führt, angesichts dessen, was wir über das Virus Stand heute wissen, weiterhin nicht für [sic] angemessen*“).⁴ This, the health minister emphasized, applied not just to travel within the EU,

¹ <https://www.youtube.com/watch?v=QDOH65qFeEc>

² As late as 19 May, in a common statement with the EU’s Foreign Policy and Security High Representative Josep Borell, Kyriakides reiterated that “the virus knows no borders, and neither should our response;” see https://eeas.europa.eu/headquarters/headquarters-homepage/79610/world-health-organisation-joint-statement-high-representativevice-president-josep-borrell-and_en

³ E.g. on 12 February in the Bundestag: <https://www.cducsu.de/themen/familie-frauen-arbeit-gesundheit-und-soziales/jens-spahn-ein-virus-macht-den-grenzen-nicht-halt>

⁴ <https://www.deutsche-apotheker-zeitung.de/news/artikel/2020/03/06/suedtirol-jetzt-risikogebiet-spahn-gegen-grenzschiessungen>

but also to entry restrictions for travellers from third countries. “The virus is there. It is in Europe” („Das Virus ist da. Es ist in Europa“), he stated, echoing the idea that once cases of the virus have been introduced, the dynamic of the epidemic is dominated by domestic spread and new contagions from abroad carry little weight. Accordingly, the minister was laconic about the need for mandatory quarantine policies, emphasizing strictly voluntary restraint instead: “if a student returns from ski holidays in Northern Italy, he should on his own accord say to himself, I should perhaps better not go to university for 14 days” („Wenn ein Student zum Skifahren in Norditalien gewesen ist und jetzt zurückkehrt, sollte der von sich aus sagen, ich gehe jetzt mal nicht in die Universität für 14 Tage“⁵). This was 6 March: by then, 4,600 infections had officially been registered in Italy and 197 people had died there.⁶

1.1 The role of the World Health Organization

When one traces the origins of the conviction that travel restrictions are powerless against the spread of the virus, one actor in particular looms large: the World Health Organization (WHO). Many experts and political actors referred to the WHO’s recommendations against travel restrictions. The WHO’s sceptical attitude towards travel restrictions is to some extent built into its remit, which is based on the International Health Regulations (updated in 2005), which state as its main aim “to prevent, protect against, control and provide a public health response to the international spread of disease in ways that are commensurate with and restricted to public health risks, and which avoid unnecessary interference with international traffic

⁵ <https://www.dw.com/de/eu-rechnet-mit-rapide-steigenden-corona-zahlen/a-52666602>; see also Kee et al. 2020: 1595.

⁶ Given that 197 had died and the death rate of the COVID-19 virus is according to what we know not higher than one percent, the true number of infections in Italy at that point in time must have been much higher than the registered 4,600. Assuming a cases to deaths ration of 100: 1, 197 deaths would imply about 20,000 cases. Moreover, given that the time between infection and death averages about two weeks this estimate of 20,000 does not refer to the situation on 6 March but to that two weeks earlier. With an initial doubling time of five days, that original number would moreover in the meantime have increased four to eightfold. One can thus infer that there may have been more than 100,000 infections in Italy already by 6 March.

and trade.”⁷ Indeed, the WHO had been emphatically advising against travel restrictions from the start of the COVID-19 pandemic, emphasizing both their presumed ineffectiveness and their economic costs. On 10 January, the WHO issued its first “advise for international travel and trade” relating to the new coronavirus. It began by stating that “From the currently available information, preliminary investigation suggests that there is no significant human-to-human transmission, and no infections among health care workers have occurred.” This “available information” clearly did not include the warning that Dr. Li Wenliang had posted on social media end of December urging his colleagues to protect themselves against a new infectious virus that was spreading in the city of Wuhan. Instead of taking Dr. Wenliang’s warning seriously and choosing to err on the side of caution, the WHO followed the official line of the Chinese authorities. The latter had accused Dr. Wenliang of “spreading rumours,” had forced him to withdraw his statement, and threatened him with further sanctions. Regarding travel, the WHO recommendation of 10 January was clear and confident: “WHO does not recommend any specific health measures for travellers. It is generally considered that entry screening offers little benefit, while requiring considerable resources. In case of symptoms suggestive to respiratory illness before, during or after travel, the travellers are encouraged to seek medical attention and share travel history with their health care provider. WHO advises against the application of any travel or trade restrictions on China based on the information currently available on this event.”⁸

On 27 January, the WHO issued an updated travel advice in which it no longer repeated the claim that there was no evidence of human-to-human transmission, but reiterated that “WHO advises against the application of any restrictions of international traffic based on the information currently available on this event.” The only issue on which the WHO changed its position somewhat was that of health screening at ports of exit and entry, which at least it did not explicitly advise against anymore. However, it fell far short of recommending such measures, pointing out that “temperature

⁷ <https://www.who.int/ihr/publications/9789241580496/en/>

⁸ <https://www.who.int/news-room/articles-detail/who-advice-for-international-travel-and-trade-in-relation-to-the-outbreak-of-pneumonia-caused-by-a-new-coronavirus-in-china/>

screening to detect potential suspect cases at Point of Entry may miss travellers incubating the disease or travellers concealing fever during travel and may require substantial investments.”⁹

The WHO’s categorical advice against travel restrictions at this point in time was remarkable, in view of the fact that four days earlier, the Chinese authorities had banned all domestic travel to and from Hubei, the province of which Wuhan is the capital. The WHO’s representative in Beijing, Gauden Galea, did not try to hide the fact that the Chinese measure flatly went against the WHO’s recommendations: “The lockdown of 11 million people is unprecedented in public health history, so it is certainly not a recommendation the WHO has made.” He ended up commending the measure nonetheless, because it is “a very important indication of the commitment to contain the epidemic in the place where it is most concentrated.”¹⁰ A month later, on 24 February, Bruce Aylward, Head of a WHO team that visited Wuhan, admitted that the Wuhan lockdown had succeeded in limiting the global spread of the virus, adding that “the world is in your debt.”

But even as it applauded the Chinese authorities for the success of their domestic travel restrictions, the WHO’s line on international travel remained unaltered. In the third version of its advisory on international travel of 29 February, the organization again left no doubt about the futility of travel restrictions: “WHO continues to advise against the application of travel or trade restrictions to countries experiencing COVID-19 outbreaks.” The document also criticized countries that had implemented travel bans:

“Travel bans to affected areas or denial of entry to passengers coming from affected areas are usually not effective in preventing the importation of cases but may have a significant economic and social impact. Since WHO declaration of a public health emergency of international concern in relation to COVID-19, and as of 27 February, 38 countries have reported to WHO additional health measures that significantly interfere with international traffic in relation to travel

⁹ <https://www.who.int/news-room/articles-detail/updated-who-advice-for-international-traffic-in-relation-to-the-outbreak-of-the-novel-coronavirus-2019-ncov>

¹⁰ <https://www.reuters.com/article/us-china-health-who-idUSKBN1ZM1G9>

to and from China or other countries, ranging from denial of entry of passengers, visa restrictions or quarantine for returning travellers. Several countries that denied entry of travellers or who have suspended the flights to and from China or other affected countries, are now reporting cases of COVID-19.”¹¹

1.2. A short history of travel restrictions during the COVID-19 pandemic

As the WHO’s irritation about countries ignoring its advice against international travel restrictions indicates, a fair number of countries chose not to follow its recommendations. By the end of January 2020, 21 countries had implemented entry restrictions for travellers from China, in the form of entry bans, mandatory 14-day quarantines or a combination of the two. These included several small island states, such as the Federated States of Micronesia (the first country to restrict travel from China as early as 6 January), Jamaica, and Trinidad and Tobago. Others included Singapore, Georgia, El Salvador, Mongolia, Morocco, and from 31 January also Italy. The fact that the latter country became from the end of February onwards the next hotspot of the pandemic seemed to prove the WHO’s point that travel bans are futile. I will come back to the case of Italy later, but for the moment it suffices to point out that from a statistical point of view, one case never suffices to accept or reject a hypothesis. What matters is whether controlling for other relevant influences on the spread and severity of the pandemic, and averaging out the random effects of sheer good or bad luck, countries that implemented early travel restrictions were able to better contain the spread of the pandemic than countries that did not implement such restrictions or did so relatively late.

¹¹ <https://www.who.int/news-room/articles-detail/updated-who-recommendations-for-international-traffic-in-relation-to-COVID-19-outbreak>. The document also reiterated the WHO’s skepticism regarding screening: “Temperature screening alone, at exit or entry, is not an effective way to stop international spread, since infected individuals may be in incubation period, may not express apparent symptoms early on in the course of the disease, or may dissimulate fever through the use of antipyretics; in addition, such measures require substantial investments for what may bear little benefits.”

Other countries followed with entry restrictions for travellers from China in early February, including Australia (1/2), the United States (2/2), Israel (2/2), the Philippines (2/2), New Zealand (3/2), India (5/2), Indonesia (5/2), and Saudi Arabia (6/2). By the end of February, the number of countries with entry restrictions applying to the whole of China (not counting many other countries with restrictions limited to Hubei province) had risen to 50. Remarkably absent were the countries of the European Union and the Schengen Area. Until the end of February, apart from Italy, only the Czech Republic (8/2) had restricted travel from entire mainland China. The Online Appendix to this paper offers a fully sourced overview of the timing of travel restrictions on mainland China, on Italy, as well as general entry restrictions for independent states across the world.¹²

While the spread of the outbreak in China was slowing down considerably as a result of the lockdown measures, Iran and Italy developed into the next hotspots during the month of February. Because of its relative international isolation, Iran was relevant as a source of early COVID-19 introductions especially in neighbouring countries such as Azerbaijan, Afghanistan and Iraq, as well as in Lebanon, which has strong ties to Iran through its Shiite minority. Italy, however, being much more central to the international travel network, succeeded China as the most important source of new virus introductions, either by Italians who travelled abroad, or by travellers who had visited Italy and returned to their home countries. For instance, the origin of the first COVID-19 infection is known for 39 out of 41 European countries represented in my dataset. In 22 countries (56% of known origins), the first case originated in travellers who were from or had been to Italy; in 9 countries (23%), the first cases were travellers from China; and in 8 countries (21%), they had arrived from other countries (including Iran, Germany, and Spain). Among second and third chains of infection with known origins, the relative importance of Italy as a source country increased further to 64 percent and 68 percent, respectively, whereas imports from China declined to 15 percent and 4 percent.¹³ A detailed study

¹² The online appendix can be downloaded at: <https://wzb.eu/de/media/63618> . There you will also find a Stata dataset containing all the variables used in the analysis.

¹³ These data were collected from publicly available sources. Coded were not the first, second, and third cases in a country, as these may all be part of the same chain of transmission. For instance, the first three cases in Germany all originated in China, and

for Romania (Hâncean, Perc & Lerner 2020) shows that of the first 147 COVID-19 cases in that country, 88 were imported, and among these 64 came from Italy.

On 21 February, the first Italian corona death was registered in Veneto province in the North of the country. Within days, first countries around the world introduced entry restrictions for travellers from Italy. Mauritius, Samoa and Kuwait were the first to do so on 24 February. The passengers of an airplane that was already on its way from Italy to Mauritius as the measure was announced were given the choice between returning to Italy or spending 14 days in quarantine in Mauritius. In the next days, among others Jordan (25/2), Iraq (25/2), Israel (26/2), Jamaica (27/2), Lebanon (28/2), and Turkey (29/2) followed. By the end of February, 17 countries had travel restrictions (entry bans and/or quarantines) for travellers from Italy. Among them was not a single European country, even though they ran the greatest risk of introduction of the virus from Italy. On 2 March, Romania – which has about 1.2 million nationals living in Italy – and Iceland were the first European nations to introduce entry restrictions. Iceland did so after the country's first infections had been registered among a group of ski holiday makers returning from Northern Italy. Russia and the Czech Republic followed on 5 and 7 March, respectively.

By 8 March, 366 people had died in Italy and infections had spread to large parts of the country. That day, the government ordered a lockdown of three zones in the North with a total of 16 million inhabitants, including the region's largest city, Milano. Because Northerners immediately began trying to escape the lockdown zone by car or public transport, the government extended the lockdown to the whole country the next day. But while domestic travel in Italy had come to a complete standstill, the country's borders to its neighbours were still open, and international flights from and to most other European countries continued. On the day of

could be traced back to an employee of a car parts manufacturer near Munich, who had recently returned from a business trip to Wuhan and was diagnosed with Covid-19 on 27 January. Instead, we coded the sources of the first, second, and third independent chains of transmission, i.e., the first cases in Germany mentioned before would be summarized as a first chain originating in China. Germany's second independent chain (diagnosed on 2 February) also originated in China, whereas the third chain (diagnosed on 25 February) originated in Italy. I thank Gizem Ünsal for coding these data.

the national lockdown, Albania – another country with a large emigrant population in Italy – closed its borders to travellers from Italy, and a day later, on 10 March, Austria and Slovenia, two of the countries that directly border Northern Italy, closed their frontiers. Spain stopped air traffic to and from Italy on the same day, but unlike Austria and Slovenia, it did so at a moment when the epidemic had begun spiralling out of control within Spain and had claimed 36 confirmed deaths already.

On 12 March, US President Trump announced that the US would restrict, effective the next day, entry from the 26 countries of the Schengen Area, in which internal border controls have been abolished. After critical questions as to why these countries were not included, the United Kingdom and Ireland were added a few days later. The US entry ban drew widespread criticism in Europe. European Commission and European Council Presidents Ursula von der Leyen and Charles Michel jointly stated that “The Coronavirus is a global crisis, not limited to any continent and it requires cooperation rather than unilateral action.” The next day, Von der Leyen affirmed the EU’s rejection of travel restrictions, and once again referred to the WHO: “Certain controls may be justified, but general travel bans are not seen as being the most effective by the World Health Organization.”

Whereas with the China entry ban of 2 February the United States were comparatively early, with its 13 March ban against Italy and the rest of the Schengen zone, it was not part of the avant garde: by 12 March, 54 countries had already implemented entry restrictions against Italy (and some also against other European countries such as France and Spain). Moreover, with the exception of Spain, almost all these countries implemented restrictions on travel from Italy before the epidemic had started to claim a significant domestic death toll. Most of them had no deaths at the time of the entry restrictions (some did not even have any registered cases), a few others like Norway, Egypt and Argentina had only one fatality, and Australia had three. The death toll in the United States, by contrast, had already reached 48 when the travel ban on the Schengen zone went into effect. The Trump administration’s mistake may not have been the Schengen travel ban as such, but the fact that it was introduced too late.

What the US decision did achieve, however, was to set in motion a cascade of further entry restrictions. In Europe, Switzerland closed its borders to Italy on the 13th. Other countries halted commercial air travel to and from Italy, among others Croatia (13/3), the Netherlands (13/3) and Greece (14/3). Some countries closed their borders entirely. Slovakia on the 13th, followed by Denmark (14/3), Cyprus (15/3) and Poland (16/3). Even Germany, which had long been a vocal opponent of the idea of travel restrictions, especially within the Schengen free movement zone, closed its borders to Austria, Switzerland and France on the 16th. A day later, the European Union made a 180 degree turn, as well. Only five days after rejecting the US entry ban categorically and after almost two months of insisting that travel restrictions were unnecessary, ineffective, or even counterproductive, the EU issued a blanket entry ban against travellers from all countries outside the EU or the Schengen Area, exempting only the United Kingdom. If ever there was a badly targeted entry ban, it must have been this one, because the epicentre of the pandemic at that moment was overwhelmingly situated within Europe rather than outside it. Nonetheless, the EU entry ban may have made an important contribution to limiting the pandemic in an unintended way, namely by protecting the rest of the world from further infections brought by travellers from Europe. Within a week after the US decision, the number of countries with travel restrictions on Italy had doubled to 109, by the end of March it had reached 121. With travel bans in place against China and Europe in the United States, against the rest of the world in Europe, and elsewhere against China, Italy, and increasingly also against other European hotspots such as Spain, France and the United Kingdom, as well as against the United States, international commercial travel had come to an almost complete standstill. Only a few countries, most importantly the United Kingdom and Mexico, never issued entry bans, but even their international air traffic was strongly reduced because of other countries' travel bans.

Much of the remaining passenger air traffic concerned repatriation flights for stranded citizens abroad. Here too, an important question arose, namely whether or not returning citizens and permanent residents (who in many countries were exempted from entry bans) should be submitted to a mandatory quarantine (usually of 14 days). Some countries introduced quarantines early on (simultaneous to entry bans or even prior to them),

others did so at a relatively late date, e.g. Germany on 16 April and the United Kingdom as late as 8 June. Still others, such as Sweden, the Netherlands, Brazil, and Mexico, never had a mandatory quarantine requirement at all.

2.Theoretical framework

2.1. The strength of cutting weak ties: what sociology teaches us about diffusion in social networks

The autonomous capacity of viruses for mobility is zero. They have neither legs, wings nor wheels, and cannot get from A to B without the help of human hosts and human-made means of transport. From an evolutionary point of view, the aim of a virus is not to harm or kill its host, but, just like every other biological organism, to make as many copies of itself as possible, and this it achieves by using the host to infect other individuals. The most successful viruses are not those that quickly cause serious illness or even death in their hosts, because hosts will then soon be immobilized and others would likely take precautions in view of the patient's serious illness or death. The most dangerous viruses are those that are infectious even if the infected person does not (yet) have serious symptoms or is entirely asymptomatic. COVID-19 is precisely such a virus. Because, at least initially, it does not constrain a host's mobility and does not give away early warning signs that would make others cautious in their social contacts with the infected person, it is perfectly equipped to spread widely by way of human interaction and mobility. This is a key difference between COVID-19 and earlier coronavirus epidemics such as SARS and MERS, which had much higher rates of serious illness and death (with fatality rates of 10% for SARS and 34% for MERS), but were not nearly as infectious, especially because transmission rarely occurred asymptotically.¹⁴ These differences between COVID-19 and earlier pandemics were known early on, and should have alerted the WHO and other responsible authorities that conclusions drawn from these earlier pandemics, especially when related to

¹⁴ <https://www.medicalnewstoday.com/articles/how-do-sars-and-mers-compare-with-covid-19>.

the virus' potential for global diffusion through international networks, were not necessarily applicable to COVID-19.

In sociology, processes of diffusion across networks also play an important role. Many social phenomena, as diverse as fashion trends, rumours, news, technological innovations, information about jobs, riots, political ideologies, and business models originate in a certain location and may from there diffuse across social space (epidemiologists have been aware of this parallel between epidemic and social diffusion; e.g., Brockmann and Helbing 2013). Some of these social innovations and information items stay confined to limited social circles, others spread globally; some spread very rapidly, others very slowly or not at all. The adoption of social innovations is in many ways much more complex than the spread of a virus because humans do not choose to adopt a virus or not, whereas in social diffusion, human transmitters and recipients play an active role in choosing which information and innovations to transmit, adopt and discard. Among other things, the perceived utility of the innovation to the adopter and the social relationship between source and adopter play important roles.

One important determinant of the course of diffusion processes operates however in very similar ways for viral and social contagion, namely the structure of social networks. If social networks were unstructured, every individual would be directly connected to everybody else. That is a far cry from the real social world, in which networks are extremely sparse, in the sense that only a small fraction of all possible direct connections between people are in fact activated. Even if we add together all our family members, friends, acquaintances, colleagues and even all the people we casually meet in the street, each of us is directly connected to only a tiny fraction of the seven billion or so people who inhabit our planet. These contacts, moreover, differ in what US American sociologist Mark Granovetter has called their "strength"; defined by aspects such as the frequency, intensity, and the degree of social exchange that occurs across a tie. In his classic paper "The strength of weak ties" (one of the most-cited – perhaps *the* most cited – paper in sociology with 56,000 citations in Google Scholar to date) Granovetter makes a number of important observations about differences between social ties that are "strong" and "weak" in the above-defined sense.

Individuals who are connected by a strong tie have a much higher likelihood of sharing contacts to the same third parties than individuals who are connected by a weak tie. Because, according to the principle of social “homophily” (McPherson et al. 2001), people prefer associating with people who are similar to themselves (regarding age, ethnicity, level of education, cultural tastes, political preferences, religion and the like), people connected by strong ties also tend to be more similar to each other than acquaintances linked by a weak tie. Further, because the strength of ties depends in part on the frequency of interaction, strong-tie networks tend to be clustered in geographical space. Strong-tie networks are clusters in social space that are internally highly integrated (in the sense that many of the possible ties between individuals are activated) and at the same time to some degree segregated from other strong-tie networks. The network links that connect strong-tie clusters to each other Granovetter calls “bridges” and they are almost always weak ties (because if the tie between A and B were strong, it is highly unlikely that none of the other strong ties of B would know A).

This bridging function makes weak ties crucially important for diffusion processes. Weak ties, Granovetter writes, “create more, and shorter, paths. Any given tie may, hypothetically, be removed from a network; The contention here is that removal of the average weak tie would do more “damage” to transmission probabilities than would that of the average strong one.” (Granovetter 1973: 1365–66)

“Intuitively speaking, this means that whatever is to be diffused can reach a larger number of people, and traverse greater social distance (i.e., path length), when passed through weak ties rather than strong. If one tells a rumor to all his close friends, and they do likewise, many will hear the rumor a second or a third time, since those linked by strong ties tend to share friends. If the motivation to spread the rumor is dampened a bit on each wave of telling, then the rumor moving through strong ties is much more likely to be limited to a few cliques than that going via weak ones; bridges will not be crossed.” (ibidem: 1366)

In the above quotation, Granovetter uses the word “damage” when referring to removing weak ties from a network, because most of the applications of diffusion and network theory in sociology and economics study things that are socially useful, such as technological innovations or valuable information. This is also what the title of his paper, “the strength of weak ties”, refers to. The empirical illustration that Granovetter gives from his own research is an example: people are much more likely to obtain information about opportunities on the labour market (e.g., about vacancies) through their weak-tie acquaintances. Most of their socially proximate strong-tie contacts share the same information (and information deficits) that they themselves already have, whereas new information is much more likely to come from more socially distant weak ties. Individuals that lack such weak ties to distant others outside their own social group are consequently seriously harmed in their labour-market opportunities. In my own field of empirical research, a typical example are the labour-market chances of immigrants, which are especially enhanced by (usually weak) social contacts to people outside their own ethnic group (in this field often referred to as “bridging ties”) whereas immigrants whose contacts are largely limited to members of the own ethnic group (referred to as “bonding ties”) tend to do worse on the labour market (see, e.g. Lancee 2010, Koopmans 2016).

However, when dealing with the diffusion of something harmful, such as a deadly virus, the strength of weak ties become a weakness, and conversely it is the cutting or removal of such ties that becomes a strength, hence the title of this section: “the strength of cutting weak ties”. Just as in Granovetter’s example of a rumour, a virus that spreads in a strong-tie network will sooner or later encounter redundancies, meaning that its chances of infecting new individuals will decline because these individuals have already been infected by others in the social network. If one member of a family household contracts an infectious virus, it is, because of the strength of ties within the family, highly likely that other members of the household will contract it as well. In the next instance, however, the persons that the newly infected individuals are in closest and most frequent contact with, namely the members of that same household, are likely to have already been infected by the first patient. To diffuse further, the virus must break out of the strong-tie cluster of the household via a

bridging weak tie, for instance when one of the family members carries the virus to work and infects one of her colleagues. There, however, a similar process repeats itself. Work colleagues may not be as closely tied as members of the same household, but the intensity of their interaction still is far greater than that between random members of society. The first infected colleague will find many potential colleagues to infect, but in the next round, opportunities for further infection start to decline because some of the most likely candidates for infection have already been infected by the first patient in the firm.

An additional reason why infections can spread more rapidly across weak-tie links is that within strong-tie networks, people regularly interact and are in a position to monitor each other. As soon as one person in a strong-tie network falls ill, the other members of the network will be alerted and – especially when this occurs in the context of an emerging deadly pandemic – will take precautionary measures to prevent that they themselves and others in the network become infected. They will also monitor their own health status and may choose to test themselves. In many ways, strong-tie networks enable the kind of monitoring that recently developed “corona apps” aim to achieve: to know whether one has been in contact with an infected person and to take measures accordingly. In weak-tie networks, such monitoring is impossible without the help of technology precisely because of the weakness of the ties. The more distant and casual the relationship, the less likely one is to ever learn that someone one has been in contact with has fallen ill or has tested positive.

This argument can be extended on and on to ever wider circles of social interaction, but the important point is that diffusion in a social network is not unconstrained but rather has to pass through a continuous succession of social bottlenecks, when whatever diffuses reaches the boundaries of strong-tie clusters and needs to find weak-tie bridges to diffuse more widely in society. This can be illustrated with the metaphor of water lilies in a pond that is often used to illustrate the accelerating speed of exponential growth. “If the number of lilies in the pond doubles every 24 hours and it takes ten days for the pond to be entirely covered with lilies, how many days did it take for half the pond to be covered?” The correct answer is of course “nine days”. The implicit assumption of the example is

that society is one big, unstructured pond, in which there are no limitations for lilies – read: a virus infection – to double every 24 hours. But what the social network perspective of Granovetter teaches us, is, that in reality, society consists of numerous small ponds (strong-tie networks), connected by relatively narrow bridges (weak ties) or sometimes not connected directly at all, but only through several intermediary little ponds, which are in turn connected by narrow social bridges. In such an environment, diffusion will not occur at a constant rate, but will rather proceed in fits and starts and at an average rate that may be much below the theoretical rate of diffusion if society were one big unstructured pool in which expansion would always be possible in all directions without any constraints.

The point of lockdowns of whatever kind is to reduce the number of active connections in a social network. What the sociological perspective on network diffusion makes clear is that a smart and effective lockdown strategy is one that cuts the weak ties, the social bridges in the social network, which link otherwise distant individuals and groups. Not all lockdown measures are equal in this regard. Banning large gatherings and events, such as concerts, festivals and fairs, is a measure that targets weak ties because at such events many people come together that are not otherwise in contact with each other and who may live dispersed over a large geographical area. Closing kindergartens is a measure that targets much stronger social ties: they not only bring together much smaller numbers of people than large events, they are also highly localized in social and geographical space. To stay in terms of our water lily metaphor, banning large events targets the bridges, closing kindergartens targets the small ponds.

In a recent simulation study, Block et al. have applied these sociological insights to different strategies for social distancing. They find that compared to a “naïve” social distancing strategy in which individuals reduce their social contacts in a random fashion, strategies that focus on cutting weak ties (e.g., to socially dissimilar individuals) or of restricting social contacts to strong-tie networks (e.g., to good friends whose friendship networks overlap strongly with one’s own) have the capacity of delaying the peak of infections by 37 percent and reducing the height of

the infection peak by 60 percent. A much greater effect, still, is achieved by the cognitively least demanding strategy of all, namely to limit contacts to geographically proximate others (Block et al. 2020, extended data Figure 1b). This social distancing strategy is effectively equivalent to policies that restrict movement and travel between localities (within cities, between regions, or between countries).

Sociological network theory therefore predicts that the most effective measures to contain social diffusion of any kind, including that of a virus pandemic, are those that target the weakest among the weak ties: those that connect people who are maximally distant from one another. “Distance” in a sociological sense has many dimensions, including segregation along lines of social class or cultural difference. However, apart from a small cosmopolitan class of frequent travellers who entertain many and relatively intense transnational network ties (e.g., Calhoun 2002), most people’s social interactions are overwhelmingly confined within the boundaries of old-fashioned nation-states. This is true even in our age of globalization and even if we consider forms of social interaction that do not require physical contact (and that are therefore irrelevant from the point of view of the spread of a virus). Of all phone calls made worldwide, only 5 percent are international, 16 percent of Facebook friends are foreign, and 17 percent of overall internet traffic crosses international borders (Ghemawat & Altman 2014: 12). Indicators of face-to-face interaction across international borders show even lower levels: foreign students make up two percent of all university students worldwide and only three percent of the world population lives in a country different from the one where they were born (and many of them never or rarely return to their country of birth). For tourism (based on overnight stays), the international share is higher, but even that does not exceed 15–18 percent according to estimates of the World Tourism Organization (*ibidem*: 53). International contacts will also, on average, be much less intense than domestic contacts for the simple reason that physical distance makes frequent and intimate interaction difficult.

But precisely because they are weak ties, international social connections offer a virus the opportunity to diffuse much more rapidly and widely than transmission through local or national social ties allows. This is true a

fortiori for international tourists, who when travelling make weak-tie connections to numerous people in many different places in the countries they visit. They stay in hotels and go out dining, they visit museums and tourist attractions and many of them do not stay put in one place, but visit different localities in the country of destination or even across several countries during one trip. While doing so, they not only have casual face-to-face interactions with many locals – waiters, shopkeepers, museum attendants, etc. – but also with many other tourists from countries all over the globe. The first COVID-19 cases that were discovered in Italy, a Chinese tourist couple, are a case in point. They were in Rome when they were hospitalized on 29 January 2020, but they had arrived in Italy at Milan’s Malpensa airport and had made visits to Verona and Parma on their way to Rome. The international airports that tourists pass through, especially if they are hubs that connect many different airports around the world, provide major opportunity structures for the spread of a virus.

2.2. International seeds and domestic growth

As a result of the virtue signalling culture that is unfortunately rampant in academia, some of my colleagues regard the following truism as a dangerous statement that, if it is not already xenophobic itself, at least encourages a xenophobic interpretation. The truism is that in a global pandemic that originates in a certain location somewhere on earth, the import of cases into another country by definition occurs by way of people who cross borders, be they tourists, businesspersons or migrants. In epidemiology, by contrast, scholars do not hesitate to state what counts as blasphemy for many in the social sciences: “When people move, they take contagious diseases with them. Their movements are thus a harbinger of the future states of an epidemic” (Jia et al. 2020: 391). Judging from what we know about the first COVID-19 infections that occurred in various countries, among the different types of border-crossers, tourists were the most important source. The truism, by the way, does not imply that the people who introduce a virus into a country are always or even predominantly foreigners in the sense of citizenship and/or residence. From what we know the virus was as often spread by residents or citizens

returning to their country after a foreign trip as by foreigners visiting another country.

From the above, it follows logically that travel restrictions must have a significant or even decisive impact when they are introduced before the first cases of a virus in a country occur. When such restrictions are well-targeted and implemented strictly enough, they can even prevent the introduction of the virus altogether. In the case of COVID-19, there are a number of countries that illustrate this. Nine independent island states in the Pacific – Samoa, Marshall Islands, Vanuatu, Palau, Nauru, Kiribati, Tuvalu, Tonga and the Federated States of Micronesia – have at the time of writing (4 October 2020) not registered a single case of COVID-19. All of them introduced entry bans and/or mandatory quarantines for people entering from China, including their own citizens, very early on, and quickly followed on with restrictions for people seeking to enter from other virus hotspots such as South Korea and Italy, or simply from any country that had registered cases of COVID-19. Micronesia was the first country in the world to ban entry from mainland China, as early as 6 January. Quarantine requirements in most of these countries moreover stipulated that the mandatory 14 days should not be spent locally but prior to arrival in a third country entirely free from COVID-19 cases. Before being allowed entry, travellers needed to show a recognized doctor's certificate of a recent negative COVID-19 test. Several of these countries additionally required ships delivering cargo to have been at sea for at least 14 days before they were allowed entry to port. Of course, in implementing these measures effectively and preventing anyone to circumvent them, these countries were greatly helped by the fact that they are island states protected by the natural borders of the ocean. Because they share no land border with any other state, they are moreover shielded from spillover effects that may occur if a neighbouring country experiences a major outbreak and/or does not police its borders sufficiently.

But what if you are not an island state and if you already have one or more COVID-19 cases, do travel restrictions then still make sense? This question is the main point of debate in the literature. For understanding this debate, it is important to distinguish two components that contribute to the spread of a virus in a given country. The first is the number of “seeds” or

importations, i.e. infections that are imported by travellers from abroad into a country. These seeds subsequently become starting points of independent domestic chains of transmission, which constitute the second component of the spread of the virus. Once cases of a virus have been introduced into a country, domestic transmission, if unchecked, proceeds exponentially, the speed depending on the contagion factor R . From the combination of R with the average incubation time, a virus's doubling time can be calculated, which in the initial phases of the COVID-19 outbreak in China was estimated to be about four days (Chinazzi et al. 2020: 2). Given such a doubling time, one single initial case would have multiplied more than thousand-fold 40 days later.

Precisely because international ties are weak ties, the number of introductions from abroad can be quickly dwarfed by domestic transmission once the latter has passed a certain threshold. A study (Pinotti et al. 2020) on international transmission in the early phase of the COVID-19 pandemic estimated that between 3 January and 13 February only 288 cases of international importation of the virus had occurred, most of them from the Chinese hotspot province of Hubei, a smaller number from the rest of China, and an even smaller, but towards the end of the period increasing number from other countries such as South Korea. By the end of that same period, China had registered 64,000 corona cases. Yet, it is from the relatively few cases of transmission out of China that all the millions of registered corona cases worldwide have sprung.

International importations of COVID-19 are thus relatively rare occurrences compared to the number of domestic transmissions. But as is typical of the dynamics of complex systems, rare events can have huge consequences. Here is something important to keep in mind: even if it is probably true that of all the COVID-19 cases in any given country at any given point in time (except in the early stages of a local epidemic) 99 percent or more are the result of domestic transmission, this does not imply that domestic transmission is the key to explaining why a certain country has more COVID-19 deaths than another. To the contrary, it may well be that the explanation to an important extent lies in the relatively few cases that were the result of international importation since they are

the seeds that lie at the root of each independent chain of domestic transmission.

In a study simulating the impact of different containment measures on the spread of a hypothetical novel influenza epidemic in the United States and the United Kingdom, Ferguson et al. (2006) calculated that border restrictions would be able to delay the domestic peak of the pandemic by 12.5 days if they would achieve a 90 percent reduction of international air traffic, and three weeks if the reduction amounted to 99 percent. A meta-analytic overview of several studies (overwhelmingly mathematical simulations) estimated the delaying effect of 90 percent effective international travel restrictions at three to four weeks, and up to four months for more extensive restrictions (Mateus et al. 2014: 877). To put these figures into perspective, during the COVID-19 pandemic, by 7 April 2020 the number of air passengers, domestic and international combined, was down by 95 percent in the United States.¹⁵ In Germany, the number of passengers in the first week of April was down by 98 percent compared to the same time the previous year, and this was still in the period when repatriation flights were bringing back stranded Germans from abroad.¹⁶ Germany and the United States were moreover among the countries that did not halt international commercial air traffic altogether. In countries that did, reductions up to 100 percent may have been achieved.

Even so, one might argue that if a few weeks or months delay is all one can gain by reducing international air traffic, it is not a very effective measure. Indeed, once a virus has been introduced into a country and assuming that nothing else would be done, closing borders would only delay but not prevent a full-scale outbreak (Mateus et al. 2014). However, in a global pandemic time is of the essence. For instance, the 12.5 days delay calculated in the British study (Ferguson et al. 2006) amounts to three times the estimated unconstrained doubling time of COVID-19, which implies that at the end of that period there would be eight times more cases than at the beginning. In the three weeks that a 99 percent reduction of international traffic would buy according to this study, the virus would even have

¹⁵ <https://www.cntraveler.com/story/coronavirus-air-travel-these-numbers-show-the-massive-impact-of-the-pandemic?verso=true>

¹⁶ <https://www.tagesschau.de/wirtschaft/corona-flughafen-passagiere-101.html>

doubled five times and increased 32-fold. In other words, two or three weeks more or less time can make a huge difference in the early stages of a pandemic. Once a virus is spreading domestically there is no way around measures that can “flatten the curve”, but, assuming the same domestic containment measures introduced at the same date and with the same level of efficiency, it is likely to make a big difference in terms of the level at which the curve is eventually flattened whether one starts the process with a curve that has already reached an eight to 32 times higher level.¹⁷ Moreover, some of the most effective ways of limiting domestic transmission, such as contact tracing and isolation, quickly become unfeasible once case numbers go up (Leung et al. 2020: 1389; Pinotti et al. 2020; Maier & Brockmann 2020: 1). Epstein et al. (2014: 2) therefore conclude. “Since the benefit of travel restrictions can be substantial while their costs are minimal, dismissal of travel restrictions as an aid in dealing with a global pandemic seems premature.”¹⁸

The first major study of the impact of travel restrictions on the spread of the COVID-19 pandemic was published on 6 March 2020, and was at the time widely cited in the media as confirming the WHO’s view that travel restrictions were not very effective.¹⁹ This study (Chinazzi et al. 2020) is based on simulated scenarios from a statistical model calibrated with real epidemic data on the diffusion of the virus before the imposition of a domestic travel ban on the city of Wuhan and Hubei province on 23 January. Further travel restrictions that followed shortly after the Wuhan lockdown were the suspension of flights to mainland China by 59 airlines (including Lufthansa) until the end of January, as well as comprehensive bans until that date on travel from and to China by a number of countries (including the United States, Australia, and Italy). Compared to the extrapolations from a model without travel restrictions the authors find a 77 percent reduction in cases internationally imported from mainland China until early February (Chinazzi et al. 2020: 2). This reduction is

¹⁷ As Epstein et al. (2007: 2) put it: “From a public health perspective, it becomes clear that the *main* purpose of travel restrictions is to delay dissemination of the disease until targeted medical and other interventions can be developed and deployed.”

¹⁸ For the United States, Epstein et al. (2014: 14-15) estimate the economic cost of entirely shutting down the airline industry to be at most one percent of yearly GNP.

¹⁹ See for instance https://www.focus.de/reisen/reise-news/reisen-trotz-corona-was-fuer-baldige-grenzoeffnungen-spricht-und-was-dagegen_id_11979732.html

particularly due to a reduction of imports from the lockdown city of Wuhan, which before 23 January was the source of 86 percent of all international importations, and after that date does not even appear among the top-10 of Chinese source cities (ibidem: 2, 10).

Chinazzi et al.'s simulations of the effect of reductions of international air traffic predict an initial tenfold reduction in the number of cases of international importation if the volume of air travel is reduced by 90 percent. This prediction indeed fits available data on international importations from China until mid-February (Pinotti et al. 2020). Yet, Chinazzi et al. concluded from their simulations that “while the Wuhan travel ban was initially effective at reducing international case importations, the number of cases observed outside Mainland China will resume its growth after 2-3 weeks from cases that originated elsewhere ... Travel limitations up to 90% of the traffic have a modest effect unless paired with public health interventions and behavioural changes that achieve a considerable reduction in disease transmissibility.... Moving forward we expect that travel restrictions to COVID-19 affected areas will have modest effects, and that transmission-reduction interventions will provide the greatest benefit to mitigate the epidemic” (Chinazzi et al. 2020: 4).

This conclusion, again, seems to underestimate the importance of gaining two to three weeks of time, which is the equivalent of 3-5 doubling times of the virus. Moreover, it hardly seems a realistic scenario that international travel is curbed but nothing else is undertaken in order to limit domestic transmission. Empirically at least, this seems to have been a very rare scenario during the COVID-19 pandemic. Chinazzi et al.'s own simulations show that a combination of a 90 percent travel reduction even with a moderate domestic transmission reduction (in which the rate of transmission is reduced by only 25%) would have the effect that “the epidemic peak [in mainland China, RK] is delayed to late June 2020,” with a greatly reduced number of international importations. How important timing of interventions can be, is demonstrated by very similar simulation models of the Chinese pandemic by Lau et al. (2020). They show that if the package of interventions that was introduced on 23 January had been introduced one, two or three weeks earlier, the number of cases in

mainland China could have been reduced by 66%, 86% or 95%, respectively, and the number of affected cities would have declined from 308 to 192, 130 or 61. Conversely, delay of the measures by one, two or three weeks would have resulted in 3 to 18-fold increases in the number of cases (ibidem: 2).

Whereas the studies discussed so far are all simulation-based and use forecasting models to derive estimates of the effects of travel restrictions, Kraemer et al.'s study, published in early May 2020, performs regression analyses (poisson, negative binomial, and log-linear) using real mobility data from the Baidu network to predict the actual development of COVID-19 case numbers across locations in China until mid-February. They find that prior to the 23 January travel restrictions, travel out of Wuhan predicted the number of cases elsewhere in China extremely well ($R^2 = .89$). After the Wuhan travel ban, growth rates elsewhere strongly declined and eventually became negative. The authors test whether this development can alternatively be explained by increased testing capacity across China over the investigated period but they find that mobility from Wuhan and its restriction after 23 January offer for the large majority of provinces a much better explanation (Kraemer et al. 2020: 3). Jia et al. (2020) similarly show that travel flows out of Wuhan to 296 prefectures across China until 24 January (the day after the Wuhan travel ban) provide by far the strongest predictor (compared e.g., to GDP or population size of a prefecture) of the cumulative case count in these prefectures. Interestingly, the effect of population flows on cumulative counts increases over time, reaching an R^2 of over .90 by the end of the observed period, 19 February. According to the authors, this suggests, "that the spreading pattern of the virus gradually converged to the distribution of the population outflow from Wuhan to other prefectures in China" (Jia et al. 2020: 391).

Similar findings for China are presented by Lau et al (2020), who in addition show that case numbers in countries outside China until mid-February were very strongly related to air traffic passenger volumes from China. For past epidemics such as the 2003 SARS and the 2009 H1N1 epidemics, air traffic flows have also been shown to be very strong and consistent predictors of epidemic arrival times across countries (Brockmann & Helbing 2013). On a much smaller geographical scale, Harris (2020) has shown how passenger flows through the New York City subway

network offer a powerful explanation for differences in COVID-19 incidence across the city's zip code areas. Given the strong explanatory power of domestic and international travel flows for virus arrival times and cumulative case counts, it would seem logical to expect that, conversely, travel *restrictions* are also an effective tool for containing epidemics. Indeed, this is what Costantino et al. (2020) find in a study on the impact of Australia's 1 February entry ban on travellers from China, which the authors estimate to have reduced the number of cases and deaths by 87 percent compared to a scenario without any travel restrictions.

However, two recent cross-national studies conclude that compared to other, domestic, non-pharmaceutical interventions such as banning mass gatherings, school and workplace closures, curfews, and stay-at-home requirements, international travel restrictions only had a minor and mostly not statistically significant impact on the development of COVID-19 case numbers. Both studies reach this conclusion by looking at the impact of various measures on the subsequent growth rate of the virus, which is driven by the reproduction number R of the virus (i.e., the number of people that an average infected person will infect). Askitas et al. (2020) investigate the development of the number of COVID-19 cases across 135 countries worldwide, using data on government responses drawn from Oxford University's Coronavirus Government Response Tracker (Hale et al. 2020; see also my analyses using these data below). The latter data source includes information on the exact timing of different types of government responses, including the following used by Askitas et al.: public transport closures, cancellation of public events, restrictions on private gatherings, school closures, workplace closures, stay-at-home requirements, internal mobility restrictions (across regions and cities), and finally international travel controls (screening upon arrival, quarantines and partial and full border closures). As I will discuss below, this dataset has some methodological problems where the measurement of international travel controls is concerned, but for now, I will take the results obtained by Askitas et al. at face value.²⁰ To investigate the effects of policies, Askitas et

²⁰ One of these problems is that international travel controls are measured on a four-point "intensity" scale. Askitas et al. state that they transformed all policy scores in a scale running from 1 to 6, but this cannot possibly pertain to the international travel control scores, which run from 1 to 4 and are coded differently from most of the other policies.

al analyse whether compared to the trend of case numbers starting 20 days before the intervention up to 35 days after its introduction, the introduction of a measure leads to a significant decrease in the growth rate of case numbers, controlling for the effects of other policy measures and a number of country characteristics such as world region, GDP, and population density. They find that the three mobility measures (public transport restrictions, domestic mobility restrictions, and international travel controls) have no significant effects, whereas banning of public and private gatherings, as well as school closures did result in significant reductions of virus growth rates. The authors attribute the lack of a significant effect of international travel restrictions on such measures not having been taken early and stringently enough. “If countries have [sic] banned all international travel soon after the outbreak in China, it would have certainly be [sic] an effective measure to seal the country from the virus. However, because most countries did not introduce such bans before the virus had started spreading domestically, or they did introduce some restrictions but not complete bans, those restrictions had limited impact on mobility and could only reduce new imported infections but not contain the spread of the virus.” While the first part of this argument, referring to the efficacy of travel restrictions introduced before first cases have arrived, is certainly accurate, the second part is less convincing. After all, we have seen above that travel restrictions did lead to very strong reductions in mobility and to strong reductions of the number of interregional and international importations.

Bonardi et al. (2020) use a somewhat different approach. They study 184 countries between 31 December and 4 May and investigate the effect of two types of policies on the development of case numbers. “Internal measures” refer to the timing of curfews, the declaration of a state of emergency, and the first instance of regional lockdowns and “partial selective lockdowns”, the latter referring to a collection of measures including bans on public gatherings and school closures. “External measures” include the timing of border closures towards either Wuhan, the whole of China, Iran or Italy (“selective border close stage 1”) or towards any other country or group of countries (“selective border close stage 2”), as well as the full closure by a country of all international borders (“international lockdown”). Data on these policies were collected by the authors from media reports and US

embassy bulletins. These data seem to be very incomplete, though, as the authors state that out of the 184 countries, 108 (59%) had implemented at least one of the internal or external measures. However, the Oxford government response dataset registers partial or full border closures in 96 percent of the countries and territories included in that dataset. If we add other measures such as school closures and restrictions of public gathering that percentage increase to almost 100 percent: only one country (Nicaragua) in the Oxford dataset has not implemented any of these measures.²¹

Because their models include country fixed effects, Bonardi et al.'s analysis looks at within-country impacts of policies only, and like the Askitas et al. study, the dependent variable is the growth rate in the number of cases, comparing rates before and after a measure was taken. Their main finding is that internal, domestic measures (particularly regional lockdowns and what they call "partial selective lockdowns") were effective in reducing growth rates, whereas external measures related to international travel were not. "Even in a globalized world, internal policies are the name of the game," the authors sum up their results. They argue that the differential effect of external and internal measures on individual behaviour might be the reason: "internal measures are effective at reducing opportunity costs for people of going out during a partial lockdown, whereas outside measures do not have this effect."

The latter conclusion seems quite obvious. Why would one expect border closures to have any impact on the frequency and intensity of social contacts within a country? Both Askitas et al.'s and Bonardi et al.'s research design and their interpretations of their results suggest that they have not adequately theorized the way in which international importations of a virus, i.e. seeds, affect the development and outcomes of an epidemic. In both studies, the dependent variable is the growth rate, which is mainly determined by the number of infections caused by the average infectious person. Domestic lockdowns, closures, contact restrictions as well as

²¹ Similarly, in my own dataset only seven countries (out of 181) had neither entry bans for China, for Italy, nor for all countries of the world until the end of May 2020. Moreover, some of these had border closures towards other countries than China or Italy (e.g., Mexico towards the USA) that the Bonardi et al. data should have picked up.

hygiene measures such as masks and the two-meter distance rule directly target the reproduction rate R by reducing the number and potential infectiousness of social interactions. As I have already demonstrated above, international contacts and interactions make up only a small fraction of all social interactions in a society and therefore a reduction in their number cannot have much of an effect on reducing R . The only exception where international importations can have a sizeable impact on the growth rate is in the early stages of an epidemic, when such importations still make up a considerable proportion (or very initially even the entirety) of all cases.

Even in these initial stages, however, the contribution of international importations to growth is not through their effect on R , because the transmission rate is an almost entirely domestic process. To see why this is so, imagine that an infectious person enters a country and that right after his or her entry, the country closes its borders so effectively that no other infected person enters the country anymore. R , if no other measures are taken, will be exactly the same as it would have been without the border closure and the number of cases will approximately double every four days. Using the approach of the previously discussed studies that only look at growth rates, one could hardly have found anything else than that international travel restrictions do not matter. Domestic contact restrictions and hygiene measures, by contrast, do have the potential to reduce R if they are stringent enough and implemented effectively.

By contrast, the number of seeds brought into a country by way of international travel does have, as we have seen in several studies cited above, important effects on the absolute *levels* of an epidemic. Because each seed establishes an independent domestic chain of transmission, the difference between one seed in the case of an early entry ban or ten seeds because of a less timely introduction of travel restrictions is not just a difference of nine cases. It is effectively the difference between having one or ten independent chains of transmission and will ultimately result in much higher case numbers and deaths.

Figure 1: Theoretical model

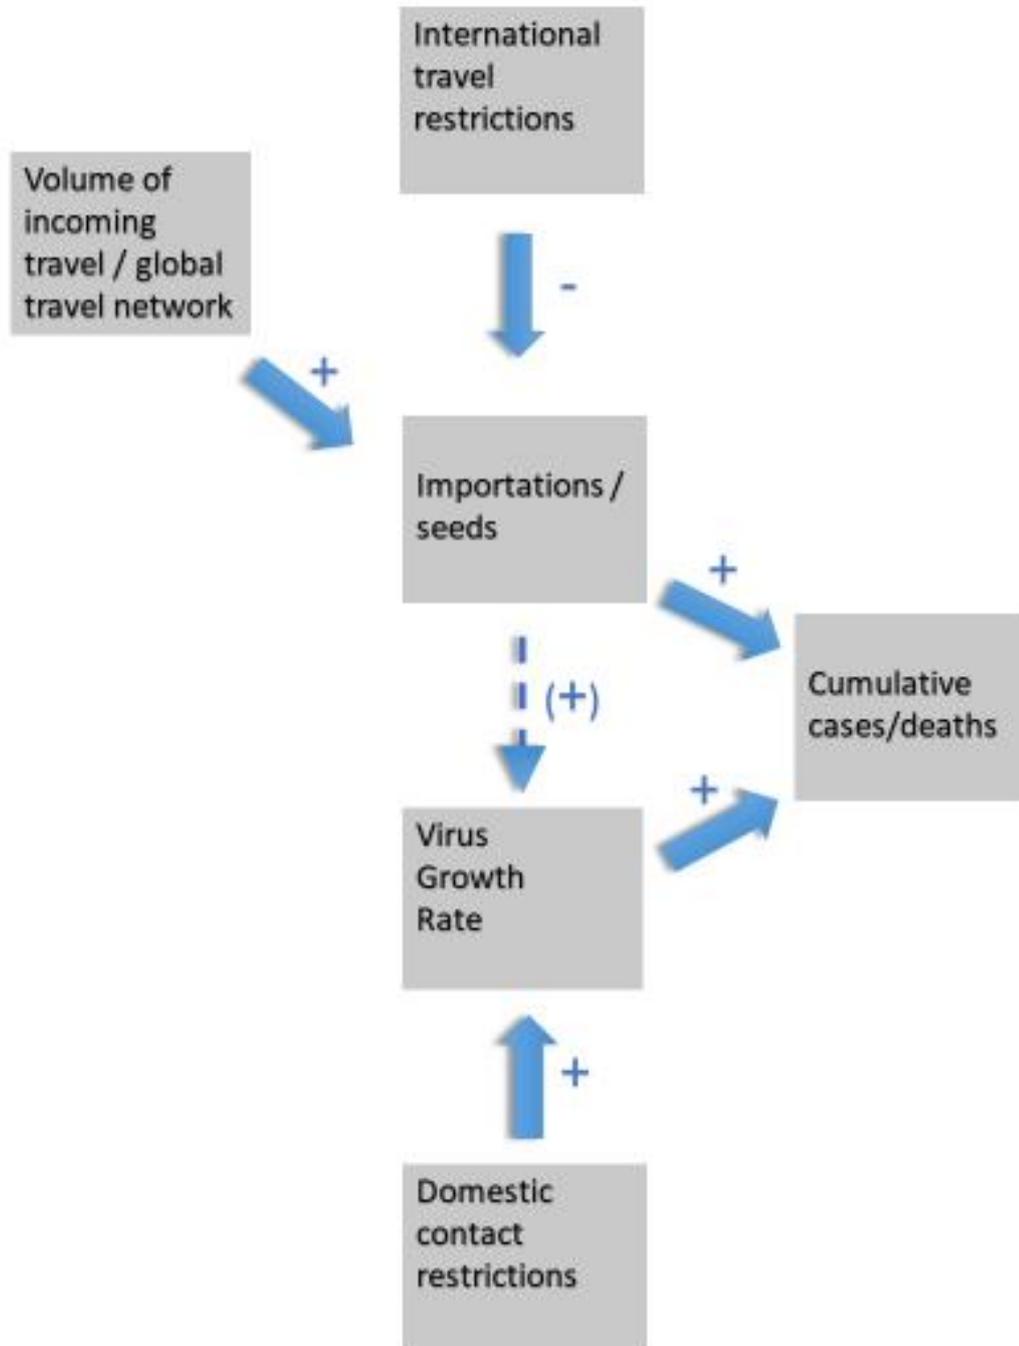


Figure 1 visualizes this theoretical argument. It shows how international travel restrictions and domestic measures affect epidemic outcomes in the form of cumulative case and death counts along two distinctive and largely independent causal paths. International travel in the form of tourism flows and centrality in air traffic networks affects the exposure of a country to international importation of a virus, both in the form of a higher likelihood of early arrival of the first case, and in the form of a higher number of seeds. International travel restrictions counter this process by slowing down, or if implemented broadly and effectively, entirely halting seeding. Only in the early stages of a national epidemic does the number of seeds have a significant direct impact on the growth rate, as long as the number of internationally imported cases is not yet dwarfed by the number of domestic transmissions – this is indicated by the dashed arrow in the figure. Domestic contact reduction and hygiene measures, by contrast, directly target the growth rate and the reproduction factor R , and if effective can succeed in “flattening the curve” and stabilizing the cumulative case or death count at a certain level. Which level that will be, is jointly determined by the two processes: the number of seeds and the timing of travel restrictions on the one hand, and the timing and efficacy of domestic control measures, on the other.

3. Data and variables

3.1. Country sample

To investigate the impact of policies restricting international travel on the course of the COVID-19 pandemic, I coded entry restriction policies and linked them to COVID-19 mortality data for 181 independent states worldwide.²² I focus on independent states and do not consider overseas (e.g., Guadeloupe or Greenland) or special territories (e.g., Hong Kong and Macau) that are not formally independent, partly because information on relevant variables is not systematically available for these entities, and partly because dependencies do not have full sovereignty over their borders. Among independent countries, I exclude a handful of tiny states

²² I thank Jasper Jansen and Aaron Lauterbach for their coding of policies of sub-Saharan African countries. The remainder of the countries were coded by the author.

with less than 50,000 inhabitants, four in Europe (Monaco, Liechtenstein, San Marino and the Vatican State), and three in the Pacific (Nauru, Palau and Tuvalu). The reasons here are also data availability (for the three Pacific states) and limited border sovereignty (the four European ministates are de-facto integrated into France, Switzerland, and Italy, respectively). Also excluded are a few countries (Kosovo, Turkmenistan, North Korea) for which no information on COVID-19 deaths is available in the sources that I used, as well as five countries that are in the midst of civil wars, where the central government controls only part of the territory, and where therefore official pandemic data cannot be reliable and complete (Libya, Syria, Yemen, Afghanistan, Somalia).

3.2. Dependent variables

My analyses aim to explain cross-national variation in the number of deaths as a result of COVID-19. Depending on the modelling strategy (see further below), I use the raw number of deaths or the number of deaths per 100,000 inhabitants. I focus on deaths rather than the number of COVID-19 cases, because the latter are a less reliable indicator of the extent of the pandemic than deaths. The reason is that many COVID-19 infections are asymptomatic or cause only light symptoms that may be similar to a common cold or flu. Such cases will only be detected if extensive testing for COVID-19 is undertaken. Testing efforts and capacities are however far from constant. Over the course of the pandemic, testing has increased, which may result in higher numbers of registered cases even in the absence of a true increase in the number of infections. Second, testing capacities vary widely across countries, and are tendentially lower in poorer and less developed countries and regions. Third, testing may vary over time and across countries as a result of policy decisions. Some countries have implemented systematic testing policies early on, others later. Some countries have implemented random testing, others have limited testing to symptomatic individuals or to particularly vulnerable groups.

Similar problems also affect the number of deaths, but to a lesser extent. The reason is that patients with serious symptoms are more likely to be tested and therefore to be correctly classified as COVID-19 related.

Nevertheless, false positives and false negatives may occur. False positives may occur because deadly cases of COVID-19 tend to be concentrated among patients who already have serious underlying conditions, such as heart problems. The attribution of a death to COVID-19 and not to co-morbidities therefore carries a certain degree of uncertainty. False negatives occur when somebody dies as a result of a COVID-19 infection but is not officially registered as such, for instance when the patient dies before being hospitalized and no post-mortem tests are undertaken. Biases resulting from these problems affect any study of COVID-19 outcomes, especially in a cross-national context where informal and formal registration and testing practices may vary widely. However, to the extent that such biases are not systematically correlated with the explanatory variables of interest – in this case policies to contain the pandemic – they need not affect our substantive conclusions.

There are however two important ways in which systematic correlations that are problematic can come about. The first relates to the quality and capacity of national health systems. Countries with dense and well-functioning health systems and high testing capacity will tendentially register more COVID-19 cases and related deaths. Conversely, one may argue that, at least where deaths are concerned, countries with better health care systems will have a lower death rate from COVID-19 because of more and earlier detection and more effective treatment. The quality of health care in turn is highly correlated with countries' level of wealth, and richer countries tend to be more integrated into the world economy and in international travel networks. Therefore, these countries may be both more exposed to importation of the virus and less inclined to impose travel restrictions. To remove these possible sources of systematic bias, I control in the analyses for per capita gross domestic product (GDP), and as an alternative measure for the Human Development Index (HDI).

Another potential source of systematic bias is political: authoritarian states have greater capacity (and perhaps also greater motivation) than democratic states to conceal high numbers of COVID-19 cases and deaths. Authoritarian states also have greater capacity to impose strict travel restrictions on their populations, although it is not a priori clear whether they are also more motivated to impose such measures. In their attempts to

deny the seriousness of the pandemic, they may also choose the opposite strategy and take fewer measures than more democratic states. Be that as it may, it seems wise to include a control measure of political authoritarianism in the analyses to control for possible biases. To this end, I use the Freedom House civil liberties and political rights index for my main analyses. As an alternative measure, I also use the Polity IV index.

As sources of COVID-19 data, I use the data reported by Johns Hopkins University (JHU) as well as the Worldometer database. These two sources are not independent as JHU partly relies on Worldometer data. While there are some minor deviations between the two sources for particular countries and particular points in time, for a cross-national study such as the current one, it does not practically matter much which source one uses because the two sources are extremely highly correlated and statistically virtually indistinguishable. For the 181 countries included in my analysis, the correlation coefficient of the numbers of COVID-19 deaths in the two datasets on 30 June 2020 is close to unity: $r = .9999$. Because of its easier accessibility and additional provision of per capita measures, I have used the Worldometer data as my primary source but results are identical when using JHU data. I use the number of deaths on 30 June 2020 as my primary reference, but provide additional results for earlier and later dates. The reason to stop the main analysis at the end of June is, that, from the beginning of June onward, increasing numbers of countries started to relax entry restrictions, particularly holiday destinations that rely for their income to a large extent on tourism. The rise in the numbers of cases in countries such as Greece and Croatia during the summer of 2020 suggests that travel relaxations may indeed have led to increases in infection rates in countries that had previously been lightly touched. A systematic analysis of the effects of relaxations of earlier travel restrictions on infection rates is beyond the scope of the current research. Here, I focus on the impact of travel restrictions over the course of the first half of 2020.

Figure 2: Top-30 countries with the highest cumulative COVID-19 mortality rate per 100,000 inhabitants as of 30 June 2020

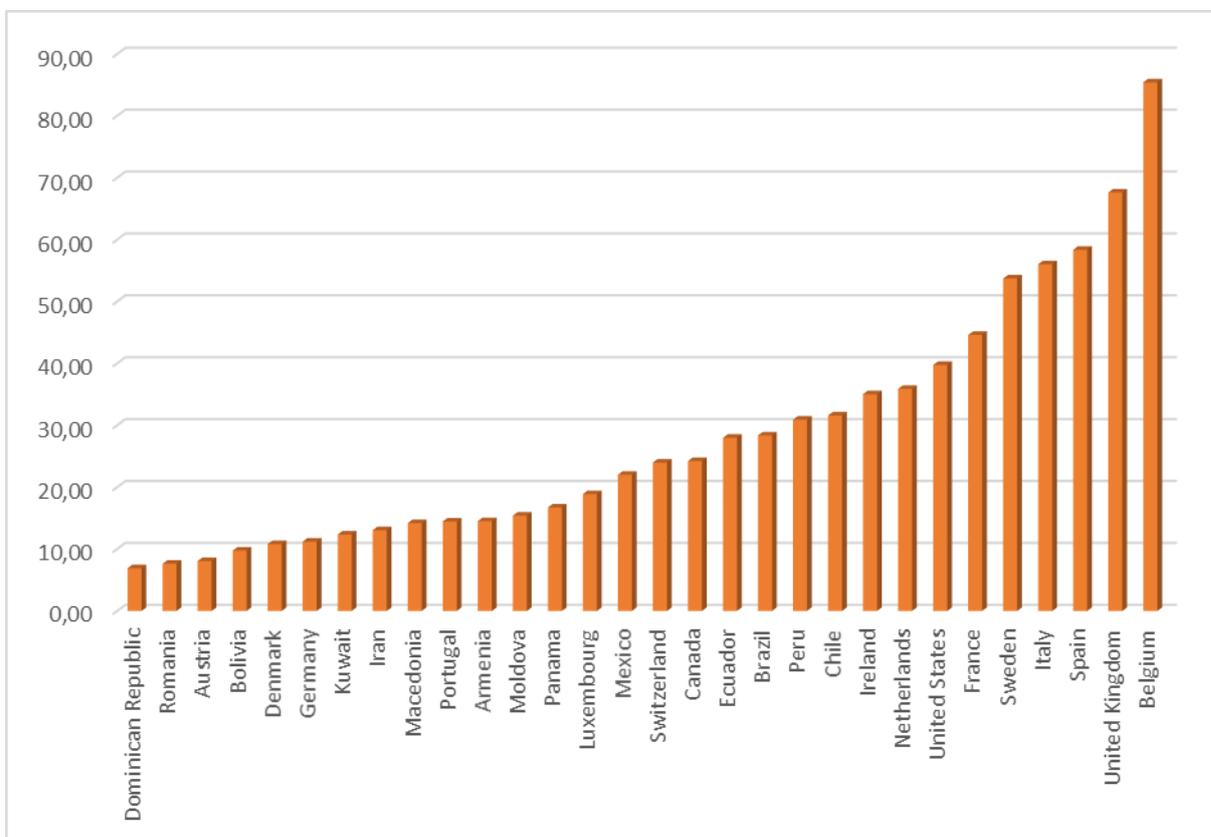


Figure 2 shows the top 30 of countries with the highest cumulative COVID-19 mortality rates as of 30 June 2020. Belgium tops the list with over 80 deaths per 100,000 inhabitants, followed by five other European countries: the United Kingdom, Spain, Italy, Sweden and France. However, high mortality rates were by no means universal among European countries. Slovakia (112th), at the other end of the spectrum, had only 0.5 deaths per 100,000 inhabitants, a similar level as New Zealand and Australia. Even neighbouring European countries sometimes differed starkly. While Sweden had a very high mortality rate, Norway's rate was more than ten times lower (4.7 per 100,000; 49th). Germany (25th with 11.2 deaths per 100,000 inhabitants) is widely considered as a European country that has mastered the pandemic well, but still its mortality rate is significantly higher than that of some of its neighbours, including Austria (8.1 per 100,000; 28th), which directly borders Northern Italy, as well as Poland (3.8; 54th) and Czechia (3.3 per 100,000; 57th).

Because of its large population size, the United States have the highest absolute number of deaths worldwide, but on a per capita basis they only come in seventh place. The top 30 also includes several Latin American countries, such as Chile, Peru, Brazil and Mexico (respectively 9th, 10th, 11th and 15th); three countries in the Middle East (Armenia, Iran and Kuwait; 20th, 23rd and 24th respectively), and one Caribbean country (the Dominican Republic; 30th). The African country with the highest mortality rate is the island state of Sao Tomé and Príncipe (32nd with 6.6 deaths per 100,000) followed by South Africa (50th with 4.7 per 100,000). Strikingly, East Asia, where the pandemic started, has very low mortality rates. The highest-ranked country is the Philippines (87th at 1.2 per 100,000), followed by Japan (98th at 0.8 per 100,000). China itself ranks only in 120th place with 0.3 deaths per 100,000 inhabitants. The countries of Oceania also have very low mortality rates, with New Zealand just before Australia as the highest-ranking countries (respectively in 113th and 114th place, both with 0.5 deaths per 100,000 inhabitants).

3.3. Exposure to international travel

In a first step of my analysis, I will test the hypothesis that the number of casualties as a result of COVID-19 is positively associated with the exposure of a country to movements of people across its international borders. I use two alternative indicators of international travel flows. The first is the annual number of tourists received by a country (in millions) based on data of the United Nation's World Tourism Organization (UNWTO).²³ These data refer to the total number of foreign arrivals in the year 2017 (or the most recent year available; see UNWTO 2018). That year, France was the country with the largest number of tourist arrivals (almost 87 million), followed by Spain, the United States, China, and Italy. Data are available for 176 of the 181 countries.²⁴ The UNWTO also provides data on outbound tourism, but for a smaller sample of independent countries (n=109). Inbound and

²³ See <https://www.e-unwto.org/doi/pdf/10.18111/9789284421152>.

²⁴ The UNWTO also provides data on outbound tourism, but for a smaller sample of independent countries (109). Inbound and outbound tourism are quite strongly correlated ($r=.65$, $p<.001$) and results are similar when outbound tourism is included instead of tourist arrivals.

outbound tourism are quite strongly correlated ($r=.65$, $p<.001$) and results are similar when we use the other measure (see below).

The advantage of the tourist data is that they encompass both overland and air travel, but they may not adequately reflect business and commuter travel, and also do not account for passengers who pass through airports in transit to another country. Therefore, as an alternative measure, I use the so-called Airport Expected Force of Infection (AEF) index, which has been calculated by Lawyer (2015, 2016) from 2014 World Airline Network (WAN) data and International Air Transport Association (IATA) data on the number of available seats on these routes. The WAN database includes 3,458 airports worldwide that were in 2014 connected by 68,820 routes. Based on these travel network data, Lawyer calculated an index with values from 0 to 100 indicating the time it takes for a virus to achieve pandemic status if it originates in the location of a particular airport. Some airports are connected to many other airports and receive large numbers of international passengers, either as a final destination or in transit. Other airports receive few passengers, are linked to few other airports and are consequently less powerful transmission belts for the spread of a virus. If we now switch the perspective from the likelihood of spreading the virus to the likelihood of receiving it, we can assume that airports that receive many passengers and are linked to many other airports are also more likely to be on the receiving end of a global pandemic. This implies that they are both more likely to receive virus infections earlier, and to receive more virus seeds overall. The index shows Frankfurt Airport, Amsterdam Schiphol, Paris Charles de Gaulle, Istanbul's Atatürk Airport and Atlanta Airport as the airports most likely to be on the receiving side of a virus infection that has its source in a random place on earth. In cases where the data include several airports for one country, I take the airport with the highest AEF score (e.g., Frankfurt for Germany or Atlanta for the USA) as the country score. AEF data are available for 180 of 181 countries (all countries but Andorra, which has no international airport).²⁵

Finally, I look at an important natural feature of international borders, namely whether a country has any international land borders at all, or

²⁵ I thank Aaron Lauterbach for the coding of these data.

whether alternatively it is an island or group of islands. Even though air travel and tourist flows cover important dimensions of exposure to the importation of infections from abroad, it may still matter whether a country has any land borders at all. Land borders are difficult to monitor, especially in more remote areas, and quite a bit of undocumented cross-border movement may therefore occur. Several legal forms of cross-border movement, such as short-distance commuting for work, leisure or visiting family and friends are moreover not covered by tourism data, which generally refer only to foreigners who stay overnight. I therefore assume that island states have lower exposure to international human movement, even when tourism and airport network centrality are accounted for. I define islands as countries that have no land borders at all and are surrounded by sea on all sides. Among the 181 countries studied, 36 are island states, including, e.g. Australia, New Zealand, Japan, the Philippines, Taiwan, and Sri Lanka, as well as a host of Atlantic, Caribbean, Pacific, and Indian Ocean island states.²⁶

3.4. Policies restricting international travel

The main explanatory variable in my analysis consists of policies that aimed at containing the spread of COVID-19 through restrictions on international travel. I distinguish two forms: entry bans and mandatory quarantines. Entry bans prohibit travellers who arrive from particular countries to enter. Usually, this took the form of rules stipulating that anyone who had been in the designated countries in the last 14 days would be refused entry. Sometimes, as in the case of Morocco²⁷, these bans also precluded citizens and permanent residents to return to their home country, but in most cases, exceptions were made for these categories, which either had to find ways to return on their own or were repatriated

²⁶ I also considered a less strict definition of what constitutes an “island state”, which additionally includes countries that have land borders, but are still located on islands. This definition additionally includes the United Kingdom and Ireland; Cyprus; Haiti and the Dominican Republic, and Indonesia, East Timor, Brunei and Papua New Guinea. Because the strict definition of island state performed better and is also theoretically preferable, I used that definition for the analysis.

²⁷ <https://www.volkskrant.nl/nieuws-achtergrond/geen-marokkaan-komt-eigen-land-nog-binnen-wereldwijd-27-850-marokkanen-gestrand~bbe87ee7/?referrer=https%3A%2F%2Fwww.google.com%2F>.

with special flights organized by travel companies or home governments. For the coding of an entry ban, it was not necessary that the ban also included citizens and/or permanent resident foreigners; the measure was coded if it targeted non-resident foreigners.²⁸ I coded three types of entry bans: those pertaining to China, to Italy, and to all countries of the world. These were the most frequent types of entry bans, although there were of course also targeted entry bans that applied to other source countries (e.g., Iran, South Korea or the United States), or to subsets of the countries of the world (e.g., the EU's entry ban for all non-EU and non-Schengen countries except the United Kingdom of 17 March). If such a ban implied a not yet existing ban on China or Italy, it was coded as such. For example, for all EU countries that did not yet have an entry ban on China in place, the 17 March ban on non-EU countries was coded as a China ban. Similarly, the 13 March USA entry ban on travellers from Schengen countries was coded as an Italy ban. Entry bans for all countries of the world were simultaneously coded as China and Italy entry bans. Coding all possible types of entry ban constellations would have been impossible and would have yielded little extra explanatory power as bans on other sets of countries often coincided with or could be coded as bans on China and/or Italy.

The second type of travel restrictions that I considered are mandatory quarantines. In most cases, they pertained to anyone who had recently (usually specified as the last 14 days) been in particular source countries. Unlike entry bans, quarantines almost always included everyone who had been in a source country, regardless of whether they held citizenship or permanent resident status of the target country. Only quarantine policies that also applied to citizens and permanent residents were coded. To be coded, quarantines had to be mandatory; mere advice to self-isolate for 14 days was not sufficient. Within the category of mandatory quarantines,

²⁸ The mere suspension of flights to/from a source country by a single national airline (e.g. Lufthansa flights to China from the end of January onward) does not count as an entry ban, as long as entry of passengers from the source country who fly with other airlines is not banned. Only if the national airline is the only direct connection and there is no land border, this is counted as a de-facto entry ban. Also not counted as entry bans are entry rules that apply only to parts of the banned source country (e.g., Wuhan/Hubei; Northern Italy). Finally, suspension of issue of "visa on arrival" alone is also not counted as an entry ban, because it continues to permit entry of travellers who already have a visa or apply for it in advance.

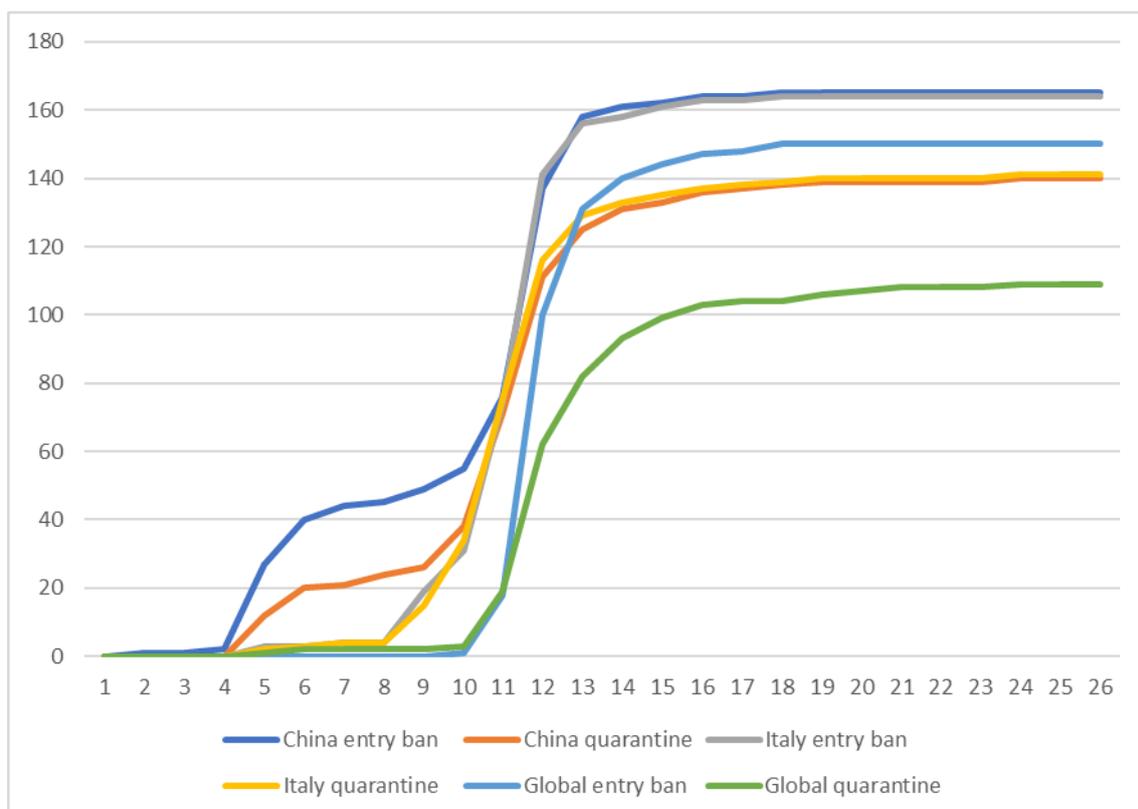
there was however wide variation. Some countries, such as several Pacific island states, required returnees to spend these 14 days in a corona-free foreign country and to subsequently obtain a doctor's certificate of a negative COVID-19 test. Others let returnees into the country but required them either to spend the 14 days in a quarantine facility set up by the state (e.g., Turkey and several Caribbean states) or to spend these 14 days in mandatory self-isolation at home. The degree of enforcement of self-isolation varied as well. Some countries (e.g., Israel) monitored self-isolation and penalized infringements with heavy fines, whereas at the other extreme, the Belgian quarantine regime from 25 March onwards consisted of handing out information leaflets to returnees at the airport with no apparent follow-up enforcement measures. While such variations are likely to have been consequential, the data sources provided for many countries too little detail for this to be coded in a systematic fashion. A mandatory quarantine can therefore mean anything from the loose Belgian to the very strict Pacific islands regimes. Like entry bans, mandatory quarantines were coded when they applied to China, to Italy, and to all countries of the world.

For each of the resulting six travel restriction policies, the exact begin data on which the policy went into force in a particular country was coded. Since my focus is on the effects of restrictions and not on those of subsequent easing or lifting of restrictions, I have not coded end dates of travel restriction policies. However, a condition for coding a travel restriction was that it had to have remained in force for at least two weeks. The resulting dataset can be downloaded as an online appendix with this paper, complete with full references to the sources used. These sources included any reliable material available online, including newspaper and press agency (especially Reuters) articles, government and embassy websites, as well as a number of sources that have compiled (each in non-exhaustive and not always error-free ways) COVID-19 related travel restrictions.²⁹ Whenever possible, information was cross-verified from several of these sources.

²⁹ These include, e.g., <https://restrictions.info/>;
<https://www.garda.com/crisis24/coronavirus-updates>;
https://en.wikipedia.org/wiki/Travel_restrictions_related_to_the_COVID-19_pandemic;

Figure 3 shows, across the 26 weeks of the first half of 2020, how many countries had introduced one of these six types of travel restrictions. The first type of travel restrictions to be introduced, starting in January, were entry bans for travellers from China, followed somewhat later by quarantine measures for travellers arriving from that country. Entry bans and quarantines on Italy lagged about a month behind those on China and took off seriously from the end of February onward (week 9 ending on 1 March). Global travel restrictions were introduced later still, from the second week of March onward (week 11). By the end of April (week 17/18) the diffusion of travel restrictions had reached a plateau phase and only very few countries still introduce new restrictions. The figure shows that although travel restrictions on Italy took off later than those on China, both ultimately became equally popular. Regardless of their scope, entry bans were more popular than quarantines.

Figure 3: Number of countries that had introduced various types of travel restrictions by week, January – June 2020



<https://www.thinkglobalhealth.org/article/travel-restrictions-china-due-COVID-19;>
<https://www.nytimes.com/article/coronavirus-travel-restrictions.html>

For the regression analyses, the dates of the six types of travel restrictions were transformed into linear variables by counting the number of days from their introduction until the end of the observation period, 1 July 2020. For instance, a travel restriction introduced on 15 March would get a value of 107, meaning that it was introduced 107 days before 1 July 2020. To prevent that cases where a particular measure had not been introduced until 30 June would become missing values, the date of the measure was set to 1 July 2020, and the linear variable accordingly to zero. Coded this way, the variables reflect the “dosage” of the treatment “travel restrictions”: when a specific restriction was never introduced, its dosage was zero, whereas if it had been introduced 60 days before the end of the observation period, the dosage was 60, and so on. In addition to the six separate variables, I created a composite scale averaging the timing of all six measures for each country. To assess the relative efficacy of entry bans and mandatory quarantines, I created scales only of entry ban and of quarantine measures (both based on 3 items). I also created separate scales summarizing targeted travel restrictions (4 items: entry bans and quarantines targeting China and Italy) and global restrictions on all foreign countries (2 items).

When using the date of travel restrictions to predict the course of local epidemics in a cross-national perspective, the problem arises that timing relative to the spread of the pandemic matters. For instance, if a country that is geographically close and has intensive exposure to international travel (say, Japan) introduces an entry ban for travellers from China on a particular date (say 1 March) this is not equivalent to a geographically distant country with low exposure to international travel (say, Ghana) imposing such a ban on exactly the same date. By 1 March, countries close to China and highly exposed to international travel are likely to have long been on the receiving end of several infection seeds from China, which will moreover have already had time to spread further by way of domestic transmission. An entry ban on 1 March will therefore likely be late relative to the local course of the pandemic. At the same time, a geographically distant country with low exposure to international travel is likely to have become by the same date the target of very few infection seeds or even none at all, and the entry ban may therefore be early relative to the course of the pandemic in this country. Therefore, the calendar date of a travel

restriction as such is an imperfect indicator of its potential effectiveness. In causal terms, the same dosage of the treatment “travel restrictions” will not have the same effect on a country that has not yet received (m)any infections as on a country that is already in the middle of a full-blown domestic epidemic. To deal with this issue, I will additionally use relative measures of the timing of travel restrictions (more detail below).

3.5. Control variables³⁰

GDP per capita: World Bank data on nominal GDP for the year 2019. For Eritrea, South Sudan, and Taiwan, which are not included in the World Bank data, I use 2019 IMF data instead. I use nominal GDP rather than GDP adjusted for local purchasing power because what matters here is less individual welfare than the collective capacity of states to handle a global pandemic, which depends on purchasing power in the world market for medical equipment, tests, vaccines and drugs.

Human Development Index: I use the 2019 scores of this index provided by of the United Nations Development Program (UNDP). This composite index encompasses measures of life expectancy, education, and gross domestic product.

Democracy: Freedom House Political Rights and Civil Liberties score for the year 2020, measured on a scale from 0 to 100, with 100 indicating the highest level of democracy and human rights.³¹ Alternatively, I use the Polity dataset for the year 2015, which codes each country on a scale ranging between -10 (full autocracy) to +10 (full democracy).³²

Median age: to control for the fact that countries with an older population may have more corona deaths, I use the median age of the population on the basis of the CIA World Factbook. As an alternative, I use World Bank data

³⁰ Except for the control variables discussed in this section, I considered population density, as well as indicators of climate (rainfall and average July and January temperatures). Because these never turned out significant in preliminary analyses, I have not included them in the analyses reported in the paper.

³¹ See <https://freedomhouse.org/countries/freedom-world/scores>

³² See <https://www.systemicpeace.org/inscrdata.html>

on the percentage of the population of 65 years and older; the latter data are not available for six of the 181 countries studied.

World region: All models use robust standard errors clustered by world region, to account both for unobserved regional heterogeneity, and for regional interdependence of infection risks.³³ From the temporal spread of the pandemic, it is clear that regional variation is important: Asian and Pacific countries with close proximity and ties to China were affected earliest, followed by Middle Eastern and European countries, then the Americas, and last sub-Saharan Africa. I cluster standard errors by the following seven world regions: Asia (n=27), Pacific (n=11); Middle East and North Africa (n=20; including Iran, Turkey and the three Caucasus countries); Europe (n=41; including Russia); North America and the Caribbean (n=13), South and Central America (including Mexico; n=22); and finally, sub-Sahara Africa (n=47). In the Online Appendix I show that results using region fixed effects, which restrict the analysis to within-regional variation, are similar.

Early epidemic severity: For the purpose of several robustness checks below, I use measures of how early and how severe countries were hit in the early phases of the pandemic. Such controls are useful for a number of reasons. First, there may be reverse causality when travel restrictions are introduced as a reaction to early epidemic development in a country rather than pre-emptively before a country has been significantly hit. Second, countries that were hit early have had more time to accumulate deaths until 30 June, and this, rather than early or late introduction of travel restrictions, may explain why some countries have higher cumulative mortality. To address these concerns, I use the timing of the first and tenth COVID-19 cases and deaths, as well as the death toll on 13 March as control variables. Adding such controls is a conservative test of the effect of travel

³³ The alternative strategy of including fixed effects for world region as a predictor variable would risk controlling away regional commonalities regarding the variables of interest, e.g. that countries within a region are on average richer or more exposed to international travel. I thank Macartan Humphreys for suggesting this way of dealing with regional variation to me. In the Online Appendix, I do however provide results using world region fixed effects. These are similar to those reported in the paper, although effect sizes for travel restrictions are slightly smaller because of the limitation to within-region variations.

restrictions because early epidemic severity may itself be a result of the presence or absence of early travel restrictions (see further below).

Domestic COVID-19 containment measures: Travel restrictions might be correlated with other policy measures that states have taken to contain the COVID-19 pandemic, and this may bias estimates of the effects of the former. To control for this possibility, I use policy data coded in the context of the “Coronavirus Government Response Tracker” developed by the Blavatnik School of Government at Oxford University. From that database, the date of introduction of the following mandatory containment policies was coded (I ignored policies that were merely non-mandatory recommendations), both on the national (“general”) and the local (“targeted”) levels: bans on public events; stay-at-home requirements (“lockdowns”); school closures; workplace closures; public transport closures; and domestic travel restrictions. Introduction dates were transformed into linear variables by counting the number of days between the introduction of the measure and 30 June 2020. These items were combined into a strong “domestic restrictions” scale with $\alpha = .87$.

3.6. Modelling strategy

I use two different ways to model effects on COVID-19 mortality:

- 1) A negative binomial count regression of the total number of deaths as of 30 June 2020 with population as the exposure variable;
- 2) A linear regression of the number of deaths per 100,000 inhabitants as of 30 June 2020.

Each of these strategies comes with its own advantages and disadvantages. The count regression is the statistically most appropriate strategy because deaths are not normally distributed but characterized by relatively few countries with very high mortality and a long tail of cases with much lower frequencies. However, similar to standard regression with a logged dependent variable, this modelling strategy weighs the difference between 1 and 100 deaths equal to the one between 100 and 10000 deaths, even in countries with the same population size. Population (logged) is not included in the negative binomial model as a predictor, but as the exposure variable.

Coefficients of predictor variables thus indicate the degree to which their influence causes the number of deaths to deviate from what one would expect merely on the basis of population size.

The linear regression of deaths per 100,000 inhabitants by contrast has the advantage that it is a direct and easily interpretable measure of the severity of the local epidemic that can be directly descriptively compared across countries of different population sizes. Whereas the count regression gives greater weight to (relatively small) country differences at the lower end of the pandemic severity scale, the linear regression gives greater weight to the (relatively large) country differences at the higher end of pandemic severity. The disadvantage of this strategy is, however, that it does not adequately capture the statistical distribution of the dependent variable and that standard errors may therefore be wrongly estimated (although I use robust standard errors to reduce this problem). Moreover, it may lead to predictions of negative death rates, which are logically impossible. By combining the two modelling strategies and focusing on results that are robust across them, I aim to combine the advantages of both. Statistically, we should however put our greatest trust in the results of the negative binomial regressions.

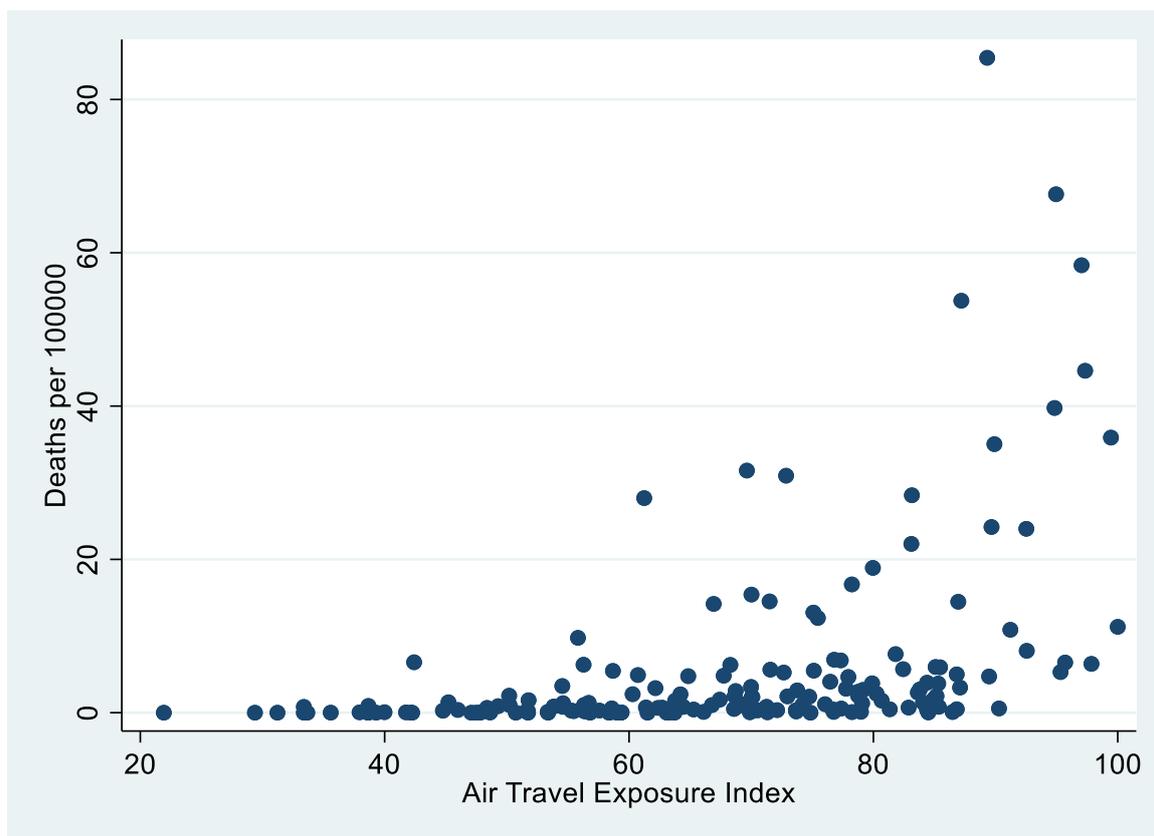
To allow comparison across models, all regressions exclude countries that have a missing value on any of the predictor variables. This affects four countries: Andorra, for which no AEF airport data are available, and South Sudan, Liberia, and Equatorial Guinea, for which no data on numbers of tourist arrivals are available. China and Italy are also excluded because they were the main sources of the pandemic and the main targets of the travel restrictions that we will subsequently analyse (differently put: Italy and China have, by definition, missing values on travel restrictions on themselves). This leaves 175 countries for the analysis. I will however discuss whether the Chinese and Italian cases fit the general pattern in a separate section below.

4. Main results

4.1. Exposure to international travel

Figure 4 gives a visual impression of the bivariate relationship ($r = .44$; $p < .001$) between air travel exposure and mortality per 100,000 inhabitants. While in the middle part of the distribution of air travel exposure there is quite a bit of variation in mortality outcomes, all countries with very low travel exposure (below 40) also have very low COVID-19 mortality, whereas at the other end of the spectrum, most countries with very high travel exposure (above 80) also have elevated mortality. A similar bivariate relationship exists for tourist arrivals ($r = .53$; $p < .001$).

Figure 4: Bivariate relationship between air travel exposure (AEF index) and COVID-19 mortality per 100,000 inhabitants



Do these results hold when we control for other country characteristics? As a first step, we look at a baseline model in which only per capita GDP, the median age of the population, and the degree of democracy are included.

Moreover, as explained above, the regional location of countries is taken into account by correcting for regional clustering of standard errors. Table 1 first shows the results of a baseline model for the negative binomial (column 1) and linear regressions (column 3), respectively. Both regressions show that more affluent countries have more deaths (1.9 percent more deaths per additional 1,000 USD per capita income according to the negative binomial regression), as have countries with older populations (8.3 percent more deaths for a one year higher median age). Democratic countries also tend to have more deaths, but these coefficients are not statistically significant.

Table 1: Regressions of COVID-19 mortality on indicators of exposure to international travel and control variables

	Negative binomial: baseline model (IRR coefficients)	Negative binomial: adding exposure to international travel (IRR coefficients)	Linear: baseline model (B coefficients)	Linear: adding exposure to international travel (B coefficients)
Dependent variable	Cumulative number of deaths on 30 June 2020	Cumulative number of deaths on 30 June 2020	Cumulative deaths per 100,000 inhabitants on 30 June 2020	Cumulative deaths per 100,000 inhabitants on 30 June 2020
Per capita GDP in 1000 USD	1.019**	1.008	.221**	.168**
Median age	1.083*	1.015	.166+	-.092
Democracy and human rights (0-100)	1.004	1.010	.039	.088+
Air travel exposure (AEF index 0-100)	-	1.042***	-	.024
Tourist arrivals in millions	-	1.004	-	.352**
Island state	-	.387	-	-6.184 ***
<i>Clustered by world region</i>	<i>yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
N=	175	175	175	175
R ²	.027	.040	.229	.406

Most of the effects of GDP and age disappear, however, when we add in columns 2 and 4 the three indicators of exposure to international travel.

The main reason why richer and older countries have more COVID-19 deaths, it seems, is that these are usually also countries with high exposure to international travel. Median age has no significant impact at all anymore, whereas the effect of per capita GDP remains significant only in the linear regression (where it indicates a modest effect of .168 more deaths per 100,000 inhabitants for an additional 1,000 USD per capita income). In the linear regression, we also find a marginally significant effect of a country's political regime, which may indicate that democratic countries are somewhat more willing to report the true extent of COVID-19 casualties.

The most important finding of the regressions in columns 2 and 4 is, however, that we find consistent and strong support for the importance of exposure to international travel. Because the three measures are fairly strongly correlated (particularly AEF and tourist arrivals: $r=.57$), which indicators become significant differs between the negative binomial and linear regressions. In the negative binomial model, we find a strong effect of centrality in the international air travel network. For every point increase in the 0-100 AEF index, the number of deaths is 4.3 percent higher. To give an impression of what this might mean in practice, Spain (AEF=97) and Austria (AEF=87) differ 10 points on the AEF index, and the model results predict that therefore the number of COVID-19 deaths will be 52 percent (1.043 to the 10^{th} power) higher in Spain than in Austria. Most countries have lower AEF's still than Austria, which has the 19^{th} ranking AEF among the 175 countries in the analysis.

In the linear regression, by contrast, the number of annual tourists and whether a country is an island state have strong and highly significant effects. Other things being equal, island states have on average six fewer deaths per 100,000 inhabitants, which is a sizeable effect given that six is the average death rate per 100,000 inhabitants across the 175 countries analysed. Every million additional annual tourists is associated with .35 more deaths per 100,000 inhabitants. As an illustration, we may take France, the country with the worldwide highest tourist inflow (87 million), and Germany (37 million). The model findings predict that this difference in tourist numbers accounts for 17.4 more deaths per 100,000 inhabitants in France, which amounts to half the actual difference in per capita COVID-19

deaths in the two countries as of 30 June 2020 (45 per 100,000 in France and 11 in Germany). Here, too, we compare two countries in the higher range of travel exposure: Germany has the 7th highest number of tourists in the world and 38 percent of the countries analysed receive less than one million tourists yearly.

Results using outbound tourism for the 109 countries for which these data are available are similar. When both inbound and outbound tourism are included, the former has the strongest effect (.27, $p < .01$), but outbound tourism also has a modest independent effect (.07, $p < .10$). These are of course averages. For several European countries, returning outbound tourists rather than incoming foreign tourists were the most important source of early contagions, particularly related to people returning from ski holidays in Italy and Austria.

4.2. International travel restrictions

To what extent were countries able to counter their vulnerability to the pandemic by introducing travel restrictions? Figure 5 shows the bivariate relationship between the summary scale of travel restrictions and the COVID-19 mortality rate. As explained above, higher values on the travel restrictions scale indicate earlier introduction dates. The bivariate correlation between the two variables is moderately strong ($r = -.40$; $p < .001$). It is suppressed by Belgium, which is the outlier in the upper right part of the figure (excluding Belgium the correlation is $-.47$; $p < .001$). One possible reason why Belgium is an outlier has been discussed above: officially, the country had a mandatory quarantine policy from 25 March onward, but enforcement seems to have been lacking. Another reason, which will be discussed below, is that the absolute timing of travel restrictions displayed here does not take into account when measures were introduced relative to the phase of the local epidemic. From the latter perspective, Belgium was a latecomer. When it introduced entry restrictions on Italy on 20 March, the country had already registered 37 COVID-19 deaths, and by the time the quarantine policy was introduced that number had increased tenfold to 372.

Figure 5: Bivariate relationship between travel restrictions (summary scale) and COVID-19 mortality per 100,000 inhabitants

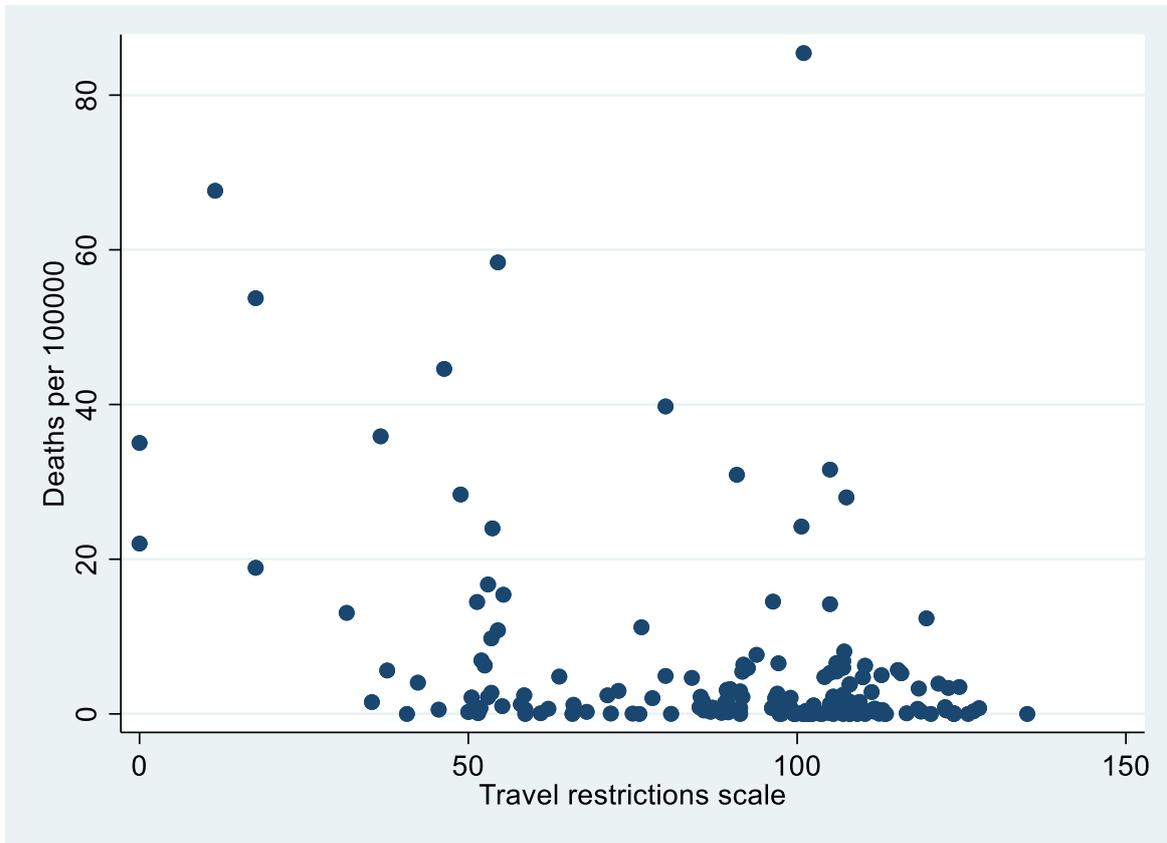


Table 2 shows to what extent the relationship between travel restrictions and COVID-19 mortality holds in a multivariate context. Columns 1 and 4 show this for the summary scale of travel restrictions. While controlling for all the variables included in Table 1, the coefficients for the travel restrictions scale are highly significant and imply that introducing these measures one day earlier, reduces the number of deaths by 0.8 percent (in the negative binomial regression) or by .10 per 100,000 inhabitants in the linear regression. The effects of GDP, democracy, and exposure to international travel that we found in Table 1 remain, but the effect sizes, particularly of GDP, become smaller, indicating that the higher death rates of richer countries are partly explained by the fact that they tended to introduce travel restrictions later, possibly as a result of considerations about negative effects on business and trade. Indeed, richer countries were tendentially later in imposing travel restrictions than poorer countries ($r=-.21$; $p<.01$).

Table 2: Regressions of COVID-19 mortality on timing of travel restrictions, exposure to international travel and control variables

	Negative binomial: Cumulative number of deaths on 30 June 2020 (IRR coefficients)			Linear: Cumulative deaths per 100,000 inhabitants on 30 June 2020 (B coefficients)		
Per capita GDP in 1000 USD	1.008	1.009	1.008	.152***	.153***	.147**
Median age	1.021	1.026	1.015	-.065	-.063	-.074
Democracy and human rights (0-100)	1.009	1.008	1.009	.075+	.075+	.077+
Air travel exposure (AEF index 0-100)	1.037***	1.036**	1.041***	.005	.004	.011
Tourist arrivals in millions	1.002	1.002	1.003	.333**	.334**	.339**
Island state	.434	.449	.444	-4.7734**	-4.769**	-4.498**
Travel restrictions scale (6 items; higher score means earlier introduction)	.992***	-	-	-.099*	-	-
Entry bans scale (3 items)	-	1.000	-	-	-.048	-
Quarantines scale (3 items)	-	.994***	-	-	-.050*	-
Targeted restrictions scale (4 items)	-	-	.990***	-	-	-.091
General restrictions scale (2 items)	-	-	1.003	-	-	.006
<i>Clustered by world region</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
N=	175	175	175	175	175	175
R ²	.041	.042	.042	.448	.448	.451

Columns 2 and 5 inquire which of the two basic types of entry restrictions were more effective: entry bans or quarantines. The timing of the two is only moderately correlated ($r=.30$; $p<.01$), indicating that some countries used these measures as alternatives, or introduced them at different points in time. Results, particularly those of the negative binomial regression, suggest that quarantines were the more effective of the two. The most

plausible reason is that quarantines almost always applied to all travellers who had been in a certain source country in the previous 14 days, regardless of their nationality or residence status. Entry bans, by contrast, in most cases did not apply to citizens and permanent residents returning to their home country. In the absence of a quarantine requirement, if infected, such returnees could therefore spread the virus.

A further question of practical policy interest is whether targeted measures against specific high-risk countries or general measures targeting all foreign countries were more effective. Of course, to the extent that the latter also include high-risk countries, they may be effective because more targeted bans are effective. In columns 3 and 6, we however include targeted and non-targeted measures in the same model, and therefore test whether global travel restrictions had any additional benefit beyond what would have been achieved by targeted travel restrictions on China and Italy. The answer is that they had no additional benefit whatsoever: both the negative binomial and the linear regression show that the coefficients of general restrictions are very small and not even in the right direction, once we control for targeted travel restrictions.

Finally, we can look separately at effects of travel restrictions on China and Italy. To this end, I combine entry bans and quarantines and take the date of whichever of these measures came first (because measures on China and Italy are too highly correlated – $r=-.68$; $p<.001$ – we cannot include them simultaneously in one model). Because the results for other variables hardly change, these results are not included in Table 2 and I only report the coefficients for the travel restrictions here. In the negative binomial model, the coefficient for restrictions on China is .990 and that for restrictions on Italy .989 (both highly significant at $p<.001$). This implies that each day that such restrictions were implemented earlier is associated with 1.0 percent, respectively 1.1 percent fewer deaths. In the linear regression, the difference between the estimated effectiveness of restrictions on China and Italy is larger: the coefficient is -.10 for restrictions on China against -.15 for Italy (because of the relatively large standard errors in the linear regression both effects are significant only at $p<.10$).

To give an idea of what these results mean in practice, we may look at the timing with which different countries introduced travel restrictions on Italy. Among the countries in the European and Mediterranean area, Israel was the first to introduce restrictions on travellers from Italy. On 26 February, when the Italian death toll was still as low as twelve, and Israel itself had registered only two cases and not a single death, the country introduced a mandatory quarantine for travellers entering from Italy, including also Israeli citizens and permanent residents. On 9 March, Israel followed up with a total entry ban for Italy. Turkey also reacted relatively early: on 28 February, it imposed an entry ban on foreigners arriving from Italy, followed by a mandatory quarantine for all travellers arriving from Italy from 6 March onward. The Czech Republic was the first European Union country to impose restrictions on travel from Italy in the form of a mandatory quarantine on 7 March. Italy's neighbour Austria was slower to react: it closed its borders to travellers from Italy on 10 March and imposed a quarantine for people returning from Italy on 17 March. Germany was later still, with a ban on flights from Italy from 18 March onwards, and a quarantine for returnees from high-risk countries including Italy that only went into force on 10 April. Some other European countries were much later still: France only closed its borders to travellers from Italy on 7 April, the United Kingdom did not impose a quarantine until 6 June, and Luxemburg, Ireland and Sweden never had any entry ban or mandatory quarantine affecting Italy at all (though because of other countries' measures, their traveller inflow was of course also very strongly reduced).

When one considers such strong policy differences across countries, an effect of 1.1 percent fewer deaths per day restrictions are introduced earlier, which may perhaps seem small at first, becomes much more meaningful. If we take the difference between Israel and Germany, for instance, the 22-day difference in the timing of the first travel restrictions on Italy implies a prediction of 22 percent (0.989 to the 22th power) fewer deaths in Israel (according to the negative binomial regression) or 3.3 fewer deaths per 100,000 inhabitants according to the linear regression. This would account for a considerable part of the actual difference between the two countries: 3.9 deaths per 100,000 in Israel and 10.2 in Germany. Needless to say, countries that were much later than Germany, such as France, the United Kingdom and Sweden have had to pay, according to these

results, a much higher price still for their reluctance to reduce travel from high-risk areas.

As I have argued earlier, while we find consistent evidence for effects of considerable size of the date at which travel restrictions were introduced, these estimates may still underestimate the true impact of travel restrictions. This is so, because, depending on the local trajectory of the pandemic, the same absolute date may be late or early in a relative sense. Some examples can illustrate this. Austria and Spain both introduced travel restrictions on Italy on 10 March. However, by that date, 36 people had already died of COVID-19 in Spain, whereas Austria had not registered a single death. Finland introduced travel restrictions on Italy on 19 March, only one day before Belgium. But, in timing relative to the epidemic, the difference was large: Finland introduced the measure before it had registered any deaths, Belgium had already registered 37 deaths by 20 March.

Recalling the theoretical considerations on international seeds and domestic contagion in the theory section above, we must assume that after domestic diffusion of the virus has passed a certain threshold, the influence of international importations, and consequently also of international travel restrictions that aim to contain such imports, will rapidly decline. Where the demarcation between “too late” and “still early enough” lies, is ultimately an empirical question, the answer to which may vary to some extent depending on case-specific circumstances (e.g., to what extent there are one or multiple epicentres of domestic spread). Here, however, we look for patterns that hold across the global population of countries. To this end, I experimented with several cut-off points to distinguish early from late interventions: before and after the first case, the tenth case, the first death, and the tenth death, respectively. While results go in a similar direction for other cut-off points, the demarcation before or after the 10th death gives the strongest results. Given what we know about COVID-19 fatality rates, which are estimated to be around 0.5 percent of those infected, 10 deaths would correspond to approximately 2,000 cumulative cases. Because deaths generally occur a few weeks after infection, the number of cases would moreover have doubled several times by the time ten cumulative deaths are reached. Germany can serve as an

example: the threshold of ten cumulative deaths was reached on 15 March, but by that day, 5,800 cases had already been registered. With such high domestic case numbers, of which many at this early stage of the pandemic are still active cases, it is not hard to imagine that the point has been reached beyond which travel restrictions will no longer have a significant impact on the diffusion of the virus.

Using the day on which the tenth death was registered as a cut-off point, I coded for all four targeted measures – entry bans and quarantines targeting China and Italy – whether they were implemented before or after. Measures implemented on the same day as the tenth death were coded as after. Table 3 shows how many countries, along with some examples, implemented all measures before; one, two or three measures after; and finally all four measures after the tenth death. One could argue that countries that introduced their travel restrictions after the tenth death were just unlucky in being hit early on by the pandemic, without having sufficient time to react adequately, whereas other countries had the luck of being prewarned by what was happening elsewhere and were thus able to react in time. The last column of Table 3 shows that this is as a general rule not what happened. Countries that implemented all four measures before the tenth death on average took these measures on 11 March, weeks or months before the average of countries that took all or some of these measures after the tenth death (or never implemented them at all). Moreover, countries that reacted with travel restrictions before they were engulfed by the pandemic included countries that were hit early on, such as Australia, or that were very close to the epicentres of the pandemic, such as Italy's neighbour Austria, and China's neighbour Taiwan. The other way around, countries that reacted only after the pandemic had hit them with full force included countries such as Brazil and Mexico, which compared to Asian-Pacific and European countries had had much more time to react.

Table 3: Timing of targeted travel restrictions relative to the tenth domestic death

	Number of countries (percent of total)	Average date of introduction across four measures	Examples
All four measures before	118 (67.4%)	11 March	Australia, Austria, Czech Republic, Israel, Taiwan
One after	11 (6.3%)	1 April	Albania, Iceland, Malaysia, Sri Lanka, Turkey
Two after	29 (16.6%)	1 May	Canada, Denmark, Latvia, Morocco, Portugal
Three after	5 (2.9%)	16 April	Indonesia, Japan, Luxemburg, Sweden, United States
All four measures after	12 (6.9%)	17 May	Brazil, Belgium, Germany, Iran, United Kingdom

To nonetheless exclude the possibility that the available time to act and the misfortune of being hit early would influence our estimates of the impact of timely travel restrictions, I control in the regressions reported in Table 4 for the timing of the tenth death.³⁴ This is a conservative procedure, because the timing of the tenth death may to some extent itself be a result of travel restrictions or a lack of them.

³⁴ In additional analyses not reported in the table, I excluded all countries that never had a tenth death from the analysis. Even though this excludes one third of the countries, the coefficients of the travel restriction variable hardly change.

Table 4: Regressions of COVID-19 mortality on timing of travel restrictions relative to the timing of the tenth local death, exposure to international travel and control variables

	Negative binomial regression	Linear regression
Per capita GDP in 1000 USD	1.011	.079
Median age	1.025	-.019
Democracy and human rights (0-100)	1.003	.056+
Air travel exposure (AEF index 0-100)	.990	-.027
Tourist arrivals in millions	.997	.151+
Island state	1.401	-3.783*
Number of targeted restrictions introduced before 10th death: None	Ref.	Ref.
One	.446	-12.878+
Two	.337**	-22.600+
Three	.256+	-25.133*
All Four	.270**	-23.921*
Timing of 10 th death	.973***	-.023
<i>Clustered by world region</i>	<i>Yes</i>	<i>Yes</i>
N=	175	175
R ²	.060	.570

The results of both negative binomial and linear regressions show that once we account for the timing of travel restrictions relative to the local timing of the pandemic, not only almost all the significant effects of GDP and democracy disappear (only democracy remains marginally significant in the linear regression), but also those of air travel exposure and tourist inflows. Only the latter remain marginally significant in the linear regression, but with a more than halved coefficient size compared to Table 2. The reason for this reduced impact of the travel exposure measures is their relationship with the likelihood of countries to implement travel restrictions at the very early stages of the local pandemic. Countries that occupy a central role in the international air travel network and that receive high numbers of annual tourist visitors were much less likely to implement travel restrictions early ($r = -.42$ and $-.45$, respectively; both $p < .001$). Because of their economic dependence on international business and leisure travel, these countries seem to have been more inclined to follow the WHO's advice in the early stages of the pandemic, which claimed that travel restrictions were not effective and economically harmful.

Table 4 shows that travel restrictions taken during the early stages of the local epidemic are associated with much lower cumulative COVID-19 mortality. Compared to countries that took all four targeted measures after the tenth domestic death had occurred (or that refrained from such measures entirely), countries that took at least three out of four measures before the tenth death had on average a 73-74 percent lower death toll. Countries that took two out of four or at least one of the four measures early still had 66, respectively 55 percent fewer deaths. The linear regression points in a similar direction, with on average 22 to 25 fewer deaths per 100,000 inhabitants in countries that took at least two of the measures before the tenth death, and 13 per 100,000 less in those that took one measure early enough. The fact that the regressions control for the exact timing of the tenth death indicates that we are not dealing here with selection effects of countries simply having had the bad luck or good fortune of being hit by the pandemic early or late.

Based on the results of Table 4, column 1, we can estimate how many deaths could have been prevented if all countries had implemented all four targeted measures before the tenth domestic death. The model predicts for countries that implemented all four measures in time an average reduction of deaths by 73 percent (95 percent confidence interval from 48 to 98 percent) compared to those that implemented none of the measures in time. Considering that among the latter group of countries 227,000 deaths occurred, this implies that even if we take the lower bound of the 95 percent confidence interval (a reduction by 48 percent), at least 109,000 deaths could have been prevented by restricting international travel early. In an average scenario of a 73 percent reduction, we would be talking about 165,000 deaths: Add to this smaller reductions in mortality for countries that implemented only one or two measures in time and we may conclude that the effect size of travel restrictions even in a very cautious interpretation of the results is very substantive.

These calculations may moreover still underestimate the true global effect of travel restrictions, because they do not account for the spillover effects of early and encompassing restrictions in some countries on the exposure to international travel of countries that took such measures later or not at all. For instance, even though beyond the EU-wide entry ban against non-

Schengen countries, Sweden did not implement any additional restrictions, many countries imposed restrictions that included Sweden. In combination with the worldwide collapse of air travel, even countries such as Sweden or Mexico, which implemented hardly any restrictions themselves, profited from the travel restrictions implemented by others.

5. Robustness checks

5.1. Reverse causality?

That travel restrictions, if they had been implemented more timely, could have saved many lives is a strong claim, which of course needs to be put to a variety of robustness checks. An important potential problem is reverse causality. The above analyses assume that travel restrictions affect epidemic outcomes, but although in some countries these restrictions were taken entirely preemptively, before any local cases of COVID-19 had occurred, in other countries, they were at least in part a *reaction* to the domestic spread of COVID-19 cases and deaths. To the extent that countries that were hit by the pandemic earlier and more severely were more likely to implement travel restrictions early on than countries that were hit late and lightly, our above estimates of travel restriction effects might therefore be biased. The correlations between the timing of travel restrictions and indicators of early pandemic severity do not point in this direction, however. If travel restrictions were primarily reactive rather than pre-emptive, we should observe that early travel restrictions were more likely to appear in countries that experienced their first COVID-19 cases and deaths early on in the pandemic. As it is, however, the correlation between the timing of travel restrictions and the timing of the first and tenth cases and deaths in a country are throughout negative (between $-.23$ with the timing of the first death – $p < .01$ – and $-.30$ with the timing of the tenth death – $p < .001$). In other words, countries that were hit early by the pandemic were actually *less* likely to implement travel restrictions or, alternatively, they had earlier outbreaks because they did not implement travel restrictions.

A conservative way to eliminate reverse causality is to control the regressions for indicators of early epidemic severity. Such a procedure is conservative because it assumes that early epidemic severity, or its opposite, being hit late and lightly, are themselves not affected by travel restrictions or their absence. Yet, even under such conservative modelling scenarios, the moderating effect of travel restrictions on mortality remains significant and of similar size. The first column of Table 5 controls for early epidemic severity by including the timing of the first and tenth cases and deaths as predictors. For this regression, I exclude countries that never experienced ten deaths until 30 June, for the simple reason that travel restrictions up to 30 June cannot be a reaction to deaths that did not occur before that date. The results in column 1 of Table 5 show that the timing of the first cases and deaths does not significantly predict cumulative mortality. More importantly, controlling for these indicators of early epidemic severity does not in any way reduce the effect of travel restrictions. What does change is that air travel exposure becomes statistically insignificant, which is, as separate analyses show, due to the exclusion of countries that never had ten deaths until 30 June, which tend to have low air travel exposure.

Table 5: Negative binomial regressions of COVID-19 mortality additionally controlling for early epidemic severity

	Controlling for timing of 1 st and 10 th cases and deaths	Controlling for death toll as of 13 March	Only countries with >= 1 deaths as of 13 March
Per capita GDP in 1000 USD	1.016+	1.009	1.021
Median age	1.029	1.019	.941
Democracy and human rights	1.001	1.010	1.015+
Air travel exposure (AEF index 0-100)	.984	1.037***	1.053*
Tourist arrivals in millions	1.006	.999	1.003
Island state	.560	.428	.102***
Travel restrictions scale (6 items)	.993***	.993**	.987**
Timing 1 st case	.985	-	-
Timing 10 th case	1.010	-	-
Timing 1 st death	1.010	-	-
Timing 10 th death	.974	-	-
Death toll as of 13 March	-	1.003***	-
<i>Clustered by world region</i>	Yes	Yes	Yes
N=	119	175	35
R ²	.039	.042	.056

Another problem related to the local timing and severity in the early phases of the pandemic might be that countries that were hit early had more time to accumulate deaths until 30 June, and that this, rather than early or late travel restrictions, explains why some countries have higher cumulative death rates. To eliminate this bias, we can control for mortality levels early in the pandemic as a predictor of the cumulative death toll until 30 June. In the second column of Table 6, I add countries' death toll on 13 March, the day of the US travel ban against Europe, as a predictor. The results show that, not unexpectedly, the number of deaths on 13 March significantly predicts cumulative mortality on 30 June. However, both the travel exposure and travel restrictions variables remain highly significant and of the same magnitude. In other words, travel exposure and restrictions explain cross-national differences in COVID-19 mortality even

when we control for mortality levels earlier in the epidemic. This is again a very conservative test because mortality levels as of 13 March may themselves have been affected by earlier travel restrictions (see also Table 3 above). Until 12 March, almost half of the analysed countries had already implemented travel restrictions on China and more than 40 percent had such measures in place for travellers from Italy. Very similar results are obtained taking the death toll for other days in March.

Yet another way of dealing with the problem of countries being hit early or late is to only look at variation among the former group. I define countries that were hit early as those that had registered at least one death as of 13 March, which applies to 35 states, 15 of which are in Europe, the others distributed across all other regions except sub-Saharan Africa. The results (column 3 of Table 6) are again remarkably consistent and become stronger rather than weaker: 1.3 percent fewer deaths per day travel restrictions are introduced earlier.

5.2 Considering other COVID-19 containment policies

While the analyses so far show that travel restrictions have consistent, substantial and statistically significant effects, these could potentially be attributable in whole or in part to correlated domestic containment policies. Table 6 shows that adding the domestic restrictions scale based on Oxford University's "Coronavirus Government Response Tracker" to the negative binomial regression (for the 156 independent countries that the Oxford data cover) does not alter the result for travel restrictions; the effect size even increases. Surprisingly, the domestic restrictions scale is significantly positively associated with COVID-19 mortality. A reinforcing effect of domestic restrictions on mortality is of course highly implausible theoretically, and the result therefore most likely means that the relationship between mortality and domestic restrictions is dominated by the reverse causal path, where countries are more likely to implement extensive domestic restrictions when they are severely hit by the pandemic. Be that as it may, the important finding that concerns us here is that the effect of travel restrictions does not diminish when we control for domestic pandemic containment measures. This applies also (not shown in the table) if we add domestic measures to the linear regression of mortality

per 100,000 inhabitants, as well as when we add single domestic measures rather than the scale.

Table 6: Negative binomial regressions of COVID-19 mortality using Oxford University data on domestic COVID-19 containment measures as an additional control and using Oxford data on travel restrictions as an alternative predictor

	Domestic containment measures as an additional control variable	Using alternative Oxford University travel restrictions measure	Using Oxford and my own travel restriction measures simultaneously
Per capita GDP in 1000 USD	1.015	1.009	1.009
Median age	1.055	1.040	1.036
Democracy and human rights	1.001	1.005	1.003
Air travel exposure (AEF index 0-100)	1.035***	1.048***	1.044***
Tourist arrivals in millions	1.001	1.004	1.001
Island state	.309*	.309*	.311*
My travel restrictions scale (6 items)	.987***	-	.993***
Oxford domestic containment measures scale	1.017***	-	-
Oxford travel restrictions scale	-	.989*	.991+
<i>Clustered by world region</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
N=	156	156	156
R ²	.049	.044	.045

The Oxford data allow a further robustness check because they also include a measure of international travel restrictions. Unfortunately, these data are coded as an additive scale that lacks theoretical or empirical justification. Introduction of screening measures (e.g., temperature checks) at points of entry is considered as the lowest level of travel restrictions, followed by quarantines as level 2. Both can refer to only one target country, to all countries of the world, or anything in between. Level 3 consists of targeted entry bans (for at least one country) and level 4 of global entry bans for all foreign countries. Only one code applies to any given point in time, and

higher-level measures override lower-level measures. For instance, quarantines are only coded if no entry bans are in place, possibly even if the quarantine measure applies to all countries and the entry ban only to one country. This way of coding also presupposes that entry bans are more effective than quarantines, whereas the above analyses indicate the reverse. An advantage of the Oxford data compared to mine is that they include screening measures, but since these are seen as the lowest-level measure, they are only coded if no quarantines or entry bans are in place. The data also seem to contain some coding errors.

These limitations of the Oxford data notwithstanding, we can assess whether we find similar effects of travel restrictions if we use these data instead of mine.³⁵ Because of the additive nature of the Oxford data, one needs to make a choice which level to use as the threshold. Results turn out to be strongest when we simply take the introduction date of the first travel restriction, regardless of its type/level. This variable correlates moderately strong with my travel restrictions scale ($r = .41$; $p < .001$). The second column of Table 6 shows that, even though the two travel restriction measures are far from identical, using the Oxford data, we find a significant suppressing effect on mortality of a similar order of magnitude as in the analyses above. Finally, in the third column of the table, we look what happens when we include both travel restriction measures simultaneously. Both measures then have effects of similar magnitudes, although the Oxford measure is only marginally significant because of higher standard errors. That the Oxford measure has an effect over and above my own travel restrictions scale may be because they measure travel restrictions in different ways – by way of the first measure introduced, respectively the average introduction date of several types of measures – and because the Oxford data include screening measures that are not included in my data. Be that as it may, the important result is that we arrive at the same conclusion regardless of whether we use the Oxford data on travel restrictions, my own data, or both.

³⁵ A direct comparison of the two datasets is difficult because they are differently structured. The most direct comparison that is possible pertains to the timing of targeted entry bans (which refer to China or Italy in my data, and to any country in the Oxford data). The timing of such measures correlates at $r = .50$ ($p < .001$) across the two datasets, which is reasonably high but also well below conventional levels of intercoder reliability.

5.3. Different observation points of cumulative mortality

Next, I explore to what extent the results for travel exposure and restrictions depend on the date at which we observe the cumulative COVID-19 death toll. As indicated above, earlier research from China (Jia et al. 2020) suggested that the explanatory power of travel flows from Wuhan to other parts of China increased over time. On the other hand, the results of mathematical modelling studies suggest that travel exposure and restrictions have their strongest effects early on in a pandemic and get overwhelmed by domestic transmission later on. Depending on when during the course of a pandemic one observes cumulative mortality may therefore affect the estimate of the impacts of travel exposure and restrictions. To investigate this, Table 7 replicates the results of the first column of Table 2 for the observation point 30 June (third column), and shows how results compare when the cumulative death toll is measured at later or earlier dates.

Table 7: Negative binomial regressions of COVID-19 mortality for different observation points of the cumulative death toll

	Death toll 31/08	Death toll 31/07	Death toll 30/06	Death toll 31/05	Death toll 01/05	Death toll 27/03
Per capita GDP in 1000 USD	.999	1.000	1.008	1.012	1.013*	1.014+
Median age	1.009	1.012	1.021	1.030	1.078*	1.080
Democracy and human rights (0-100)	1.008	1.007	1.009	1.016+	1.016+	.997
Air travel exposure (AEF index 0-100)	1.040**	1.042***	1.037***	1.037***	1.046***	1.053**
Tourist arrivals in millions	.997	.999	1.002	1.003	1.002	1.011
Island state	.326*	.341+	.434	.513	.506	1.698
Travel restrictions scale (6 items; higher score means earlier introduction)	.993*	.992***	.992***	.988***	.979**	.917**
<i>Clustered by world region</i>	Yes	Yes	Yes	Yes	Yes	Yes
N=	175	175	175	175	175	175
R ²	.026	.031	.041	.062	.105	.117

The results show that the impact of travel flows as measured by the AEF index is remarkably stable over time. Being an island state becomes more relevant over time, and becomes statistically significant for the two most recent observation points (as it already was in the linear regressions and in some of the robustness checks above). The travel restrictions variable is significant across all observation periods, but its effect size declines the later we measure the death toll. This may either mean that over time, as suggested by simulation models, domestic factors become more important in differentiating outcomes among countries, or that my measures of travel restrictions become less efficient and valid later on in the pandemic. The latter may well be the case, because China ceased to be a major source of international seeding by the end of February, and the same was probably true for Italy by the end of March or mid-April. From then on, other countries, such as the United States or Brazil, are likely to have become more important sources of seeding. Since I only coded targeted entry bans and quarantines pertaining to China and Italy and not to other countries, the explanatory power of my travel restriction measures may have waned as the pandemic spread to the Americas and Africa. An additional reason may be that I only coded the introduction date of restrictions and not their later cancellation or relaxation. New travel restrictions after 30 June were moreover not considered because my coding ended on 30 June. More complete and refined measurement of travel restrictions later in the corona pandemic might therefore reveal stronger travel restriction (and relaxation) effects at later observation points in the pandemic. Nonetheless, the results of this robustness check show that, at least until the end of August 2020, the main results regarding the importance of travel exposure and restrictions do not depend on the date at which we observe the COVID-19 death toll.

5.4. Further robustness checks

Several additional robustness checks can be found in the Appendix. They show that results are similar when we use world region fixed effects; when we use alternative operationalizations of countries' affluence, age composition, and democracy; and we restrict the scope of the analysis to various subsets of countries.

6. A quasi-experimental approach

The preceding analyses have shown that the results are stable across a wide range of model specifications and controls for possible confounding factors. Yet, these analyses ultimately remain correlational and make a causal effect of travel restrictions plausible, but cannot prove it. In a methodologically ideal world, we would have assigned travel restrictions randomly to the world's states. This would have allowed us to exclude all possible confounders and measure the pure causal effect of travel restrictions. Of course, such an approach is neither practically feasible nor ethically defensible and we therefore have to work with the natural variation that governments' decisions during the COVID-19 pandemic have provided us with. Multivariate regressions controlling for plausible confounding variables are one way of approaching this kind of data.

Another way is to approach the problem from a quasi-experimental angle and try to approach an experimental design as much as possible. This implies viewing travel restrictions as the treatment variable and to create a treatment group that receives the treatment and a control group that does not. The treatment and control groups must moreover be balanced on all other relevant variables, i.e. they should differ only on the treatment variable. In the case at hand, the treatment cannot simply be defined as the absence or presence of travel restrictions because only two out of 175 countries included in the analyses – Ireland and Mexico – never introduced any of the six types of restrictions that I distinguish. Where countries differ, however, is in whether they introduced travel restrictions early or late.

I therefore constructed a dichotomous variable akin to the one used above based on the Oxford University data, which measures the day on which the first travel restriction, regardless of its scope or type, was introduced. Based on the distribution of this variable, I then construct a treatment group of countries that introduced travel restrictions early and a control group that introduced them late or, in a few cases, not at all. In a first variant ("travel restriction treatment 1"), all countries that introduced travel restrictions up until the modal introduction day are classified as the treatment group of early adopters, and the other half of the sample with

later introduction dates as the control group of late adopters. Obviously, this dichotomization throws away a lot of information in the data on the scope (from one target country to the whole world), type (quarantine or entry ban) and timing (from marginally earlier or later than the modal value to very much earlier or later) of travel restrictions. Moreover, it entirely disregards the timing of restrictions relative to the phase of the local epidemic. To reduce at least the loss of information related to timing and the somewhat arbitrary nature of the modal value as the cut-off point, I constructed a second treatment and control group (“travel restrictions treatment 2”), defined respectively as the one third of countries with the earliest and latest introduction dates (the middle third of the distribution is disregarded in this case). Countries in the earliest one third include for instance Australia, the Czech Republic, Israel, Russia, Singapore, Turkey, the United States, and Vietnam. While some of these countries, such as Israel, were generally early with encompassing travel restrictions, others, such as the United States, had both fewer restrictions and were early with only one of them (in the case of the USA the entry ban on China of 2 February). This indicates that even the earliest one third treatment group is quite heterogeneous, which leads to relatively high standard errors of the estimates.

Different from a real experiment, selection into these natural treatment groups is not random. Some countries (especially island states, countries with low travel exposure, and countries in Asia and the Pacific) are more likely to be in the treatment group. To correct for this selection bias, I employ a variant on propensity score matching called inverse probability of treatment weighting (see Thoemmes and Ong 2015). It consists of first calculating from a logistic regression predicted probabilities of belonging to the treatment group for each case. These probabilities are then used to weight the data in such a way that the treatment and control groups do not differ significantly anymore on observed confounder variables.³⁶ As a result, the estimate of the treatment effect is no longer biased by observed confounders and can be more confidently interpreted in causal terms. Unlike a real experiment, unobserved confounders may still affect selection

³⁶ The treatment group is weighted by the inverse of the predicted probability: $1/\text{predprob}$; the control group by the inverse of its probability: $1/(1-\text{predprob})$. I thank Max Schaub for suggesting this method.

into the treatment, but if there are no plausible confounders that have not been measured and if measurements of the observed confounders are reasonably adequate, this method closely approximates random assignment.

Table 8: Quasi-experimental treatment effects of early introduction of travel restrictions

	Negative binomial: Cumulative number of deaths on 30 June 2020 (IRR coefficients)		Linear: Cumulative deaths per 100,000 inhabitants on 30 June 2020 (B coefficients)	
Travel restrictions treatment 1 (early vs. late half of sample)	.593+	-	-2.907	-
Travel restrictions 2 (early vs. late third of sample)		.383*	-	-4.863*
N=	175	119	175	119
R ²	.002	.011	.016	.044

I use all variables that were used in the above regressions as predictors of selection into the treatment groups, including the region in which a country is situated. When weighting by inverse probabilities, the treatment and control groups are indeed balanced and do not differ significantly on any of the possible confounder variables anymore. Table 8 shows how the treatments affect the outcome variables, the absolute number of deaths and the number of deaths per 100,000 inhabitants. Columns 1 and 3 show that the treatment definition that divides the country sample in half around the mode of the distribution is only marginally significant in the binomial regression, and just above the $p < .10$ threshold in the linear regression. This may indicate that too much information is lost by bringing back the measurement of travel restrictions to a dichotomy. In addition, the mode may not be the optimal dividing line (especially because a relatively large number of countries is situated, on both sides, close to the mode; see Figure 3 above).

In line with this interpretation, the treatment-control definition that contrasts the earliest one third of countries with the latest third performs much better. The effects are now significant in both regressions, and they are of substantial size. In the negative binomial regression, the effect size indicates a 62 percent (95% confidence interval between 33 and 90 percent)

lower mortality rate ($p < .05$) in countries that were among the earliest one third to introduce travel restrictions, compared to the latest third of countries. The linear regression results indicate almost five (4.9) deaths ($p < .05$; 95% confidence interval from 0.9 to 8.9) per 100,000 inhabitants fewer in early than in latecomer countries. This effect size is almost equal to the distribution mean of per 100,000 capita deaths of 5.6.

7. What about China and Italy?

Because they were the targets of entry bans and quarantines that were central to this study, I have excluded China and Italy from all preceding analyses. But obviously, these are two important cases in their own right. China was the starting point of the pandemic, and Italy was its most important accelerator as the most important source country for early seedings in most European countries (see above), as well as in places such as Brazil and the East Coast of the United States. How do these two countries fit into the story?

In terms of deaths per 100,000 inhabitants, China and Italy are on opposite ends of the spectrum. With 56 deaths per 100,000 inhabitants Italy was, as of 30 June 2020, the fourth-hardest hit country, whereas China registered only 0.3 deaths per 100,000 inhabitants. Even in absolute terms, the number of deaths in Italy (34,744) was more than seven times higher than in China (4,634). The two countries also differed markedly in their response to the early onset of the pandemic, particularly where domestic travel restrictions were concerned.

The first death as a result of the novel coronavirus was registered in China on 10 January in the city of Wuhan in Hubei province. Although initially the Chinese government denied – as we saw, echoed by the WHO – that there was evidence of human-to-human transmission and suppressed critical voices, by the second half of January it implemented strict containment measures that went beyond the WHO's recommendations. On 23 January, China imposed an internal travel ban, which precluded all travel in and out of Wuhan and other cities in Hubei province. At the time, 17 COVID deaths

had been recorded, all of which had occurred within Hubei province.³⁷ As several studies we have discussed in detail above indicate (Chinazi et al. 2020; Lau et al. 2020; Kraemer et al. 2020; Pinotti et al. 2020), these domestic travel restrictions strongly reduced seeding both to other parts of China, and to other countries of the world.

If we look at pandemic spread in China from the perspective of the current study, we can conceive of the travel restrictions imposed on 23 January as an entry ban by the rest of China against Hubei province. This entry ban was both timely – it occurred before a single death had been registered outside of Hubei province – and very effective. Against the argument that travel restrictions have little power to mitigate the spread of an epidemic, the isolation of Hubei province from the rest of the country had the effect of freezing the Wuhan/Hubei vs. the rest of the country differential: by 2 August, according to Chinese government data, no less than 97 percent of all COVID-19 deaths had occurred within Hubei province (57 million inhabitants).³⁸ The low per capita mortality in China as a whole is entirely due to the area outside Hubei province. In Hubei itself, the mortality rate as of 2 August was 7.6 per 100,000 inhabitants (in Wuhan city even 35 per 100,000), which is close to the global average on that date and higher than in three quarters of the countries of the world. The rest of China, by contrast, had a mortality rate of only 0.009 per 100,000, by which it ranges among the 13 percent least affected countries worldwide. Even if we factor in the possibility of underreporting of deaths by the Chinese government – but recall that none of the above analyses pointed in the direction of strong underreporting of mortality in authoritarian states generally – this would not explain the stark difference between Hubei and the rest of China. Everything points in the direction that the strict travel ban from 23 January onwards played a major role in containing the epidemic and in saving the rest of China from the fate of Wuhan city and Hubei province. The Chinese story is therefore fully compatible with the findings of the current study.

³⁷ <https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200123-sitrep-3-2019-ncov.pdf>.

³⁸ https://en.wikipedia.org/wiki/COVID-19_pandemic_in_mainland_China.

But what about Italy? Italy seems to be a prominent contrary case because it belonged to the countries that imposed an entry ban on travellers from China early on, beginning on 31 January. The easy way out would be to say that Italy is just an exception to the rule, and that first cases had been registered already in Italy before the entry ban. The first officially registered cases in Italy were a Chinese tourist couple from Wuhan, who had entered Italy via Milano's Malpensa airport on 23 January, visited Verona and Parma as well as various tourist sites in Rome, where they were hospitalized and diagnosed with COVID-19 on 31 January. A study of traces in waste water suggests that COVID-19 was already circulating in the Northern Italian cities of Milano and Torino as early as December 2019.³⁹

However, there are also more structural reasons why the Italian entry ban on China may not have been as effective as entry bans elsewhere. Like many other European countries, Italy is part of the Schengen zone in which internal border controls have been abolished. Therefore, the entry ban on travellers from China practically only affected passengers directly arriving to Italy by air. Passengers arriving across the land borders would not have been checked, nor were air travellers from China who arrived via a stopover in another Schengen country, where they would have cleared customs and been able to continue unhindered towards Italy. Moreover, because, with the exception of Czechia, which introduced an entry ban on China on 8 February, no other Schengen countries introduced travel restrictions on China until well into March, infection chains originating in China could spread unhindered from other Schengen countries to Italy. The source of the infection of „patient zero“ of the Northern Italy outbreak (a manager of a local Unilever branch) remains unknown, but genetic evidence suggests that the virus may have arrived in Northern Italy via Southern Germany.⁴⁰

The first two deaths in Italy occurred on 22 February, one in Lombardy region, the other in Veneto. From the next day, Italy prohibited all travel to and from the two municipalities where the victims originated as well as a

³⁹ <https://www.bbc.com/news/world-europe-53106444>.

⁴⁰ <https://www.reuters.com/article/us-health-coronavirus-italy-scientists-idUSKBN20Y35B>.

few surrounding ones.⁴¹ That area was very small though and had a total population of only 54,000. Moreover, cases of coronavirus had already been registered in other parts of Lombardy and Veneto as well as in the regions of Emilia-Romagna, Piemonte and Lazio. It was not until 8 March that Italy implemented more wide-ranging travel restrictions, banning travel in and out of three non-contiguous areas across fourteen Northern provinces in five regions (with a total of 16 million inhabitants). By then, however, 366 people had already died, at least 13 of which outside the quarantine zone.⁴² Compared to China, which implemented a travel ban on a region with as many inhabitants as the whole of Italy at a moment when only 17 deaths had been registered, which moreover all had occurred within the quarantine region, Italy imposed significant internal travel restrictions much too late, at a moment when the epidemic had spiralled out of control and had already caused 366 deaths, some of which as far away from the Northern epicentre as the capital region Lazio and Puglia in the far South. Italy too, then, upon closer examination conforms to the general pattern of the findings of this study.

8. Conclusions

This study has investigated the roles of exposure to and restrictions of international travel in explaining cross-national differences in COVID-19 mortality across the global population of independent states. Theoretically, the study draws on sociological research on diffusion and social networks, which has highlighted the importance of so-called “weak ties” for the spread of social innovations and information. Weak ties connect people whose social networks hardly overlap, whereas “strong ties” connect people who share many network ties. Looking at world society from this perspective, international connections are typical weak ties, because the

⁴¹ https://de.wikipedia.org/wiki/COVID-19-Pandemie_in_Italien.

⁴² The figure of at least 13 deaths is based on deaths that had occurred in regions other than the five affected by the 8 March travel ban; see https://www.corriere.it/salute/malattie_infettive/20_marzo_08/coronavirus-italia-dati-8-marzo-bollettino-regionali-provinciali-8729bd9c-6157-11ea-8f33-90c941af0f23.shtml. However, not all provinces of the five regions were included in the quarantine zone. I have not been able to find provincial-level data to establish how many people had died in provinces of these regions that remained outside the quarantine zone.

density of social networks is much higher within than between countries. This led to the expectation that cross-national connections would be particularly conducive to the spread of a contagious virus such as COVID-19.

The sociological perspective on diffusion also allows one to derive which network connections should be the preferred target of attempts to slow down the spread of an unwanted social innovation or, by extension, of a contagious virus. The cutting of weak ties will, other things being equal, be much more efficient in containing viral spread than the cutting of strong ties. This led to the expectation that restrictions of international travel would be an efficient strategy for containing viral spread in a global social network.

Previous epidemiological research, especially simulation and modelling studies, has often been sceptical regarding the merits of travel restrictions. This conclusion is understandable in view of the fact that because of the exponential nature of viral growth, domestic spread will eventually overwhelm the contribution of seeding from abroad. However, I argue that this reasoning underestimates the importance of gaining time in the initial phase of a pandemic and keeping the number of cases at a low enough level for contact tracing and isolation to be feasible strategies. The more seeding from abroad occurs, the more independent chains of transmission will be created, and the more difficult domestic containment will become. At any rate, the ultimate test of the relevance of travel restrictions cannot rest on model simulations but needs to be conducted on the basis of a confrontation with real-world data.

This is what the current study has aimed to do. To this end, I used publicly available data on COVID-19 deaths and exposure to international travel, as well as data that were newly gathered for the purpose of this study on the exact timing of the introduction of six types of travel restrictions: entry bans and mandatory quarantines, respectively targeting China, Italy, or all foreign countries. The findings show that higher exposure of countries to international travel (as measured by the air travel exposure index AEF or the yearly number of tourist arrivals) is strongly and consistently associated with higher COVID-19 mortality. Island states, by contrast, which

have lower exposure to cross-national travel because of their lack of land borders, have much lower mortality rates.

A further key finding was that early introduction of policies that restrict international travel (entry bans and quarantines) strongly and consistently reduces the COVID-19 death toll. “Early” here means both early in absolute calendar time, and relative to the local timing of the pandemic. Travel restrictions were especially powerful when countries introduced them before the local pandemic had passed a certain threshold, which seems to lie around the time of the 10th domestic death. Among different types of travel restriction policies, mandatory quarantines were more effective than entry bans. The reason is likely to be that entry bans in most cases contain exceptions for returning citizens and permanent residents and therefore exclude an important part of traveller inflow. Quarantines, by contrast, usually apply to all incoming travellers, regardless of nationality or residence status. Targeted travel restrictions (here measured through entry bans and quarantines for travellers from China and Italy) turn out to be more efficient than global restrictions that target all foreign countries. While general restrictions are effective to the extent that they encompass restrictions on high-risk countries, they have no measurable added value beyond what targeted travel restrictions can achieve.

These results for travel restrictions hold across a wide range of model specifications and robustness checks (including additional controls for domestic containment policies such as school closures and bans on public gatherings), as well as in a quasi-experimental design that compares treatment and control groups that differ only regarding the timing of travel restrictions.

The effect sizes of travel restrictions are substantial. Comparing in the quasi-experimental design the early adopter group of countries that were among the one third of the sample that introduced travel restrictions the earliest, to the latecomer group consisting of the one third that introduced travel restrictions the latest (or not at all), we find that early adopters have an estimated 62 percent lower COVID-19 mortality rate. Regression results indicate a 0.8 percent reduction in cumulative mortality per day that travel restrictions were introduced earlier.

Beyond the effects of exposure to international travel and of travel restrictions, the study provides some evidence that more affluent countries and democracies have higher COVID-19 mortality. The reason likely is not that these countries have higher actual numbers of deaths, but that they report more deaths because of their more developed health systems and greater willingness to admit the true extent of the pandemic. These effects are however relatively small compared to the effects of travel exposure and restrictions.

The findings of the study have important policy implications. Countries that are highly exposed to international travel because of their centrality in airline networks and high tourist flows should be aware that they run a much-increased risk of early and multiple seeding from pandemic source regions. Contrary to the common wisdom during the early phases of the COVID-19 pandemic and the recommendations of the World Health Organization at the time, restrictions of international travel are an efficient means of pandemic containment, especially if they are implemented when domestic case and fatality numbers are still low.

9. References

- Askitas, Nikos, Konstantinos Tatsiramos & Bertrand Verheyden. 2020. Lockdown strategies, mobility patterns and COVID-19. *Covid Economics*. Issue 23, 28 May 2020.
- Block, Per et al. 2020. Social network-based distancing strategies to flatten the COVID-19 curve in a post-lockdown world. *Nature Human Behaviour* 4: 588-596.
- Bonardi, Jean-Phillippe. Quentin Gallea, Dimtrija Kalanoski & Rafael Lalive. 2020. Fast and local: How lockdown policies affect the spread and severity of COVID-19. *Covid Economics*. Issue 23, 28 May 2020.
- Brockmann, Dirk & Dirk Helbing. 2013. The hidden geometry of complex, network-driven contagion phenomena. *Science* 342: 1337- 1342.
- Chinazzi, Matteo et al. 2020. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* 368: 395-400.
- Costantino, Valentina, David J. Heslop & C. Raina Macintyre. 2020. The effectiveness of full and partial travel bans against COVID-19 spread in Australia for travelers from China during and after the epidemic peak in China. *Journal of Travel Medicine* 27, 5.
- Epstein, Joshua M. et al. 2007. Controlling pandemic flu: the value of international air travel restrictions. *PLOS ONE* 2(5): e401.
- Ferguson, Neil M. et al. 2006. Strategies for mitigating an influenza pandemic. *Nature* 442: 448-452.
- Ghemawat, Pankaj & Steven A. Altman. 2014. *DHL Global Connectedness Index 2014. Analyzing global flows and their power to increase prosperity*. Bonn: DHL.
- Granovetter, Mark S. 1973. The strength of weak ties. *American Journal of Sociology* 78: 1360-1380.
- Hale, Thomas, Noam Angrist, Emily Cameron-Blake, Laura Hallas, Beatriz Kira, Saptarshi Majumdar, Anna Petherick, Toby Phillips, Helen Tatlow, Samuel Webster. 2020. *Oxford COVID-19 Government Response Tracker*. Oxford: Blavatnik School of Government.

- Hâncean, Marian-Gabriel, Matjaz Perc & Jürgen Lerner. 2020. Early spread of COVID-19 in Romania: imported cases from Italy and human-to-human transmission networks. *Royal Society Open Science* 7: 200780.
- Harris, Jeffrey E. 2020. The subways seeded the massive coronavirus epidemic in New York City. *National Bureau of Economic Research Working Paper* No. 27021.
- Jia, Jayson S., Xin Lu, Yun Yuan, Ge Xu, Jianmin Jia & Nicholas A. Christakis. 2020. Population flow drives spatio-temporal distribution of COVID-19 in China. *Nature* 582: 389-394.
- Kee, Kelley, Catherine Z. Worsnop, Karen A. Grépin & Adam Kamradt-Scott. 2020. Global coordination on cross-border travel and trade measures crucial to COVID-19 response. *The Lancet* 395: 1593-1595.
- Koopmans, Ruud. 2016. Does assimilation work? Sociocultural determinants of labour market participation of European Muslims. *Journal of Ethnic and Migration Studies* 42: 197-216.
- Kraemer, Moritz U.G. et al. 2020. The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science* 368: 493-497.
- Lancee, Bram. 2010. The economic returns of immigrants' bonding and bridging social capital: The case of the Netherlands. *International Migration Review* 44: 202-226.
- Lau, Hien et al. 2020. The association between international and domestic air traffic and the coronavirus (COVID-19) outbreak. *Journal of Microbiology, Immunology and Infection* 53: 467-472.
- Lawyer, Glenn. 2015. Understanding the influence of all nodes in a network. *Scientific Reports* 5: 8665.
- Lawyer, Glenn. 2016. Measuring the potential of individual airports for pandemic spread over the world airline network. *BioMed Central Infectious Diseases* 16: 70.
- Leung, Kathy, Joseph T. Wu & Gabriel M. Leung. 2020. First-wave COVID-19 transmissibility and severity in China outside Hubei after control measures, and second-wave scenario planning: a modelling impact assessment. *The Lancet* 395: 1382-1393.

- Maier, Benjamin F. & Dirk Brockmann. 2020. Effective containment explains subexponential growth in recent confirmed COVID-19 cases in China. *Science* 368: 742-746.
- Mateus, Ana L.P. et al. 2014. Effectiveness of travel restrictions in the rapid containment of human influenza: a systematic review. *Bulletin of the World Health Organization* 92: 868-880.
- McPherson, Miller, Lynn Smith-Lovin & James M. Cook. 2001. Birds of a feather: homophily in social networks. *Annual Review of Sociology* 27: 415-444.
- Pinotti, Francesco et al. 2020. Lessons learnt from 288 COVID-19 international cases: importations over time, effect of interventions, underdetection of imported cases. MedRxiv preprint doi: <https://doi.org/10.1101/2020.02.24.20027326>.
- Thoemmes, Felix & Anthony D. Ong. 2016. A primer on inverse probability of treatment weighting and marginal structural models. *Emerging Adulthood* 4: 40-59.
- UNWTO. 2018. *UNWTO Tourism Highlights. International Tourism Trends, 2018 edition*. Madrid: United Nations World Tourism Organization.

10. Appendix: Additional robustness checks

In the paper, I control for the spatial clustering of both the dependent and the predictor variables by employing standard errors clustered by seven world regions. This strategy is preferable to using world region as an additional predictor in a fixed-effects design, because the latter modelling strategy removes from the analysis any variation between regions. Nevertheless, some readers may wonder what the results would look like if we adopt a fixed-effects approach and only look at within-region variation.

Table A1 shows the results for the negative binomial and linear regressions paralleling Table 2 in the paper with added world region fixed effects (with Asia as the reference category). Compared to the reference category Asia, all other world regions, except in the negative binomial regression the countries of the Pacific, have higher mortality rates. This holds particularly strongly for Europe, the MENA region and South and Central America. Even though these world region fixed effects absorb a substantial part of the cross-national variation, the main findings for the variables of interest remain. In the negative binomial regression, the effect of air travel exposure remains significant and of similar size as in Table 2. The effect of travel restrictions is reduced somewhat from .992 in Table 2 to .993 here and is now only marginally significant. In the linear regression, tourist arrivals remain highly significant and of similar size. However, being an island state is no longer significant, which is due to the confounding role of the Pacific region, which with only one exception (Papua New Guinea) is almost entirely composed of island states. The travel restrictions' effect, finally, remains significant and is only slightly reduced in size from -.099 in Table 2 to -.087 here. In conclusion, even with world region fixed effects that remove an important part of the variation, the main findings remain standing.

Table A2 explores whether the results, taking the first column of Table 2 as a reference, change when we use alternative control variables or restrict the range of observations to subsets of the sample. In the first column of Table A2, I employ alternative control variables: instead of GDP, the Human Development Index (HDI); instead of median age, the population aged 65 and

over; and instead of the Freedom House index of democracy, the Polity IV score. The number of cases is lower for this regression especially because of the Polity score, which is only available for countries with at least 500,000 inhabitants. None of these alternative control variables is significant and the results for travel exposure and restrictions are robust. This is also the case if the alternative control variables are implemented one at a time, rather than all three in the same regression.

Next, I consider what happens if we limit the sample to countries for which we can have greater confidence in the accuracy of officially reported COVID-19 deaths. As I have argued in the paper, poorer countries may not have the capacity to test and treat all affected patients and may thus underreport the extent of the pandemic. Underreporting may also occur in authoritarian regimes, where rulers may want to hide the extent of the pandemic from the citizenry. Although the regressions control for these variables, this may not be sufficient to remove biases in the estimates entirely. In Table A2 (columns 2 and 3), I therefore restrict the regression sample to relatively affluent countries (with a per capita GDP of 10,000 USD or higher; column 2) or to democracies (with a Freedom House score of 70 or higher; column 3). The results for the travel restrictions scale are virtually identical, and even somewhat stronger, than in Table 2. Results for exposure to international travel are also similar, but unlike Table 2, the coefficients for island states are statistically significant (as in the linear regressions of Table 2), and indicate that island states had a 72 percent lower death toll than countries with land borders. Column 4 of Table 8 further tests to what extent the results are driven by very small countries and restricts the sample to countries with at least 500,000 inhabitants. This excludes a range of Caribbean, and Pacific island states as well as other small nations such as Belize and Brunei. The results again remain substantively unchanged.

Table A1: Regressions of COVID-19 mortality with world region fixed effects

	Negative binomial regression	Linear regression
Per capita GDP in 1000 USD	1.017*	.153**
Median age	1.014	-.277+
Democracy and human rights (0-100)	1.000	.045+
Air travel exposure (AEF index 0-100)	1.033*	.016
Tourist arrivals in millions	1.009	.358***
Island state	1.291	-2.563
Travel restrictions scale	.993+	-.087*
World region: Asia	Ref.	Ref.
World region: Pacific	.189*	3.392
World region: Middle East & North Africa	6.902***	3.991*
World region: Europe	7.041***	8.266**
World region: Sub-Saharan Africa	3.354**	1.314
World region: North America and Caribbean	2.961+	4.784*
World region: Central and South America	18.319***	7.563**
N=	175	175
R ²	.069	.480

Table A2: Negative binomial regressions of COVID-19 mortality with alternative control variables and various scope restrictions

	Alternative control variables	Only affluent countries	Only democracies	Only countries with > 500,000 inhabitants
Per capita GDP in 1000 USD [HDI 0-100]	3.395	1.009*	.998	1.009
Median age [% of population age 65 and over]	1.013	.953	.892***	1.052+
Democracy and human rights (0-100) [Polity IV score -10 to +10]	1.005	1.004	1.056+	1.005
Air travel exposure (AEF index 0-100)	1.051***	1.038	1.082***	1.031*
Tourist arrivals in millions	1.001	1.009*	1.004	1.000
Island state	.332	.284***	.285***	.325**
Travel restrictions scale (6 items)	.992*	.991**	.991**	.991*
<i>Clustered by world region</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
N=	151	63	76	154
R ²	.037	.051	.065	.039