ESSAYS IN APPLIED MACROECONOMICS

by

Olivér Miklós Rácz

Sumbitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy at
Central European University

Supervisor: Ádám Szeidl

Budapest, Hungary

10.14754/CEU.2022.06

CENTRAL EUROPEAN UNIVERSITY
DEPARTMENT OF ECONOMICS AND BUSINESS

Author:	Olivér Miklós Rácz
Title:	Essays in Applied Macroeconomics
Degree:	Ph.D.
Dated:	September 15, 2022

Hereby I testify that this thesis contains no material accepted for any other degree in any other institution and that it contains no material previously written and/or published by another person except when appropriate acknowledgements are made.

Signature of the author:

Olivér Miklós Rácz

DOI:10.14754/CEU.2022.06

CENTRAL EUROPEAN UNIVERSITY
DEPARTMENT OF ECONOMICS AND BUSINESS

The undersigned hereby certify that they have read and recommend to the Department of Economics and Business for acceptance a thesis entitled **"Essays in Applied Macroeconomics"** by Olivér Rácz.
Dated: September 15, 2022

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
Chair of the Thesis Committee:

DocuSigned by:
A. Weber
____
Andrea Weber

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
Advisor:

DocuSigned by:
Ádám Szeidl
____
Ádám Szeidl

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
Internal Examiner:

DocuSigned by:
Zsófia Bárány
____
Zsófia Bárány

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
External Examiner:

DocuSigned by:
Paul Pichler
____
Paul Pichler

I certify that I have read this dissertation and in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.
External Member:

DocuSigned by:
Judit Krekó
____
Judit Krekó

Abstract

This thesis applies micro-econometric methods to macroeconomic questions. The first two chapters use country-level data to estimate the effect of distancing policy interventions on the effective reproduction number of COVID-19 and a set of economic outcomes. The third chapter identifies the average time it takes for an industry-level shock for getting transmitted to customer industries in the US economy.

Chapter 1

Distancing policies became the primary preventive intervention during the COVID-19 pandemic. This paper estimates the effect of such interventions on the effective reproduction number (R_t) of this virus on a daily panel of 109 countries. Distancing interventions affect COVID infections indirectly through the regulation of social behaviors, which are also a function of voluntary decisions. The main contribution of this paper is the separation of policy-compliant and voluntary distancing effects. I identify the policy-compliant component of distancing behavior as rapid changes in social activity immediately after an intervention. This allows me to isolate the voluntary component as residual changes in activity. I use the isolated voluntary component as a control in the main estimation of distancing policy effects on R_t. I distinguish between (i) place restrictions: restricting destinations and (ii) mobility restrictions: regulations on inland movements. I find strong and permanent effects for both types of restrictions. Place restrictions that target specific destinations are found to be less effective than general mobility restrictions. The effect of voluntary distancing is also significantly negative but weaker than that of policy restrictions. These results suggest that governments can use distancing restrictions effectively in pushing the effective reproduction number below the containment threshold: $R_t = 1$.

Chapter 2

Distancing policy interventions (DPIs) were aimed at containing the COVID-19 pandemic, but they also had severe effects on economic activity. This

paper estimates the effects of DPIs on selected indicators of monthly economic activity, such as industrial and manufacturing production, construction output, retail trade, inflation, and unemployment. The main contribution of this paper is the isolation of the causal effects of distancing interventions from the effects of voluntary distancing. I use mobility data as a measure of distancing to identify DPI effects as immediate changes in distancing shortly after the first intervention. I then use residual changes in distancing as a control in the estimation of the economic effects of DPIs. I find significant output losses due to DPIs, but no evidence for inflationary or unemployment effects. Results also show that although voluntary distancing caused significant output losses, their effect was an order of magnitude smaller than that of DPIs.

Chapter 3

While there is growing evidence for the network origins of aggregate volatility, this paper investigates the potential for the network origins of aggregate dynamics. This paper builds on the predictions of the production network model of Long and Plosser (1983). In this model, the average propagation time – the time a shock needs to get transmitted between producers – is undefined in calendar units. This paper identifies the average propagation time of the US economy using annual data series of 66 industries. I find that it was between 4 and 8 months in past decades. That means the effect of a TFP shock propagates through the production network beyond the time horizon of a year, generating auto-correlated economic aggregates even without the help of auto-correlated shocks. This finding provides evidence for the network origins of aggregate dynamics at annual frequencies.

Acknowledgements

I am indebted to my advisor, Ádám Szeidl for the constant support and guidance he offered me throughout my studies. I strongly believe that the discussions with him shaped my way of thinking well beyond the scope of this thesis.

I am very grateful to the examiners of my thesis, Paul Pichler, Zsófia Bárány and Judit Krekó who gave lots of constructive comments that helped me improve the quality of my papers substantially.

My work was hugely facilitated by the Empirical Research seminars invented and organized by Miklós Koren, Arieda Muço, and Andrea Weber. This seminar provided weekly opportunities for discussing the development of my chapters.

I am also grateful to László Mátyás, and Róbert Lieli who gave me substantial insights when discussing the empirical strategies of my chapters. They were always very encouraging and supportive and their doors were open for me to share my ideas or concerns.

I spent a semester at the Faculty of Economics at the University of Cambridge as a visiting student. I am grateful for their hospitality and for the inspiring research environment, I could be part of there.

I am incredibly thankful to Vasco Carvalho, who was my hosting advisor at the University of Cambridge. Our discussions were very encouraging and gave a lot of insights into the theoretical side of the third chapter of this thesis.

From my fellow Ph.D. students, I want to thank Ákos Aczél, Anna Bárdits, Luca Drucker, Attila Gáspár, Boldizsár Juhász, Balázs Kertész, Olivér Kiss, Judit Krekó, Balázs Krusper, Lili Márk, Gábor Révész, and Lajos Szabó for all the ideas, discussions and friendship we could share together. I will miss those lunches at the CEU canteen from the pre-covid times.

I am grateful to Márta Jombach, Lilla Nagy, Veronika Orosz, and Katalin Szimler, staff members at the Department of Economics and Business, who have always been very helpful.

I am also grateful for the support and friendship I was able to enjoy from my colleagues during my employment at the Central Bank of Hungary.

10.14754/CEU.2022.06

I am especially grateful for the support I received from my boss, Balázs Világi, during the years I worked for him before submitting this thesis.

I am thankful to Attila Madaras and Oscar Wahltinez at Google for their selfless and always quick help in accessing and navigating the databases I used extensively in the writing of this thesis.

It would have been impossible to write this thesis without the support of my family. I am very grateful to my parents and my girlfriend who stood by my side any time I needed encouragement.

Contents

10.14754/CEU.2022.06

List of Tables

10.14754/CEU.2022.06

List of Figures

Chapter 1

The Effect of Distancing Policies on the Reproduction Number of COVID-19

1.1 Introduction

Distancing policies, such as school closures, gathering limits, or stay-at-home orders were the primary preventive interventions in essentially all countries during the COVID-19 pandemic of 2020–2021. The logic behind these policies is to reduce the chances of already infected people infecting others. The number of new infections an infected person is expected to cause during her illness is the effective reproduction number, R_t.

The main objective of this paper is to quantify the effect of distancing policies on the effective reproduction number of COVID-19.[1] The main contribution of this paper is the separation of policy-induced and voluntary distancing effects. I identify the policy-induced component of distancing behavior as rapid changes in social activity immediately after an intervention. This allows me to isolate the voluntary component as residual changes in activity. I use this isolated voluntary component as a control in the main estimation of distancing policy effects on R_t. Because holding voluntary distancing effects fixed allows for the identification of unbiased policy effects in a comparison of countries with different policy interventions.

[1] COVID-19, officially known as SARS-CoV-2, is a virus spread by human droplets like the regular flu. It has a higher basic reproduction number and mortality rate than the regular flu, according to Petersen et al. (2020). Neither vaccines nor designated medical treatments were available until the end of 2020.

In Section 2 I start with the description of the data. I use four datasets: (i) daily preventive policy interventions from Hale et al. (2020), (ii) reported COVID cases, deaths, recoveries, and (iii) Google's publicly available mobility reports from Wahltinez et al. (2020), and (iv) Google's COVID-19 Aggregated Mobility Research Dataset, which is available with permission from Google. I build a daily frequency cross-country panel database covering 109 countries and spanning calendar days between February 2020 and April 2021. The population of the countries involved in the sample is 5.4 billion, representing 70 percent of the world's population in 2020.

After the data description, I define the most important variables in this paper: distancing policies, reproduction numbers, social activity, and imported cases. This study focuses on the effects of two different distancing policy types: *place* and *mobility* restrictions. A place restriction targets specific destinations or events where people are not allowed to go. These are school and workplace closures; cancellations of public events; and gathering limits. A mobility restriction controls how and when people are allowed to move around within their countries, regardless of their destination. These are restrictions on public transportation, stay-at-home orders, and within-country travel restrictions.

Because of its policy relevance, I chose the effective reproduction number R_t as the outcome variable of this study. All preventive measures aim to achieve $R_t \leq 1$, which defines the containment of an epidemic. Knowing the effects of distancing interventions in terms of R_t is therefore useful information for decision-makers. I proxy R_t by the instantaneous reproduction number R_t^I. The advantage of using R_t^I is that it is much easier to calculate and proportional to R_t. Therefore, any proportional effects measured on R_t^I can be interpreted as effects on R_t.

Social activity proxies distancing behaviors. It is an indicator derived from Google mobility indicators, which measure the frequency of Google users in public spaces relative to pre-COVID levels. I use this indicator to isolate its voluntary component, which is the most important control in the main estimation. Finally, imported cases are proxied by an indicator that I created using the proprietary Google COVID-19 Aggregated Mobility Research Dataset[2].

[2]This dataset is only available with permission from Google LLC.

10.14754/CEU.2022.06

In Section 3, I present my empirical strategy. I carry out my estimation in a two-stage design. The first stage is the separation of the voluntary and policy-compliant components of social activity. The second stage is the main estimation of distancing policy effects on the effective reproduction number. In the first stage, I identify the policy-compliant component of distancing behavior as rapid changes in social activity immediately after an intervention. I isolate the voluntary component as residual changes in activity. This allows me to identify policy-compliant and voluntary distancing effects separately in the second stage by using this isolated voluntary activity component as a control variable. In the second stage, the effects of distancing policies are identified from a comparison of countries that have introduced a particular restriction to those that have not, holding voluntary activity, other preventive policies, and covariates fixed.

At the end of this Section, I discuss possible threats to the identification. A policy intervention can work as a signal, inducing voluntary distancing. This kind of voluntary distancing does not harm identification because it is a direct consequence of the interventions.[3] Countries differ in demographics, population density, and the quality of political and healthcare institutions, which are likely to correlate with interventions, social activity, and reproduction numbers. I address these differences by including country-fixed effects in both stages, assuming the invariability of these factors on daily frequencies. Countries also differ in the timing of their interventions, which is addressed by the inclusion of time-fixed effects.

Different countries provided different levels of economic support, which might have worked as incentives to leave workplaces for sick people. I address these differences by controlling for all available information on economic support. I control for daily weather conditions to address the effects of the climate on the reproduction numbers of the virus. Finally, I control for weekly seasonality in both stages of my design.

I present all results in Section 4. I find that place restrictions reduce R_t by 29 percent and mobility restrictions by 61 percent on average. These are strong effects on the reduction of the effective reproduction number,

[3] It has to be noted, though, that this kind of induced voluntary distancing is also accounted for in policy effects in this study.

suggesting that distancing policies were an effective tool for reducing the impact of the pandemic. Place restrictions that target specific destinations are found to be less effective than general mobility restrictions. A one standard deviation drop in voluntary social activity is found to decrease R_t by 17 percent. The effect of voluntary distancing is also significantly negative but weaker than that of policy restrictions. Based on these results, I calculate the contribution of distancing policies and voluntary distancing to the average decline of R_t observed in the first wave. I find that distancing policies contributed 6.5 times more than voluntary distancing to the decline in reproduction numbers.

These findings suggest that although voluntary distancing behaviors help to slow down the reproduction of the virus, any kind of distancing policy measures are much more effective in stopping a pandemic. In the second part of Section 4, I investigate heterogeneous policy effects. The first of these exercises analyse the strength of the policy effects on different time horizons. I am interested in how long the effects identified in the main design last. I do that because it is useful to know how long a government can rely on a place or a mobility restriction. To do that, I modified my second stage design into an event study design, allowing for heterogeneous effects on different time horizons. I find similarly strong effects on shorter and longer horizons for both restriction types. These results suggest that governments can rely on these distancing restrictions on longer horizons when fighting longer waves of infections.

In the second exercise, I break down the larger restriction categories into their components: place and mobility restrictions. I also allow for heterogeneity in the different stringency levels of these policies. I do this to provide comparative results for more delicate policy interventions. I found that school and workplace closures, gathering limits, and stay-at-home orders were effective restrictions in the reduction of reproduction numbers. I cannot find supporting evidence, however, for the effectiveness of the cancellation of public events, restrictions on public transportation, and inland travel restrictions.

School closures are found to be effective only if they are mandated. Workplace closures are found to be effective already when they were only a recommendation. Their efficiency only marginally increases with stringency. Gathering limits become effective at the 100+ limit and gain effectiveness at more restrictive limits. Stay-home orders are found to be effective when they are just recommended. They also gain effectiveness as they become more stringent. Overall, these findings suggest that there was heterogeneity between the effectiveness of different policies, implying that different policy mixes could have led to very different outcomes.

Conclusions are discussed in the final section of this study. Based on my results, I conclude that governments can use distancing restrictions effectively to push the effective reproduction number below the containment threshold of $R_t \leq 1$. They can rely on these effects for as long as these measures are in place. Considering the heterogeneous effects of particular distancing policies suggests that a careful selection of these policies and their stringency levels is recommended before their implementation.

Literature

This paper belongs to the empirical evaluation of non-pharmaceutical interventions (NPI) during the COVID-19 pandemic, surveyed exhaustively by Perra (2021). This literature already provides strong qualitative evidence for the effectiveness of NPIs. The quantitative comparison of these papers is difficult, however, because of the high variability in the chosen outcomes and treatments.

Within this literature, this paper is a contribution to cross-regional studies. These studies encompass a set of countries or states within a federation such as the US or Germany. Islam et al. (2020) study the effect of five physical distancing interventions on a sample of 149 countries and regions on estimated incidence rate ratios. They found that any physical distancing intervention reduced COVID-19 incidence by 13%. This finding is qualitatively in line with the findings of this study, as I also find significantly negative effects of distancing policies on case reproduction. It is much more difficult to contrast these results quantitatively because the outcome variable chosen for this study is new incidence per total number of active infections. Askitas et al. (2021) estimates the effect of different NPIs non-parametrically in an event study design controlling for overlapping interventions. They found that closing schools and workplaces had significant effects on reducing COVID-19 infections, while later installed restrictions on inland travel and public transport had no effects. When comparing different NPIs I find that school and workplace closures were much more effective in the reduction of the reproduction number than restrictions on inland travel and public transportation. As lower reproduction implies lower incidence, these findings are in line. This paper considers other NPIs as well, finding that stay-at-home orders and gathering limits set at 100+ people are found to be similarly effective to school and workplace closures. Ullah and Ajala (2020) contrasts the effects of distancing measures on testing policies on a very similar sample. They find that a unit change in their lockdown index decreases the total number of confirmed cases by 0.19 percent, which becomes significant after 7 days of its implementation and stays intact even after 21 days. This study takes into account testing policies, but the outcome variable is so different I dispense a comparison.

There are papers that choose the effective reproduction number as their outcome variable, just like this paper. Haug et al. (2020) rank 46 different NPIs by their impact on R_t on a sample of 79 territories. Overall they find that less stringent NPIs are just as effective as more drastic ones. They find that the most effective NPI is a small gathering limit, which reduces R_t by about 9 % on average.[4] They found the impact of school closure on R_t at about 7.5%. These results are about 1/3 of the effects found in this study. They find weaker, but significantly negative effects for individual movement restrictions, lockdowns. These findings are qualitatively comparable to stay-at-home orders of this study. They evaluate many other NPIs that are not directly comparable NPIs studied in this paper.

Koh et al. (2020) confirms that all forms of lockdown interventions effectively reduce average R_t regardless of stringency levels, adding that earlier implementations are associated with stronger results. They discover that, depending on the timing of the intervention, the gathering limits reduce R_t by 15 to 41%. This interval contains the results found in

[4]They report their main results in absolute reductions in R_t, whereas this study estimates percentage reductions; thus, their results can be directly compared to those found in this study by assuming some basic reproduction number, R_0. Liu et al. (2020) estimates COVID-19's basic reproduction number to be between 3 and 5. I translate their findings on absolute reductions to percentage reductions by taking the middle point of this range at 4.

this paper for the effect of gathering limits.[5] They find that "lock-down-type" measures to reduce R_t between 14 and 44 %. These numbers are almost the same in size as the findings of this paper: the slackest stay-at-home order is found to significantly reduce R_t by 18.5 %, while the most restrictive type by 35.6 %. Castex et al. (2021) find that the effectiveness of NPIs is negatively correlated with population density, country surface area, employment rate, and proportion of elderly in the population, and positively correlated with GDP per capita and health expenditure.

There are papers that estimate the effect of NPIs on the mobility of people similarly to the first stage estimation of this study. Gupta et al. (2020a) and Gupta et al. (2020b) are focused on the mobility effects of NPIs, while Castex et al. (2021) and Askitas et al. (2021) use their similar estimations as supporting evidence for their main conclusions.

A common limitation of these works is that they do not address the confoundedness of policy-compliant and voluntary distancing effects. This is where the main contribution of my paper lies relative to this strand of the literature. I address this problem by separating voluntary and policy-compliant distancing behaviors in a first-stage estimation and using the voluntary component as a control in my main specification that estimates the effects of distancing policies on the reproduction number of COVID-19.

The only paper I am aware of that addresses this confoundedness problem is Chernozhukov et al. (2021). They estimate the effect of NPIs on the growth rate of COVID cases and related deaths on a daily panel of US states by instrumenting NPIs and observing distancing behavior with the past history of their outcome variables. They find evidence for both policies and information on transmission risks having a significant influence on COVID-19 cases and deaths and show that policies explain a large fraction of social distancing behaviors. They exploit the homogeneity of past cases and deaths in the separation of voluntary and policy-compliant effects. This study leverages the discontinuity in distancing behaviors after intervention in contrast.

[5]Except for the slackest type of a gathering limit of above 1000 people, which was found to be ineffective in this paper.

Figure 1.1: Geographical Coverage of the Sample

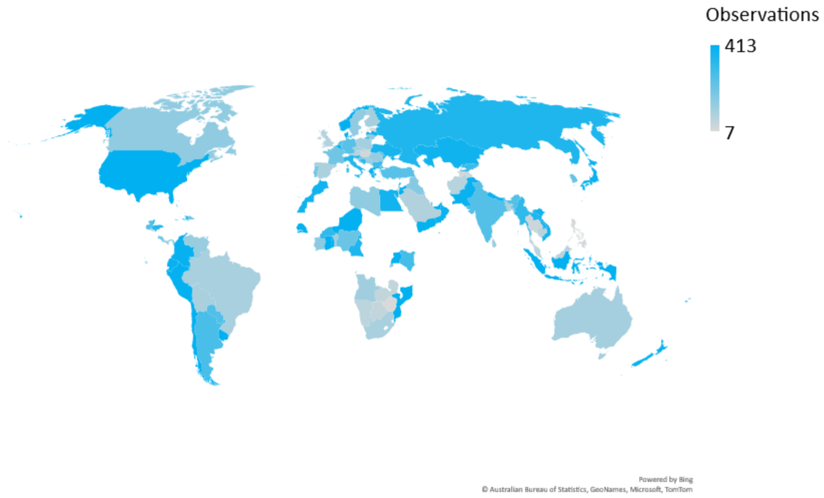

Notes: One observation per country is a daily observation. Countries are colored if included. Brighter colors show more observations.

1.2 Data and Variable Definitions

In this section, I start with a brief description of data sources and the estimation sample. Then I present the definitions of the most important variables of this study: distancing policies, reproduction numbers, social activity, and imported infections.

I use four datasets: (i) daily preventive policy interventions from Hale et al. (2020), (ii) reported COVID cases, deaths, and recoveries, (iii) Google's publicly available mobility reports from Wahltinez et al. (2020), and (iv) Google's COVID-19 Aggregated Mobility Research Dataset, which is available with permission from Google.

I build a country-day panel dataset covering 109 countries and spanning every calendar day between the 15th of February 2020 and the 3rd of April 2021. The sample covers 5.4 billion people, representing 70 percent of the world's population in 2020. Figure 1.1 shows the geographical coverage of the sample on a world map. Countries are colored if they are included. More intensive colors show more observations. The sample includes countries from all populated continents cover most of Europe Australia and both Americas. China is excluded, where the first outbreak preceded the beginning of my sample.

My primary data source is Google's COVID-19 Open-Data platform by

Wahltinez et al. (2020), which is a ”repository attempting to assemble the largest COVID-19 epidemiological database in addition to a powerful set of expansive covariates. It includes open, publicly sourced, licensed data relating to demographics, economy, epidemiology, geography, health, hospitalizations, mobility, government response, weather, and more.” I extend this data by daily country-level reports of the Johns Hopkins University about recoveries from COVID infections.[6] I employ the Google COVID-19 Aggregated Mobility Research Dataset to calculate a proxy of imported COVID infections for each country. This database is available with permission from Google LLC.

1.2.1 Distancing Policies

This study focuses on the effects of distancing interventions implemented during the COVID-19 pandemic. These interventions are collected and reported in a daily regional dataset by Hale et al. (2020) called the Government Response Tracker. They cover all sorts of government interventions related to the COVID-19 pandemic including distancing measures, e.g. gathering limits, other types of preventive policies, e.g. mask wearing mandates, and different kinds of economic support, e.g. debt reliefs.

I form two groups from the seven different distancing interventions and label them as place and mobility restrictions:[7]

- **Place Restriction**: lock-down of schools, workplaces, cancellation of public events, plus gathering limits,
- **Mobility Restriction**: restrictions on public transportation, inland travel restrictions and stay-at-home orders.

The primary reason for this grouping is statistical. Many governments introduced these measures in bundles reducing the likelihood to identify

[6]The reliability of these reports was questioned in the Summer of 2021. Therefore, these figures are no longer reported in the Johns Hopkins dataset. I use these numbers for the calculation of instantaneous reproduction numbers (R_t^I), my primary outcome variable. I provide some country-level validity checks of R_t^I in the Appendix.

[7]My interests are limited to inland restrictions. Therefore, I exclude international travel controls from distancing policies.

the effect of each distancing indicator in isolation. Collecting these measures into groups might allow for more powerful estimates. My grouping is based on the pairwise time distance between the introductions of a pair of policies. Table 1.1 reports the fraction of countries that had introduced a pair of policies within at most seven days and highlights the shares that are greater than 50 or 66,7 percent.[8] The larger fractions concentrate in two different groups which give the basis for my grouping.

Table 1.1: Percent of Countries Implementing a Policy Pair within 7 Days.

	Place Restriction				Mobility Restriction		
	School	Event	Gather	Work	Stay H	Move	Transp't
School Closure		76.15	67.59	71.56	50.00	53.70	39.00
Events Cancelled	76.15		70.37	55.96	44.44	49.07	29.00
Gathering Limit	67.59	70.37		65.74	57.94	58.88	44.00
Workplace Closure	71.56	55.96	65.74		62.96	60.19	54.00
Stay Home Order	50.00	44.44	57.94	62.96		67.29	58.59
Movement Restricted	53.70	49.07	58.88	60.19	67.29		64.65
Public Transport Closed	39.00	29.00	44.00	54.00	58.59	64.65	

Notes: highlight: $\geq$ 50%, strong highlight: $\geq$ 66.7%

These two types of policies have a qualitative difference as well that motivated their labels. Places restrictions are targeted interventions. They define specific locations or events where people are not allowed to go. Mobility restrictions on the other hand control when and how people are allowed to go regardless of where they are headed to.

Governments implemented these distancing orders with different levels of stringency and generality. A school closure can be a recommendation or a strict mandate, and it can cover different levels of education or geographic locations. Hale et al. (2020) define several stringency levels for each distancing intervention and flag if the intervention was country-level or regional.[9] To retain estimation power I use the following defini-

[8]The same grouping can be confirmed by setting different thresholds on day distance. Find similar tables for 3,5 and 9 days in the Appendix.

[9]A state-level intervention is flagged as regional in a federal state such as Germany or the US, which is treated as a single unit in this estimation.

tion for my policy indicators:

$$P^{p}_{it} = \min\left[1, \sum_{j \in \text{type}} D^{j}_{it} F^{j}_{it}\right], p \in \{\text{place}, \text{mobility}\} \tag{1.1}$$

where D^j_{it} is the category variable for distancing policy j, e.g. school closures, which is $D^j_{it} = 0$ if restriction j is not in action in country i on day t, and $D^j_{it} > 0$ codes the level of stringency in country i on day t using consecutive integer values starting from 1. Type can be either place or mobility restrictions. F^j_{it} is a binary indicator of a distancing measure j being a country level order in country i on day t or only regional. This formula defines a binary variable, therefore, for each distancing policy type. $P^{\text{place}_{it}}$ takes the value 1 if there was at least one country-wide place type restriction in action in country i on the day t, and 0 if there was none. $P^{\text{mobility}_{it}}$ defines another binary variable on the same grounds for mobility restrictions.

These definitions of policy indicators have the benefit of a binary treatment: their coefficients are easy to interpret. This advantage, however, comes at a cost: P^{place}_{it} and P^{mobility}_{it} indicate the first-ever countrywide distancing interventions, thus stay blind to later changes in those interventions. They also overlook the cross-country heterogeneity in the stringency and the number of interventions of these first interventions, as they are normalized to 1 from day 0. That means these heterogeneities and later changes are absorbed by other variables or the error terms unless they are controlled for. The current version of this paper lacks this control, which is a serious limitation.

An important limitation of the data sources is that they only provide information about the implementation of distancing restrictions but not on their announcements. There is anecdotal evidence that some of these restrictions were announced earlier in some countries, allowing people to adjust their behaviors beforehand. Panic shopping for basic goods could be a good example of such anticipatory responses. The effects of earlier announcements are discussed in the Results section.

1.2.2 Reproduction Numbers

The effective reproduction number R_t is the chosen outcome of this study. It is the number of new cases a single infection is expected to cause. When $R_t > 1$ the number of infections grows exponentially, which is the definition of an epidemic. But, when $R_t \leq 1$, the growth is linear and the contagion is considered to be contained. It is, therefore, the most useful indicator to judge the efficiency of any preventive interventions: prevention is successful if it is able to push R_t below 1.

R_t can be decomposed the following way:

$$R_t = R_t^I \cdot \mathrm{E}_t[\text{duration of infection}], \tag{1.2}$$

where R_t^I is the number of new infections an infected individual is expected to cause within a day and is usually referred to as the instantaneous reproduction number.[10] Albeit simple this decomposition is useful for two reasons.

First, in contrast to R_t, the calculation of R_t^I from daily COVID incidences is feasible. R_t^I can be calculated by dividing the number of new infections discovered on the day t by the number of known infections from the previous day:

$$R_t^I = \text{New Infections}_t / \text{Infected}_{t-1} \tag{1.3}$$

where New Infections_t are reported, thus can be observed. The number of infected individuals cannot be directly observed, but can be calculated from reported figures: $\text{Infected}_{t-1} = \text{Total Cases}_{t-1} - \text{Total Deaths}_{t-1} - \text{Total Recoveries}_{t-1}$.

Second, this decomposition allows me to identify the effect of distancing policies on R_t even if I use R_t^I as the outcome, because I assume that distancing policies cannot affect $\mathrm{E}_t[\text{duration of infection}]$ only R_t^I. The intuition is that once one has the virus its duration is independent of the frequencies at she meets other people. This strategy also requires me to estimate proportional effects, because of $R_t \propto R_t^I$.

Figure 1.2 shows the evolution of R_t^I around the days of place and mo-

[10] All the formulas presented here are consistent with and can be derived from the commonly used compartment models of epidemics, e.g. SIR models.

bility restrictions smoothed by a seven days backward-looking moving average. A turn in the trend of R_t is apparent on both graphs. R_t is in a decline after both place and mobility restrictions, which suggests a strong effect of distancing policies.

Figure 1.2: Instantaneous Reproduction Numbers Around Distancing Interventions

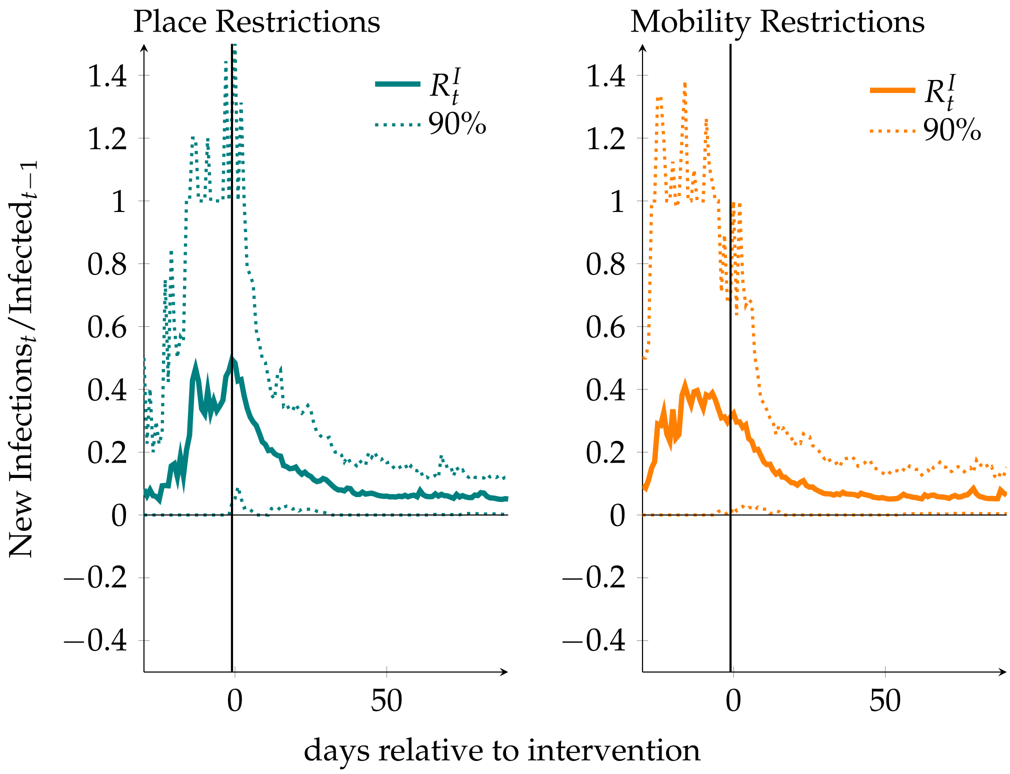

Notes: R_t^I - instantaneous reproduction number, solid line: 7 days backward-looking moving average of cross-country mean of R_{it}^I , dotted lines:7 days backward-looking moving averages of the 5th and 95th percentiles of the cross-country distribution of R_{it}^I.

Looking at the 90 percent boundaries we can see a rather wide distribution of R_t across countries ranging from 0 to 1.5 new infections by a single infected individual every day during her infection. This upper bound is huge considering the expected length of the infection is around 10 days according to Liu et al. (2020), suggesting an effective reproduction number close to 15 in some countries on some days.

1.2.3 Social Activity

The main contribution of this paper is the separation of policy compliance and voluntary distancing effects. To be able to do that I need an indicator that measures overall distancing behaviors. I call this indicator social activity and use the notation a_{it}.

I define a_{it} as the first principal component of Google's six mobility indicators. These are publicly available daily indicators published for countries and sub-regions from February 15, 2021. A mobility indicator is recording differences in the frequency of Google users relative to a five-week period from before the pandemic in a specific location category, which are:

- **groceries**: grocery markets, food warehouses, farmers markets, specialty food shops, drug stores, and pharmacies,
- **retail**: restaurants, cafes, shopping centers, theme parks, museums, libraries, and movie theaters,
- **parks**: local parks, national parks, public beaches, marinas, dog parks, plazas, and public gardens.
- **transit stations**: public transport hubs such as subway, bus, and train stations,
- **workplaces**: places of work,
- **residential**: places of residence.

Table 1.2 shows the results of the principal component analysis. The first principal component captures 86 percent of the total variance of the six mobility measures. It is loaded almost equally by all indicators with the same signs except for residential locations. This pattern parallels the intuition that social distancing resulted in fewer people in public spaces and more people at their homes relative to pre-COVID levels. The rest of the components are all dominated by one or two of the mobility indicators supporting the choice of the first component as my proxy for social activity.

Table 1.2: Principal Components of Google's Mobility Indicators

Component	1st	2nd	3rd	4th	5th	6th
Groceries	0.4029	0.0491	0.9006	-0.1201	-0.0518	0.0837
Retail	0.4295	-0.0846	-0.0387	0.5563	0.2122	-0.6726
Parks	0.3439	0.8916	-0.2172	-0.1330	0.1379	0.0534
Transport Stations	0.4275	-0.1375	-0.1886	0.5227	-0.2263	0.6621
Workplaces	0.4140	-0.3844	-0.2010	-0.4402	0.6478	0.1645
Residential Areas	-0.4252	0.1697	0.2534	0.4375	0.6800	0.2690
Share in Total Variance	0.8596	0.0793	0.0328	0.0151	0.0077	0.0055

1.2.4 Imported Infections

Imported infections are an important control variable of this study. I create a proxy for imported infections using the proprietary Google COVID-19 Aggregated Mobility Research Dataset, which is only available with permission from Google LLC. It contains anatomized mobility flows aggregated over users who have turned on the Location History setting, which is off by default. This is similar to the data used to show how busy certain types of places are in Google Maps — helping identify when a local business tends to be the most crowded. The dataset aggregates the flow of people from region to region, which is here further aggregated at the level of NUTS3 areas, weekly.

First, I keep only the flows that connect cells from different countries.[11] Second, I aggregate these flows by countries and match the epidemiological indicators by departure countries. Third, I take cross-country flows and multiply them by the number of infected individuals per 1000 citizens in departure countries. This yields me the expected flows of COVID infections by source and receiver country pairs.[12] Finally, I aggregate these expected infection flows by the receiver country to get the expected number of imported infections.[13]

This process has some minor limitations, as it is based on google user accounts, Therefore, it might be less accurate or totally missing for un-

[11] I geolocate all cells using Picard (2015).

[12] I shift back infection data by 14 days to account for the presumed delay in epidemiological reports.

[13] It is weekly frequency information. I use the first day of the week as its time index and date back by 6 days. Therefore, it codes the expected inflow of infections into a country in the past calendar week. Then I interpolate missing data points within a week by a quadratic spline developed using the csipolate Stata module developed by Cox (2009).

Figure 1.3: Causality Map

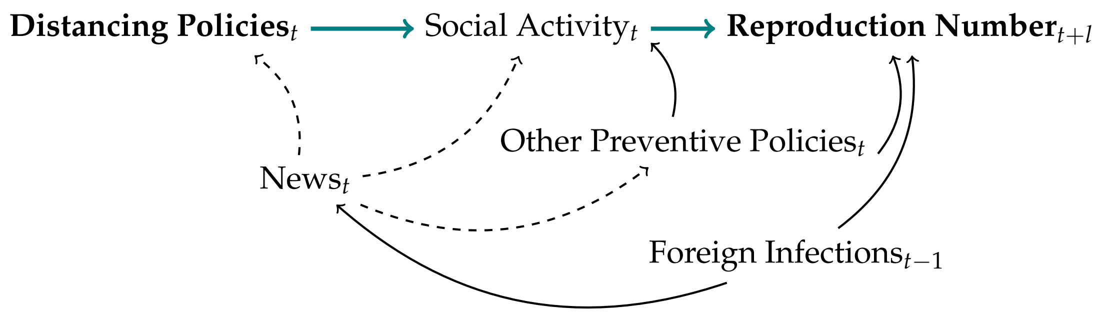

Notes: Arrows point in the direction of causality. Solid line: observed, dashed: unobserved effect. Bold font: focus variables, thick arrows: the path to be identified, Δt is one day.

derdeveloped nations. And flows are missing for some microstates, such as Lichtenstein or Andorra.

1.3 Empirical Strategy

In this section, I develop a two-stage empirical design to identify the effect of distancing policies on the reproduction number of COVID-19. Identification relies on a comparison of countries that have introduced a particular restriction to those that have not. The central question here is: which are those factors that have to be held fixed to make sure this comparison is a valid identification of the effect of that policy? In this Section, I am working towards an empirical design that takes into account all these factors and is feasible to implement.

A good way to start is to map out all the relevant causal links connecting distancing policies to reproduction numbers on a graph. Figure 1.3 shows this causality map, where each arrow shows a causal link pointing in the direction of causality. The main path connecting distancing policies to reproduction numbers is drawn by thick arrows. The first thing that meets the eye is that distancing policies are not connected to reproduction numbers directly. The reason is that a distancing order can only reduce social activity directly. This reduction in social activity is what governments expect to reduce the reproduction number by diminishing the number of new infections.

I added time indexes to explicitly show that these effects can only be observed with a significant delay l on daily frequencies. This delay happens, because it takes time (typically 10-14 days) for an infected individual to start producing symptoms, get tested, and end up reported.[14]

This causality map is not only helpful in showing a clear and comprehensive picture of all relevant causal relations, but it also informs identification. This method is known as the directed acyclic graph (DAG) method and is described in detail in Cunningham (2021). What is sufficient to know about the DAG method here is that any backdoor paths connecting policies with reproduction numbers are signaling possible omitted variable bias.

I recognize three such backdoor paths in this context. The first one connects distancing policies with social activity through news, which is a set containing any bits of information about COVID-19 that has the potential to alter government and individual decisions about distancing.[15] For example, a discovery of a large number of infections raises the probability of a distancing intervention and it can also make people decrease their social activity voluntarily. I will refer to the latter channel as voluntary distancing in this paper onward.

The second backdoor path is the channel of other preventive policies, such as mask-wearing mandates, contact tracing, testing, or vaccination. These policies have an effect on reproduction numbers and their implementation was also likely to be influenced by news. Their effect on reproduction number can be direct, e.g. mask wearing reduces the transmission probability of the virus, while it might also lead to greater distancing according to Seres et al. (2021).

The third path is the channel of imported cases. It is the number of infections in neighboring countries affecting interventions indirectly through news and domestic reproduction numbers directly. For example, if the

[14]This delay mechanism is different for traced contact persons, however, most countries did not choose to do any contact tracing or only tested the contact persons who were showing symptoms. I have information about whether a country is practicing and what kind of contact tracing, which I control for in the second stage. It is also known that a large fraction of COVID cases never gets tested, thus reported, which surely has an effect on the outcome. This effect can be addressed by fixed effects and controls for testing, which are elaborated on in the subsection. All these issues are addressed in Section 1.3.2.

[15]The arrow connecting News to Distancing Policy and Preventive Policy acknowledges the fact of endogenous selection of the treatment of this study: distancing policies. Closing backdoor paths containing this link simultaneously eliminates the endogenous selection bias.

number of new infections shoots up in a neighboring country, that might influence more restrictive policies and also increases the likelihood that new infections will or are already arriving from that neighbor by infected travelers.

Fortunately, all backdoor paths go through news, it would be sufficient therefore to control only for news to eliminate the omitted variable biases caused by them. That means that comparing countries with the same news components but different policies identifies the effect of those policies.[16] This observation is captured by the following design:

$$R_{i,t+l} = \beta P_{it} + \eta' \mathbf{N}_{it} + \mu + \varepsilon_{it}, \tag{1.4}$$

where i indicates a country, t a day. R is the reproduction number, P is the distancing policy, and $\mathbf{N}$ is a set of news components containing reported infections and deaths from $t-1$. Assuming that a distancing policy is a binary treatment, this design identifies β by comparing R in countries that have introduced policy P to those that have not, but were otherwise identical in all components of $\mathbf{N}$, i.e. received the same news.

The simple design in equation (1.4) is not feasible, however, because $\mathbf{N}$ is not completely observable. For example, I have no information about local media influencers or politicians informing the public about COVID developments. Neither about country-specific behavioral reactions to the news, such as compliance with government policies. These factors are also correlated with distancing policies and virus reproduction. I address this problem by closing the three backdoor paths separately.[17] Controlling for the channels of other preventive measures and foreign infections is simple because these are observable factors. Closing the voluntary distancing backdoor path is challenging because I can only observe social activity a_{it}, which pools policy compliant and voluntary distancing motives.

[16] Holding news fixed implies that the effects of voluntary distancing, other preventive interventions, and imported cases are the same.

[17] Alternatively, I could close all three backdoor paths by modeling the endogenous selection of policies based on Heckman and Sedlacek (1985). That would require the credible exclusion of exogenous variables, for which the share of distancing policies in neighboring countries up today $t-1$ could be a valid candidate. The exploration of this possibility however is beyond the limits of this paper. Another alternative approach is of Chernozhukov et al. (2021), which uses the observable components of $\mathbf{N}$ as an instrument for both policies and social activity.

1.3.1 Voluntary Distancing

Here I present the first stage of my estimation, which identifies the policy-compliant component of social activity a_{it} in a regression discontinuity in time (RDiT) design. The voluntary component, called voluntary activity v_{it} is then defined as the residual of the first stage regression.[18]

The effects of a distancing intervention are identified as sudden changes in a_{it} after the intervention. The key identifying assumption is that changes in social activity due to voluntary distancing are slow, while the response to a distancing intervention is quick, on daily frequencies. Distancing interventions, e.g. lockdowns, prescribe a coordinated and sudden reduction in social activities after an intervention. Voluntary distancing responses on the other hand are likely to be much less coordinated considering the heterogeneous attitudes towards COVID infection risks, e.g. virus skeptics and overly cautious people. Changes in social activity due to voluntary motives are presumably much smoother and slower, therefore, on daily frequencies when aggregated to the level of a country.

Figure 1.4 provides visual motivation for the RDiT strategy. It shows the deviation of the social activity indicator from its within-country pre-intervention mean and normalized by the full sample standard deviation in the close neighborhood for the two types of distancing interventions defined in Section 1.2.1. Darker regions show more observations.

It looks like both types of policies reduced social activity by between 1 and 3 standard deviations in most countries just within 10 days. These changes seem to be more rapid in the case of mobility restrictions. It is also apparent that social activity remained constant in most countries before the interventions. Overall the rapid drop and negligible pre-trends observed on in social activity in the close neighborhood of distancing interventions supports the identification strategy of the first stage estimation.

[18] RDiT is described in detail in Hausman and Rapson (2018). A regular RD exploits a discontinuous change in the close neighborhood of a border separating the treated and untreated samples. RDiT is a special case when the running variable is time, which is usually a discrete variable in empirical exercises. This discreteness allows us to identify the effect by event time dummies rather than a discontinuity in a continuous polynomial like in regular RD designs. This design is related to event study designs, but it lacks a control group.

Figure 1.4: Social Activity in the Neighborhood of Distancing Policy Interventions

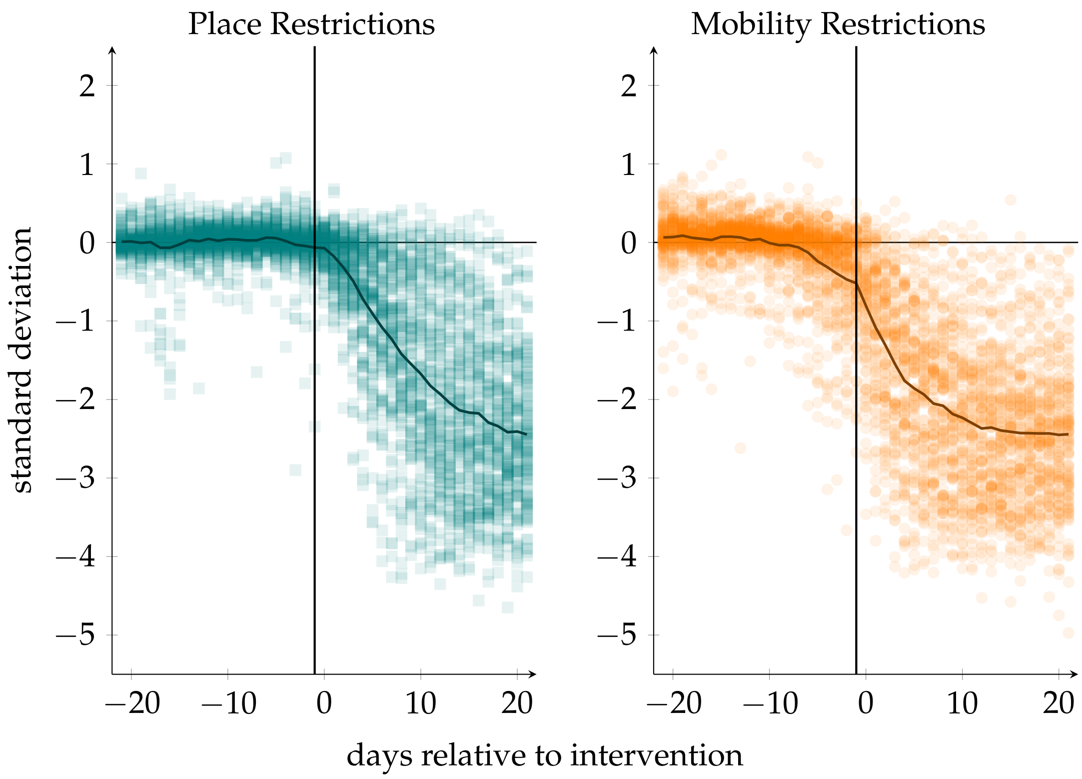

Notes: cloud: country-day observations of social activity a_{it} in the neighborhood of distancing interventions, darker regions show overlapping observations. Solid line: within day averages. Left: place restrictions, right: mobility restrictions. All figures are cleaned from their within-country pre-intervention means and normalized by the full sample standard deviation of social activity.

1.3.2 Second Stage

The second stage estimates the effect of distancing policies P^p_{it}, of type $p \in \{\text{place}, \text{mobility}\}$ on the reproduction number R_{it}. Identification is based on the comparison of countries that have introduced a distancing policy $P^p_{it} = 1$ to those that have not $P^p_{it} = 0$, holding voluntary activity v_{it}, other preventive policies, imported infections, and other covariates fixed.

1.3.3 Threaths to Identification

In this subsection, I review the possible threats to the identification. The effect of distancing interventions is conveyed by two channels: policy-compliant and policy-induced voluntary distancing. People might increase their distancing after the implementation of a restriction because of compliance, but might also because they perceive it as a signal of a worsening epidemic. This kind of voluntary distancing does not harm the separation of unconditional voluntary distancing effects, because it is a direct consequence of the interventions.

Countries differ in demographics, population density, and the quality of political and healthcare institutions, which are likely to correlate with interventions, social activity, and reproduction numbers. I address these differences by including country-fixed effects in both stages assuming the invariance of these factors on daily frequencies. Countries also differ in the timing of their interventions, which is addressed by the inclusion of time-fixed effects absorbing a common trend that tracks the days after the first reported infection within a country. Different countries provided different levels of economic support, which might work as incentives to leave workplaces for sick people. I address this by controlling for all available information on economic support.

I control for daily weather conditions to address the effects of the climate on the reproduction rate of the virus. Finally, I control for weekly seasonality in both stages of my estimations.

1.4 Results

In the first part of this Section, I specify the empirical designs and present estimation results. I start the first stage. From the results of that, I calculate and analyze voluntary activity v_{it}. I then continue with the second stage of estimation. In the second part of this Section, I investigate heterogeneous distancing policy effects by a simple modification to the second stage design.

1.4.1 First Stage

In the first stage I model social activity a_{it} as a function of event time indicators $\delta^{p}_{t-d(i,p)}$ centered around the last day before a distancing intervention $d(i,p)$ of type $p \in \{\text{place}, \text{mobility}\}$ in each country i:

$$a_{it} = \delta^{\text{place}}_{t-d(i,\text{place})} + \delta^{\text{mobility}}_{t-d(i,\text{mobility})} + \zeta' X_{it} + \mu_i + \gamma_t + \nu_{it}, \tag{1.5}$$

where X_{it} are different covariates. The first components of X_{it} are observable news components covering four sets of variables. The first set contain reported domestic COVID cases and COVID-related deaths per population from 1, 2-7, and 8-14 days before. The second set includes the average of the same per capita reports with the same time lags in neighboring countries. The third set is the share of neighboring countries that had already implemented a place or a mobility restriction in the past 1, 2-7 or 8-14 days.[19] The final set are two indicators indicating if there was ad hoc public urging or an organized public information campaign about COVID-19 in place on day t. Table A.6 in Appendix A.5 shows the estimation results for this set of variables.

X_{it} includes also other preventive policies, such as the level of international travel controls, testing policies, quantities and share of positive tests, level of contact tracing, debt reliefs, fiscal aids, and if there were income supports as an incentive for staying home when someone was sick, mask-wearing mandates, vaccination share, and different indicators of daily weather conditions (temperature, rainfall, snowfall, dewpoint, humidity) plus weekly seasonality.

[19]Neighbors are defined by land borders.

I allow for country-fixed effects μ_i to capture country fixed (e.g. cultural, demographic) differences that possibly affect social activities. I include time-fixed effects γ_t setting $t = 0$ to the day the first COVID case was reported in a country to absorb a global trend of distancing response to the evolution of the epidemic that was common across countries.

Event time indicators $\delta^{p}_{t-d(i,p)}$ are included to capture the common trend in social activity around the days of a type p intervention. Because they are intended to capture the effects of the intervention relative to the previous day, δ^{p}_{-1} is omitted for both policy types, as δ^{p}_{0} represents day 0 of intervention. Figure 1.5 show the estimation results for the event time coefficients $\delta^{\text{place}}_{t-d(i,\text{place})}$ and $\delta^{\text{mobility}}_{t-d(i,\text{mobility})}$ of the first stage equation (1.5). Circles represent the point estimates in case of place, and squares for mobility restrictions. Both sets of estimates are graphed with 99 percent error bands. I pool periods more than one week distant from the intervention into three categories, i.e. $\delta^{p}_{t-d(i,p)}$ is a single dummy if $t \in (-\infty, -8]$, or $[7, 20]$, or $[21, \infty)$, keeping the focus in the close neighborhood of the intervention.

It is apparent that social activity a_{it} decreases significantly in the first seven days of distancing interventions, while there are only weak and marginally significant trends in a_{it} preceding the interventions.[20] The rapid response after interventions and negligible pre-trends before is consistent with the main identification assumption of the first stage estimation, i.e. interventions caused sudden changes in distancing behaviors.

A place restriction reduces activity by almost half, and a mobility restriction by close to one standard deviation on day 6. It stays low on longer horizons suggesting a long-lasting effect of both policies. Both restriction types decrease social activity by roughly 1 standard deviation after one week. The effects of both restrictions are gradual in the first seven days. People seem to react to a mobility restriction already on day 0, while significant responses to a place restriction come with a roughly 4 days delay.

[20]One possible explanation for pre-trends is the anticipatory effects of earlier announced restrictions. These early announcements are not observed in Hale et al. (2020) only on the day of implementation for each distancing restriction.

Figure 1.5: Effects of a Place (●) and a Mobility (■) Restriction on a_{it}

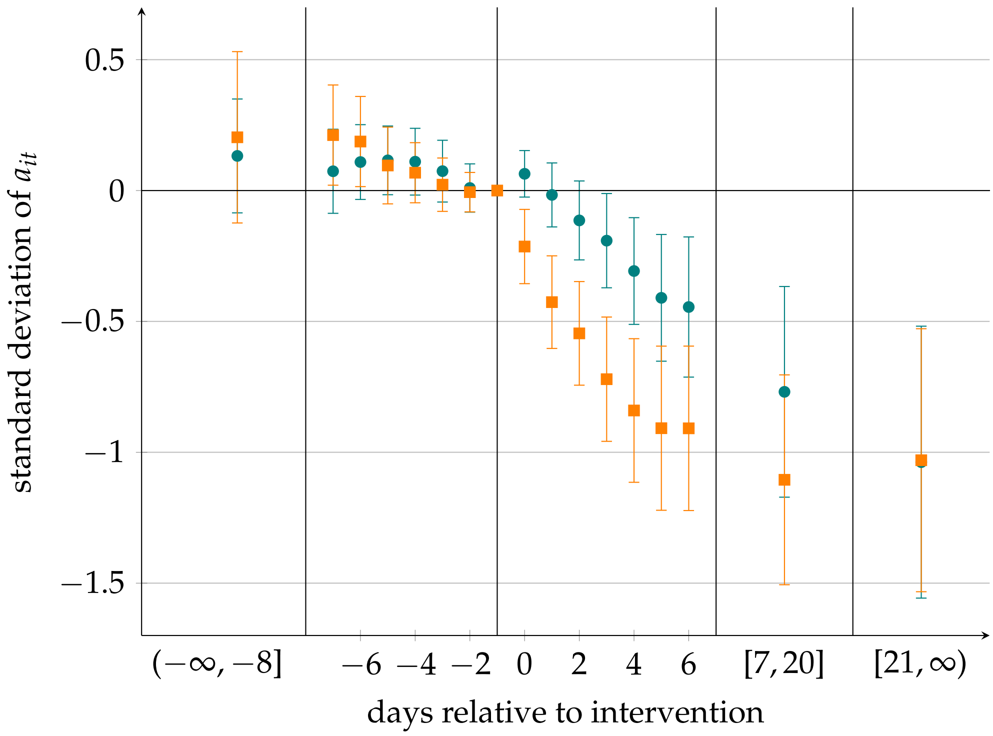

Notes: a_{it} - social activity. Point estimates of $\delta^{\text{place}}_{t-d(i,\text{place})}$ and $\delta^{\text{mobility}}_{t-d(i,\text{mobility})}$ coefficients of equation (1.5) with 99% confidence intervals. Standard errors are allowed to cluster within countries. Reference period: last day before the intervention. 50,070 daily observations within 120 countries.

1.4.2 Voluntary Activity

The goal of the first stage estimation is the isolation of the voluntary component of social activity. I use the results of the first stage estimation, therefore, to break down social activity a_{it} into three components: the effect of distancing policies $\widehat{p}_{it}^{D}$, the effect of other policies $\widehat{p}_{it}^{O}$, and voluntary activity $\widehat{v}_{it}$. Distancing policy effects are defined as changes in a_{it} from day 0 to 6 after intervention and fixed for later days as the effect on day 6. Consistently with the identifying assumption that changes only shortly after an intervention are attributed to that intervention. It is set the same way for both place and mobility restrictions.

The effect of other preventive policies is defined as changes in a_{it} due to other policies. Voluntary activity $\widehat{v}_{it}$ is then defined as residual changes in a_{it} that are not attributed to either distancing or any other policy interventions. These definitions are summarized in the following equations:

$$\widehat{p}_{it}^{D} = \sum_{p\in\{\text{place,mobility}\}} \left[\sum_{j=0}^{6} \widehat{\delta}_j^p + \widehat{\delta}_6^p \mathrm{I}_{t-d(i,p)>6}^p\right] \tag{1.6}$$

$$\widehat{p}_{it}^{O} = \theta' P_{it}^{O} \tag{1.7}$$

$$\widehat{v}_{it} = a_{it} - \widehat{p}_{it}^{D} - \widehat{p}_{it}^{O} \tag{1.8}$$

where P_{it}^{O} is a set of binary indicators indicating if a particular policy with a specific stringency level is in place in country i on the day t. By this definition, in equation (1.8) the error term of the first stage estimation of equation (1.5) is attributed to voluntary activity. This way it might be a more powerful control, than a standard instrumental variable, because that error term contains the effects of all the unobserved factors that might induce changes in voluntary distancing behaviors.[21]

Figure 1.6 shows cross-country averages of social activity with its three components. The time axis is adjusted, such that day 0 is the day when the first COVID case was reported within a country. The solid line is a social activity. Voluntary activity is pictured by a dashed line, the effect of other policies by the dotted dashed line, and the effect of distancing policies by a dotted line. The social activity started to drop soon after day 0 and leveled out roughly on day 20 approximately 2 standard devotions below its pre-COVID levels.

[21]One could interpret $\widehat{v}_{it}$ as counterfactual social activity a_{it} of a no-intervention scenario.

Figure 1.6: Decomposition of Social Activity at the beginning of the Pandemic

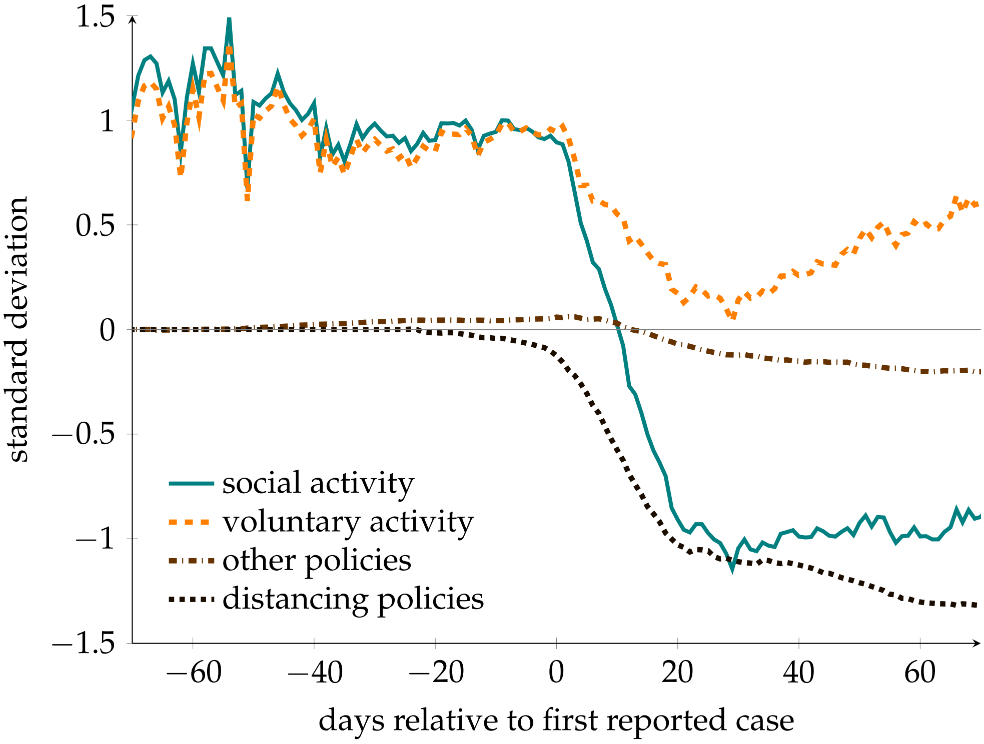

Notes: cross-country averages of social activity and its three components.

All three components are in a decline in the same period between days 0 and 20. Distancing policies dropped the most, with more than one standard deviation, and it kept on declining in the following days. Voluntary distancing dropped almost on standard deviations as well, but it started to rise again after day 30. Other preventive policies had a much smaller effect on social activity.

This decomposition suggests that distancing interventions and voluntary distancing both had a major role in the global reduction of social activities. On longer horizons policies seems to have a more prolonged effect, while voluntary distancing was less permanent. This suggests that distancing interventions might have a more important role in the containment of the epidemic. Investigating this possibility is the primary purpose of the second stage estimation, which is presented in the next Section.

1.4.3 The Effect of Distancing Policies on the Effective Reproduction Number of COVID-19

In this subsection I develop the empirical design of the second stage first, then I analyze its results. The second stage of estimation aims to identify the effect of distancing policies on the reproduction number of COVID-19 controlling for the effects of voluntary distancing, other preventive policies, and covariates. I model the effective reproduction number $R_{i,t+h}$ as a function of distancing policies, voluntary activity $\widehat{v}_{it}$, other preventive policies and covariates:

$$R_{i,t+l} = \exp\left[\beta_p P^{\text{place}}_{i,t-4} + \beta_m P^{\text{mobility}}_{it} + \beta_v \widehat{v}_{it} + \xi' X_{it} + \mu_i + \kappa_t\right] + \varepsilon_{i,t+l}, \tag{1.9}$$

where $R_{i,t+l}$ is proxied by the instantaneous reproduction number in country i observed on day $t + l$. P^p_{it} is an indicator of a distancing policy intervention being 1 on days when any components of that policy type were in action in the country i. Based on the first stage results in Section 1.4.1 place restrictions started to affect social activity after 4 days. Therefore, I include place restrictions in the second stage with a four-day delay: $P^{\text{place}}_{i,t-4}$.

By setting the functional form to exponential, this model is a Poisson regression identifying proportional effects. Identifying proportional effects allows me to use the easily calculable instantaneous reproduction number R^I_t instead of R_t and get equivalent results because $R^I_t \propto R_t$ as it has been shown in Section 2. It is less restrictive than a log transformation because it allows for zero observations in the outcome. This is useful because R^I_t is zero each day when there are 0 new infections are reported.[22] I use $\widehat{v}_{it}$ that resulted from the first stage estimation as a control to eliminate the effect of voluntary distancing.

I allow for country fixed effects μ_i to capture time-invariant differences among countries, e.g. population density, demographics, and the quality of the healthcare system, which possibly affect the reproduction number of the virus. I include also a time-fixed effect κ_t setting $t = 0$ to the day the first COVID case was reported in a country to absorb a

[22]This is a Poisson model on non-integer outcomes. For details see for example Silva and Tenreyro (2006).

global trend in the evolution of reproduction numbers that were common across countries.

Covariates X_{it} include other preventive policies, addressing the third backdoor path. These preventive policies are the level of international travel controls, type of testing policies, testing quantities and share of positive tests, level of contact tracing, debt reliefs, fiscal aids, and if there are income supports as an incentive for staying home when someone is sick, mask-wearing mandates, vaccination share. X_{it} contains also different indicators of daily weather conditions (temperature, rainfall, snowfall, dewpoint, humidity) to capture the patterns of infections in different weathers. It contain also controls for weekly seasonality. Finally, I include the expected number of imported infections and their interaction with international travel controls into X_{it} to control for the channel of imported infections.

Table 1.3 shows the main results of this paper. These are the results of different specifications of the second stage equation (1.9) setting the latency parameter $l = 10$ days.[23] The table starts with the most basic specification that includes only distancing policy interventions besides controls for weather conditions, weekly seasonality, and country and time-fixed effects. The next columns add important omitted factors: voluntary activity, other preventive policies, and imported cases one by one. This way one can judge the relevance of these omitted factors by comparing the point estimates for place and mobility restrictions across columns.

Column (1) shows strong correlations for both policy types with $R_{i,t+l}$. Based on Column (1) place restrictions reduce the reproduction number by 33 percent, while a mobility restriction by 74 percent. This is a misspecified specification however, only included as a benchmark for the better-specified models that control for different sources of omitted variable biases: voluntary distancing, other preventive policies, and imported cases.

In the second column, I add the voluntary activity indicator that has been isolated in the first stage. Controlling for this factor reduces the coefficients of both interventions substantially. This finding confirms

[23]A sensitivity analysis of l can be found in the Appendix. Results show little sensitivity to the choices of $l \in \{7, 9, 11, 13\}$

Table 1.3: Effect of Distancing Policies on Reproduction Numbers 10 days later.

	(1)	(2)	(3)	(4)
Place Restrictions $_{t-4}$	-0.415***	-0.331**	-0.287**	-0.287**
	(0.155)	(0.142)	(0.142)	(0.139)
Mobility Restrictions $_{t}$	-0.737***	-0.690***	-0.621***	-0.610***
	(0.163)	(0.139)	(0.123)	(0.118)
Voluntary Activity $_{t}$		0.154***	0.165***	0.167***
		(0.041)	(0.040)	(0.040)
Imported cases $_{t}$				0.310***
				(0.105)
Import × Screening $_{t}$				2.067
				(9.768)
Import × Quarantine $_{t}$				-1.952*
				(1.100)
Import × Selective Ban $_{t}$				1.080
				(2.119)
Import × Total Ban $_{t}$				-1.670***
				(0.553)
Observations	26,566	26,566	26,566	26,566
Countries	109	109	109	109
Other Preventiv Pol's	○	○	●	●
Country and Day FE's	●	●	●	●

Notes: *** p<0.01, ** p<0.05, * p<0.1, ● = included ○ = excluded. Standard errors in parenthesis allow for within-country clustering. The dependent variable is the instantaneous reproduction number 10 days forward: $R^{I}_{i,t+10}$. Controlled for daily weather conditions and weekly seasonality.

the importance of controlling for this factor. This result is consistent with first-stage results, which already suggest an important role for voluntary distancing in observed distancing behaviors.

In the third specification, I add other preventive policies to the set of controls. These are included as a set of different variables for which the parameter estimates are not shown in the table. The effect of distancing policies is marginally smaller compared to specification (2) suggesting that another preventive policies are also important controls to include similarly to voluntary activity. Finally, in column (4) I add imported cases and their interaction with different levels of international travel controls. The difference in the parameter estimates of place and mobility restrictions is negligible between columns (4) and (3) suggesting that international travels were not strongly correlated with distancing policy interventions.

Column (4) is my most complete specification, therefore, I refer to it as my main result. In column (4) I find that place restrictions reduce the effective reproduction number of COVID-19 by 29 percent, while mobility restrictions by 62 percent. These are strong effects suggesting that distancing policies were an effective tool for reducing the impact of the pandemic. Place restrictions that are targeting specific destinations are found to be roughly half as effective as general mobility restrictions. This finding suggests that there was much heterogeneity between the effectiveness of different policies, implying that different policy mixes could have led to very different outcomes

Now let us turn to the estimation results for the other important factors, voluntary distancing and imported cases in this most complete specification of column (4). In the case of voluntary activity, I find that a one standard deviation drop in voluntary social activity decreases R_t by 17 percent. The effect of voluntary distancing is also significantly negative but weaker than those of the policy restrictions. This finding suggests that, although voluntary distancing behaviors help to slow down the reproduction of the virus, any kind of distancing policy measures are much more effective in stopping a pandemic.

In the case of imported cases, I find that a one standard deviation rise in imported cases significantly increases R_t by 31 percent. This effect is more than offset by travel restrictions if it takes the form of quarantines or a total ban. Although this offsetting effect is only marginally significant for quarantines. I have found no evidence for the effectiveness of screening and selective travel bans in the reduction of R_t via imported cases.

1.4.4 Comparing Compliant and Voluntary Distancing Effects

It is crucial to compare the consequences of voluntary and distancing policy-induced distancing when forming policy conclusions about the relative efficiency of policies. Voluntary mobility and distancing policy-induced components are measured in different units, so Table 1.3 coefficients are not directly comparable between rows. One possible way to address this issue would be to use the first-stage estimates for place and mobility with their coefficients from equation (1.5) directly on the right-hand side of equation (1.9), instead of the policy variables. Although this strategy appears simple and straightforward, it is impossible to implement because policy variables, P_{it}^{place} and P_{it}^{mobility}, are not included in the first-stage equation. The reason they are not included is that the effect of the policies is captured by the RDiT design that builds on the key identifying assumption of sudden responses to policy changes. Giving this design up is considered to be a greater cost than the gain of the comparison that would emerge from a different design would provide.

I work around this problem by picking a different strategy to make the effects of distancing policies and voluntary distancing comparable. It is a decomposition of the changes in reproduction numbers around the time of the first global wave. I start that by taking the cross-country averages of R_{it} in the estimation sample. Reproduction numbers peaked in late February and declined until the summer. Assuming an average duration for an infection to be 12 days, the seven-day backward-looking average of the effective reproduction number peaked at 6.32 on February 26 in 2020, and fell below 1 for the first time on April 11. This was an 85 percent drop in R_{it} on average across countries in the first wave. I decompose this decline into four suggested factors by calculating the changes in the cross-country averages of distancing policies and the voluntary activity indicator in the same time period and then multiplying them by their coefficients of the most complete specification in Table 1.3.

Table 1.4: Comparing Compliant and Voluntary Distancing Effects

	Change (%)	Contribution
Effective Repr' Number R_t	-0.85	100%
Place Restrictions	-0.25 (0.102)	29.4 %
Mobility Restrictions	-0.59 (0.088)	69.4 %
Distancing Policies	-0.84	98.8%
Voluntary Distancing	-0.13 (0.032)	15.3 %
Other Factors	0.12	-14.1 %

Notes: Effects are calculated as the change in cross-country averages multiplied by the coefficients of column (3) of Table 1.3. Standard errors in parenthesis are calculated similarly, using the s.e. of the corresponding coefficient.

I proceed similarly with standard errors.

Taking cross-country averages of binary policy indicators gives the share of countries that are introducing that policy on that day, therefore its change in the period shows the change in the sample coverage of these policies. This change in coverage for place and mobility restrictions were 84 and 95 percent in this period. That means most countries in the sample implemented these types of distancing interventions within these roughly two-month periods. In the same period, voluntary activity declined by 0.8 standard deviations.

The results of these calculations are summarized in table 3. It shows that the drop in R_{it} was mostly due to restrictions, which all together contributed almost 100 percent of the total decline in reproduction numbers. It was mostly mobility restrictions that were responsible for this effect. Their sole contribution was nearly 70 percent. This suggests that mobility restrictions were much more effective than place restrictions.

Voluntary distancing on the other hand contributed only a little more than 15 percent, which was counteracted almost completely by other unexplained factors. This result suggests that voluntary distancing had only a marginal role in the containment of the COVID pandemic in the first wave.

1.4.5 Heterogeneous Effects

In this section, I investigate heterogeneous policy effects. The first of these exercises analyse the strength of the policy effects on different time horizons. I am interested in how long the effects identified in the main design last. In the second exercise, I break down the larger restriction categories: place and mobility restrictions, into their components. I also allow for heterogeneity in the different stringency levels in this exercise. I do this to provide comparative results for more specific policy interventions.

Heterogeneous Dynamics

Here I investigate how long the effects found in the main specification lasted. It is useful to know how long a government can rely on a place or a mobility restriction when they are fighting more and more waves of an epidemic. To address this question I modify equation (1.9) by allowing for time heterogeneity in the policy effects by the following modification to the main estimation equation (1.9):

$$R_{i,t+10} = \exp\left[\beta^{\text{place}}_{t-w(i)} + \beta^{\text{mobility}}_{t-w(i)} + \beta_v \widehat{v}_{it} + \xi' X_{it} + \mu_i + \kappa_t\right] + \varepsilon_{i,t+l}, \quad (1.10)$$

where $\beta^{\text{place}}_{t-w(i)}$ and $\beta^{\text{mobility}}_{t-w(i)}$ are event time dummies indicating seven day periods and relating the effect of a type of intervention to the last seven day period ($w(i) = -1$) before the intervention. This equation is other than this modification is equivalent to equation (1.9).

10.14754/CEU.2022.06

Figure 1.7: Effects of Distancing Policies on the Reproduction Number 10 days later

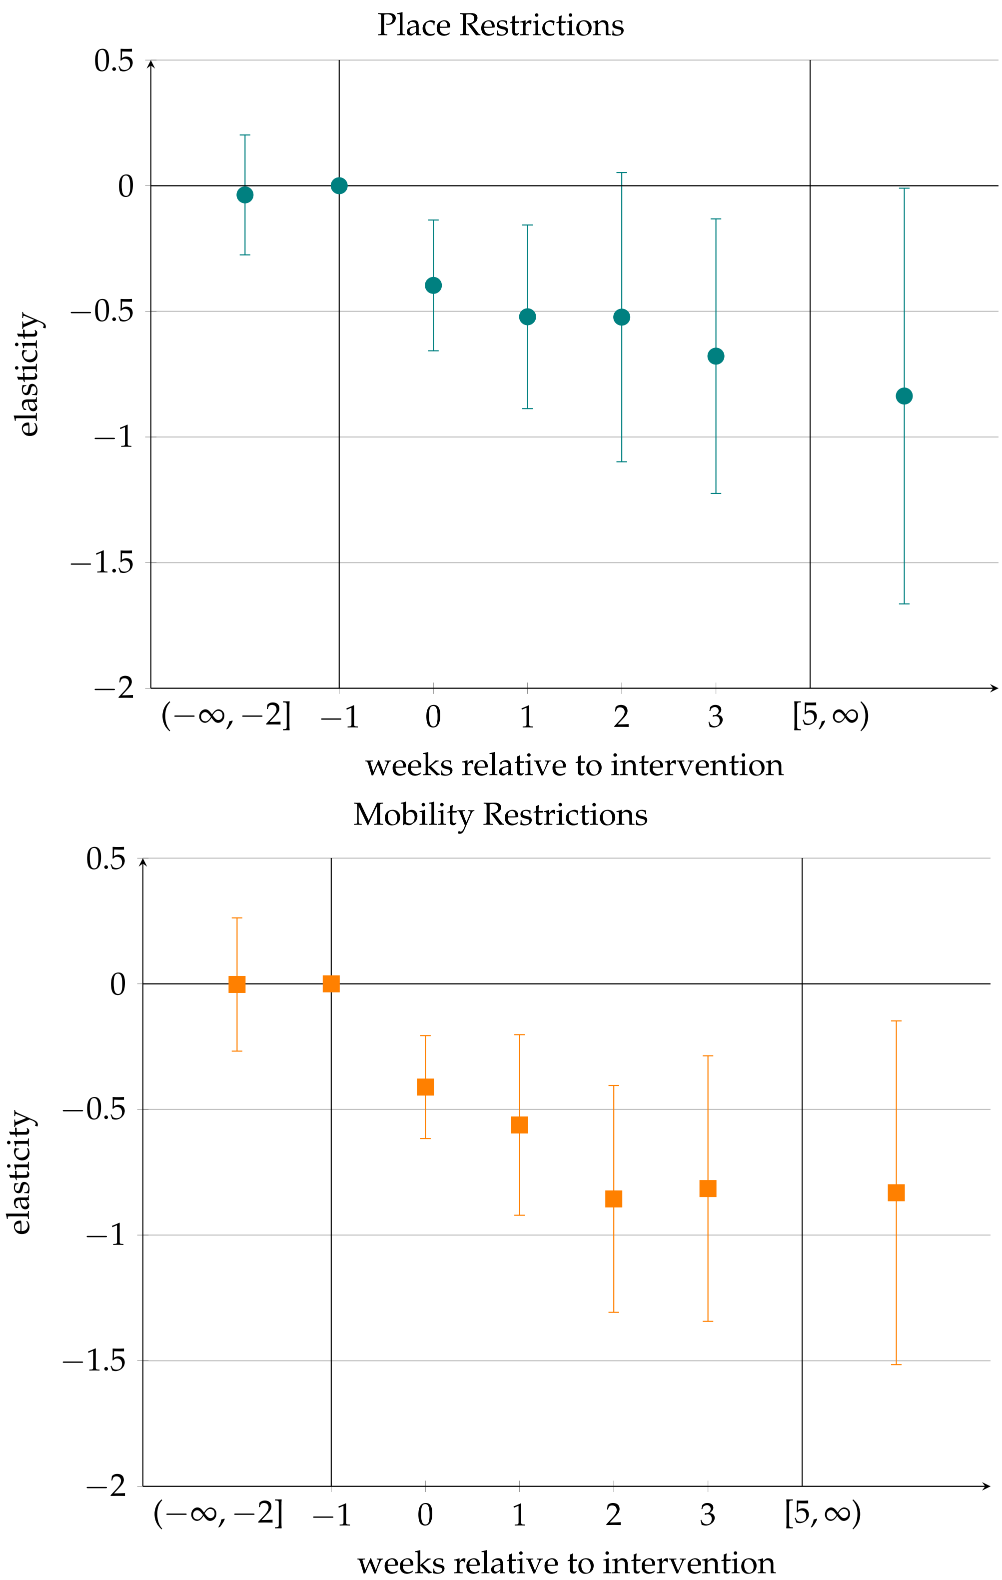

Notes: point estimates of $\beta^{\text{place}}_{t-w(i)}$ and $\beta^{\text{mobility}}_{t-w(i)}$ coefficients and 95% confidence intervals of equation 1.10. Standard errors are allowed to cluster within countries. Reference period: last 7 days before the intervention. 26,566 daily observations within 109 countries.

Figure 1.7 shows the results for the β coefficients of equation (1.10). It looks like both policies produce a significant drop in R_{it} already in the first seven days after their implementation. These effects get somewhat stronger in later weeks. In the case of a mobility restriction, these effects stay significant throughout the entire horizon. In the case of place restrictions, some longer-horizon effects are only marginally significant. Overall these results suggest that both policy types have a significant long-lasting effect on the reproduction number. Governments can rely on these distancing restrictions on longer horizons when fighting longer waves of infections based on these results.

Heterogeneous Policies

In this section, I estimate another variant of the second stage, where I break down the comprehensive place and mobility restriction indicators into their components, and estimate the effect of those components and their stringency levels separately. I do this to provide comparative results for more delicate policy interventions.

I estimate the following variant to equation (1.9), where I include all the different distancing interventions with all their different stringency levels reported by Hale et al. (2020):

$$R_{i,t+l} = \exp\left[\beta_{jk} D_{it}^{\text{jk}} + \beta_v \widehat{v}_{it} + \xi' X_{it} + \mu_i + \kappa_t\right] + \varepsilon_{i,t+l}, \tag{1.11}$$

where $D_{it}^{\text{j,k}}$ is an indicator indicating if a distancing policy j at stringency level k was in action in country i on day t. For example $j =$ workplace closures, which are $k =$ recommended. All other parts of the equation are equivalent to equation (1.9).

The results for equation (1.11) are reported in Table 1.5. Results show that four of these restriction types were found to be generally effective in my global estimation sample. For the effectiveness of cancellation of public events, restrictions on public transportation, and inland travel restrictions on the other hand I found no evidence.

Looking at policies one by one I find that school closures seem to be effective only if they are mandatory even if they are partial in terms of education levels. A required school closure reduces the effective reproduction number by between 17.6 to 24.1 percent on average depending on the coverage. Closing workplaces seems to reduce R_{it} significantly, no matter if it is recommended. The effect becomes 1.5 times stronger: -31 percent when it is required. There is no significant difference though between a partial requirement and if it includes all non-essential businesses.

Table 1.5: Effect of All the Different Distancing Policies on $R_{i,t+10}$

School Closures:		Public Transport Restrictions:	
Recommended	-0.160	Recommended	0.021
	(0.109)		(0.061)
Required Partial	-0.241**	Required	0.029
	(0.104)		(0.081)
Required All Levels	-0.176*	Stay At Home Orders:	
	(0.100)	Recommended	-0.185***
Workplace Closures:			(0.059)
Recommended	-0.209**	Required with exceptions	-0.149**
	(0.093)		(0.070)
Required Partial	-0.310***	Minimal exceptions	0.356*
	(0.080)		(0.205)
All Non-essential B's	-0.265***	Inland Travel Restrictions:	
	(0.092)	Recommended	-0.067
Public Events Cancellations:			(0.072)
Recommended	0.102	Required	0.093
	(0.111)		(0.077)
Required	0.015		
	(0.096)		
Gathering Limits:			
1000+	-0.073	10+	-0.331***
	(0.133)		(0.112)
100+	-0.268**	1+	-0.345***
	(0.112)		(0.124)
Observations		38,704	
Countries		111	
All Controls		•	
Country and Day FE's		•	

Notes: *** p<0.01, ** p<0.05, * p<0.1, • = included ◦ = excluded. Standard errors in parenthesis allow for within-country clustering. The dependent variable is the instantaneous reproduction number 10 days forward: $R^I_{i,t+10}$. Controlled for daily weather conditions and weekly seasonality.

Gathering limits show strong effects when the limit is at most 100 persons. The introduction of a 100+ limit reduces R_{it} by 26.8 percent. A less stringent 1000+ limit shows no effects, while the more stringent 10+ or 1+ limits seem to have stronger effects but with strongly diminishing gains. Stay-at-home orders, i.e. curfews are also found to be effective already if they are recommended. A recommended home staying reduces R_{it} by 18.5 percent. When it is mandatory with minimal exceptions by 35. percent, which is the strongest effect among all policies, but only weakly significant however only at 10 percent. Overall these findings suggest that there was much heterogeneity between the effectiveness of different policies, implying that different policy mixes could have led to very different outcomes.

1.5 Conclusions

In this study, I estimated the effect of distancing interventions on the effective reproduction number R_t of COVID-19. I was focusing on the effects of two types of such policies, place restrictions, that target specific destinations, and mobility restrictions which are general restrictions on inland movements. The main contribution of this study is the separation of voluntary and policy-induced distancing. I have found that distancing interventions had a strong and permanent effect on R_t. General mobility restrictions are found to be roughly two times more effective than targeted place restrictions. These policy effects were found to be much more dominant than the effects of voluntary distancing. These results suggest that governments can use distancing restrictions effectively in pushing down the effective reproduction number below the containment threshold of $R_t \leq 1$. Although these policies need time to exert their effects on reported case numbers, governments can rely on their effects for as long as these measures are in place.

Comparing specific interventions I have found significant differences. Based on these results school closures are better if they are mandated, in contrast with workplace closures, which were found to be effective already, when they are just a recommendation. Stay-at-home orders are similarly effective already, when they are only recommended, but more effective when mandated with only minimal exceptions. Gathering limits become effective below 100 people and only get marginally more effective at more restrictive limits. I have found no supporting evidence

10.14754/CEU.2022.06

for the effectiveness of cancellation of public events, restrictions on public transportation, and inland travel restrictions. These results suggest, therefore, that a careful selection of particular distancing policies and their stringency levels is recommended before their implementation.

Chapter 2

Economic Costs of Distancing Policy Interventions

2.1 Introduction

When a new virus bursts into an epidemic and no vaccines are available, the primary containment strategy is distancing policy interventions (DPIs). COVID-19 was no exception. DPIs limit social interactions in order to prevent virus transmission, but they also impose significant costs on economic activity. This paper quantifies the economic costs in terms of sector output losses, inflationary effects, and unemployment responses caused by DPIs during the COVID-19 pandemic on a country-level weekly panel dataset.

Distancing happens not only in compliance with DPIs but also as a voluntary response to threatening news about the new virus. The disentanglement of the effects of policy-induced and voluntary distancing behaviors is the main contribution of this paper. The main empirical challenge of the paper is that these two effects are strongly correlated. I tackle this challenge by observing that, although an overall downward trend in social interactions started after COVID hit a country, there was also a sudden and substantial drop in social interactions right after DPIs were implemented. I leverage this stark salience of behavioral responses to DPIs to separate the policy-induced and voluntary distancing components of overall distancing behaviors. This separation is accomplished through the use of a regression discontinuity in time design, which identifies the policy-induced component as the discontinuity and the voluntary component as the residual. Once I have this voluntary distancing component, I can use it as a control in the main estimation of DPI effects on economic activity. These effects are identified from the change in seven selected economic outcomes after DPI interventions compared to their changes from the same months of five pre-COVID years holding other correlated factors, such as voluntary distancing, fixed. Having this voluntary distancing component identified allows me to hold that factor fixed in the main estimation of DPI effects on economic activity. This empirical strategy has been developed in an earlier paper by the author: Rácz (2022).

The next section of the paper describes data sources and their transformation into two estimation samples. This study uses three main data sources: (i) one for economic outcomes, (ii) another for distancing policy interventions and other COVID-related interventions, (iii) and the third one for social mobility proxying social distancing behaviors. This third data source is Google's COVID-19 Aggregated Mobility Research Dataset.[1] The estimation samples cover 44 countries. The first sample used in the identification of voluntary distancing patterns is a weekly-country level panel dataset spanning every week between November 2019 and December 2020. The second sample used in the identification of the economic effects of DPIs covers every month between November 2015 and October 2020. The second sample contains seven selected economic indicators, such as industrial and manufacturing production, construction output, retail trade, consumer prices, producer prices in manufacturing, and the unemployment rate. Intuitively, service sectors, such as personal services, accommodation, food and drink services, or the entertainment sectors, must have suffered the most losses under DPI restrictions. These sectors are omitted because of a lack of data availability in the current version of this paper. The omission of these specific service sectors is a clear limitation of the paper, and it suggests the underestimation of the output effects of DPIs.

I define three different DPI indicators using data from Hale et al. (2020): DPI treatment, DPI intensity, and DPI extensity. The DPI treatment captures the average level of the first DPIs within each country. Treatment is held constant until the end of the sample as if these interventions were

[1]For details, see Section B.1 of the Appendix.

kept on all along, even if they were not. This way, DPI treatment resembles a more conventional treatment, making it easier to interpret its effects. Later changes in DPIs are absorbed into two other indicators: DPI extensity and intensity. Because of how DPI treatment is designed, these two indicators capture deviations of DPIs relative to the original interventions. DPI intensity measures deviations in stringency levels, and DPI extensity in DPI types, such as school closures or stay-at-home orders.

In Section 3, I present the empirical strategy. I carry out my estimation in a two-stage design. The first stage is the separation of the voluntary and policy-compliant components of social mobility. The second stage is the estimation of the economic effects of DPIs. The first stage identifies DPI effects on a social mobility indicator calculated from mobility data from carry-on devices. It is carried out in a regression-discontinuity-in-time design. The main identifying assumption is that DPI effects must appear suddenly as they impose restrictions on the whole society from one day to the other, while voluntary distancing effects are realized much slower as voluntary decisions are heterogeneous within a society. Patterns observed in raw mobility data give visual support for the discontinuity design. I define a voluntary mobility indicator by residualizing social mobility to first-stage predicted DPI effects. The second stage estimates the economic effect of DPIs, controlling for the first stage's predicted voluntary mobility, other COVID-related policy interventions, and other factors. The second stage design includes controls for distancing patterns at trading partners, anticipating the possibility of international spillovers of economic effects.

Section 4 presents the main results. It starts with first-stage results, then it describes the identification of the voluntary mobility indicator. Second stage results about the economic effects of DPIs are presented after that. Second-stage results are presented in three separate subsections for output losses, inflationary costs, and unemployment effects.

I find that distancing behaviors that were either voluntary or DPI-compliant generated substantial output losses. I found significant output losses due to DPIs, but no evidence for inflationary and unemployment effects. Findings suggest that DPIs caused substantial output losses. Results also show that although voluntary distancing caused significant losses to sector outputs, its effect was an order of magnitude smaller than that of DPIs. Only 70% of total losses in industry and manufacturing is explained by either voluntary or DPI-induced distancing, implying that other factors, such as other COVID-related interventions contributed substantially to output losses in these sectors. In construction and retail trade, on the other hand, distancing factors altogether predicted more losses than was observed. This finding suggests that other factors, such as fiscal and monetary support programs, could mitigate the short-term costs of distancing in these two sectors.

I find no evidence of significant economic costs resulting from the introduction of new types of DPIs or from mobility spillovers from trade partners.[2] Changes in the intensity of distancing interventions, such as decreasing the limit for allowed gathering sizes, were found to increase the output costs of the first interventions.

These findings provide evidence that, while DPIs were implemented to contain COVID infections, they imposed substantial costs on economic activity in terms of output loss in industrial production and retail trade. Voluntary distancing induced an order of magnitude lower output losses than DPIs did. Inflationary and unemployment effects were not detected. These results provide both qualitative and quantitative guidance for governments to consider when implementing distancing interventions in times of an epidemic. These findings also contribute to a more complete cost-benefit analysis of distancing policy interventions on the cost side.

Literature

This paper belongs to the literature on empirical evaluations of non-pharmaceutical interventions (NPI) during the COVID-19 pandemic, surveyed exhaustively by Perra (2021). Within this literature, this paper is a contribution to the assessment of the economic costs of NPIs. There are papers that provide correlative evidence between such interventions and economic outcomes. Chen et al. (2020) find that European countries and U.S. states that experienced larger outbreaks also suffered larger

[2]Except in construction.

economic losses. They find no evidence of NPIs making significant contributions to these losses. Carvalho et al. (2020) consider billions of transactions from Spanish card data and find strong consumption responses to business closures, but smaller effects for capacity restrictions; a steeper decline in spending in rich neighborhoods. Arnon et al. (2020) find that NPIs explain nearly 15 percent of the decline in employment around 3 million jobs over the first three months of the pandemic. Bodenstein et al. (2021) finds that distancing being it either voluntary or policy-compliant had significant economic effects. The main contribution of this study relative to these papers is the identification of the causal effects of NPIs.

There are studies that identify causal effects of NPIs similarly to the aim of this paper. There are papers that find strong voluntary effects. Deb et al. (2021) find that containment measures had a significant impact on economic activity for example. Their findings suggest that industrial production losses were around 10% in the 30 days following their implementation, which is very close to the findings of this paper. Berry et al. (2021) find minor but negative economic effects of NPIs. They also stress the importance of voluntary distancing behaviors, when they claim that ″Many people had already changed their behaviors before the introduction of shelter-in-place orders.″ The contribution of this study relative to these papers is the separation of NPI-induced and independent voluntary distancing effects.

An earlier study of the author of this paper is Rácz (2022). In this paper, I employ the empirical strategy used in this paper as well to estimate the causal effects of distancing policy interventions on the effective reproduction number of COVID-19.

Goolsbee and Syverson (2021) Compare ″consumer behavior over the crisis within the same commuting zones but across state and county boundaries with different policy regimes.″ They find that NPIs account for only a modest share of the documented consumption decline. This comparison, however, identifies the effect of NPIs decoupled from NPI-induced voluntary effects, which this study considers as relevant consequences of NPIs. Kong and Prinz (2020) find no evidence of unemployment effects of NPIs in line with this study but in a sample of US states.

Bodenstein et al. (2021), and Goolsbee and Syverson (2021) stress the importance of voluntary distancing in US states. This study finds that voluntary distancing effects were less important in a global sample. This comparative assessment of the role of voluntary distancing effects is supported by Maloney and Taskin (2020), who use mobility data to identify the effects of NPIs.

This study takes into account international spillovers of distancing effects as a confounding factor of NPIs and economic outcomes. There are studies that document such spillovers. For example, Boranova et al. (2022) provides evidence of international spillovers of distancing effects on car manufacturers. Barrot et al. (2021) find GDP responses to distancing through supply chains.

2.2 Data

There are three main sets of variables that are used in this study: (i) economic outcomes, (ii) distancing policy interventions and other COVID-related interventions, and (iii) social mobility. These sets of variables are derived from their own three different data sources. These three sets of variables are presented in more detail in the following subsections of this section. Besides these datasets, I also use two more auxiliary data sources. First, I used international trade data from the OECD in 2019 to calculate export and import shares by country pairs. Second, I calculate country-week level averages of various weather indicators using data from the National Oceanic and Atmospheric Administration (NOAA). I merge all this data into two estimation samples: a weekly frequency country-level panel used in the identification of DPIs on social mobility and a monthly frequency country-level panel used in the estimations of the economic effects of DPIs. This section describes data sources and how they are transformed into estimation samples.

The estimation samples cover the following 44 economies:

1. Argentina
2. Australia
3. Austria
4. Belgium
5. Brazil
6. Canada
7. Chile
8. China

9. Colombia
10. Costa Rica
11. Czechia
12. Denmark
13. Estonia
14. Finland
15. France
16. Germany
17. Greece
18. Hungary
19. Indonesia
20. India
21. Ireland
22. Israel
23. Italy
24. Japan
25. Lithuania
26. Luxembourg
27. Latvia
28. Mexico
29. Netherlands
30. Norway
31. Poland
32. Portugal
33. Russia
34. Saudi Arabia
35. Slovakia
36. Slovenia
37. South Africa
38. South Korea
39. Spain
40. Sweden
41. Switzerland
42. Turkey
43. UK
44. USA

The second-stage sample spans every month between November 2015 and October 2020, while the first-stage sample starts from the first week of November 2019 and covers every week until December 2020.

2.2.1 Economic Outcomes

I estimate the effects of distancing policy interventions on the following seven monthly economic indicators:

- industrial production,
- manufacturing production,
- construction output,
- retail trade,
- consumer price index (CPI),
- producer price index (PPI) in manufacturing, and
- unemployment rate.

The first four of these are measuring the output of different sectors: industrial, and manufacturing production, construction output, and retail trade. Using these as outcome variables in the main estimation addresses the question of output losses due to distancing policy interventions. Intuitively, service sectors, such as personal services, accommodation, food and drink services, or the entertainment sectors, must have suffered the most losses under DPI restrictions. These sectors are omitted because of a lack of data availability in the current version of this paper. The omission of these specific service sectors is a clear limitation of the paper, and it suggests the underestimation of the output effects of DPIs.

The next two outcome variables are price indexes addressing inflationary costs resulting from DPIs. Finally, the unemployment rate indirectly addresses job losses as a result of DPIs.

I normalize all seven indicators by their values from the latest November in order to get numbers that are comparable across years.[3] Figure 2.1 shows the seven economic outcomes in each country and each month around the first DPI by circle marks. X marks show their cross-country means, highlighting general tendencies. These observations are contrasted with a 5-year benchmark from the pre-COVID years of the same indicator, indicated by a thick orange line. Observations stay close to this benchmark before the first intervention in all seven graphs supporting the choice.

The figures also reveal that all four sector output indicators declined substantially from their benchmarks after the first DPI in almost all countries. In the first month following the first intervention, the average decline is around 20%, and it is consistent across all four output indicators. Even though this decline was temporary, as all four indicators converge to their benchmarks after roughly 6 months, they do not rise above them, meaning that this decline represents a permanent output loss in these sectors. The main scope of this paper is to quantify what share of this output loss can be causally linked to distancing policy interventions.

The two price variables show a much greater heterogeneity across countries in their price responses to DPIs compared to sector outputs. This observation is not too surprising as distancing disrupts both supply and demand, which have opposite effects on prices. Therefore, the overall effect of DPIs on prices can have varying signs across countries depending on the relative strength of demand and supply disruptions. Although there is this considerable heterogeneity across countries, both price indicators show a slightly negative and permanent deviation from their benchmarks, which is a little more pronounced in manufacturing PPI than in CPI. This suggests that distancing was relatively more demand-disruptive, especially for products of the manufacturing industries.

[3]Original data is seasonally adjusted for all indicators and is fixed price volume indices in the case of sector outputs.

Figure 2.1: Seven Economic Indicators around DPIs

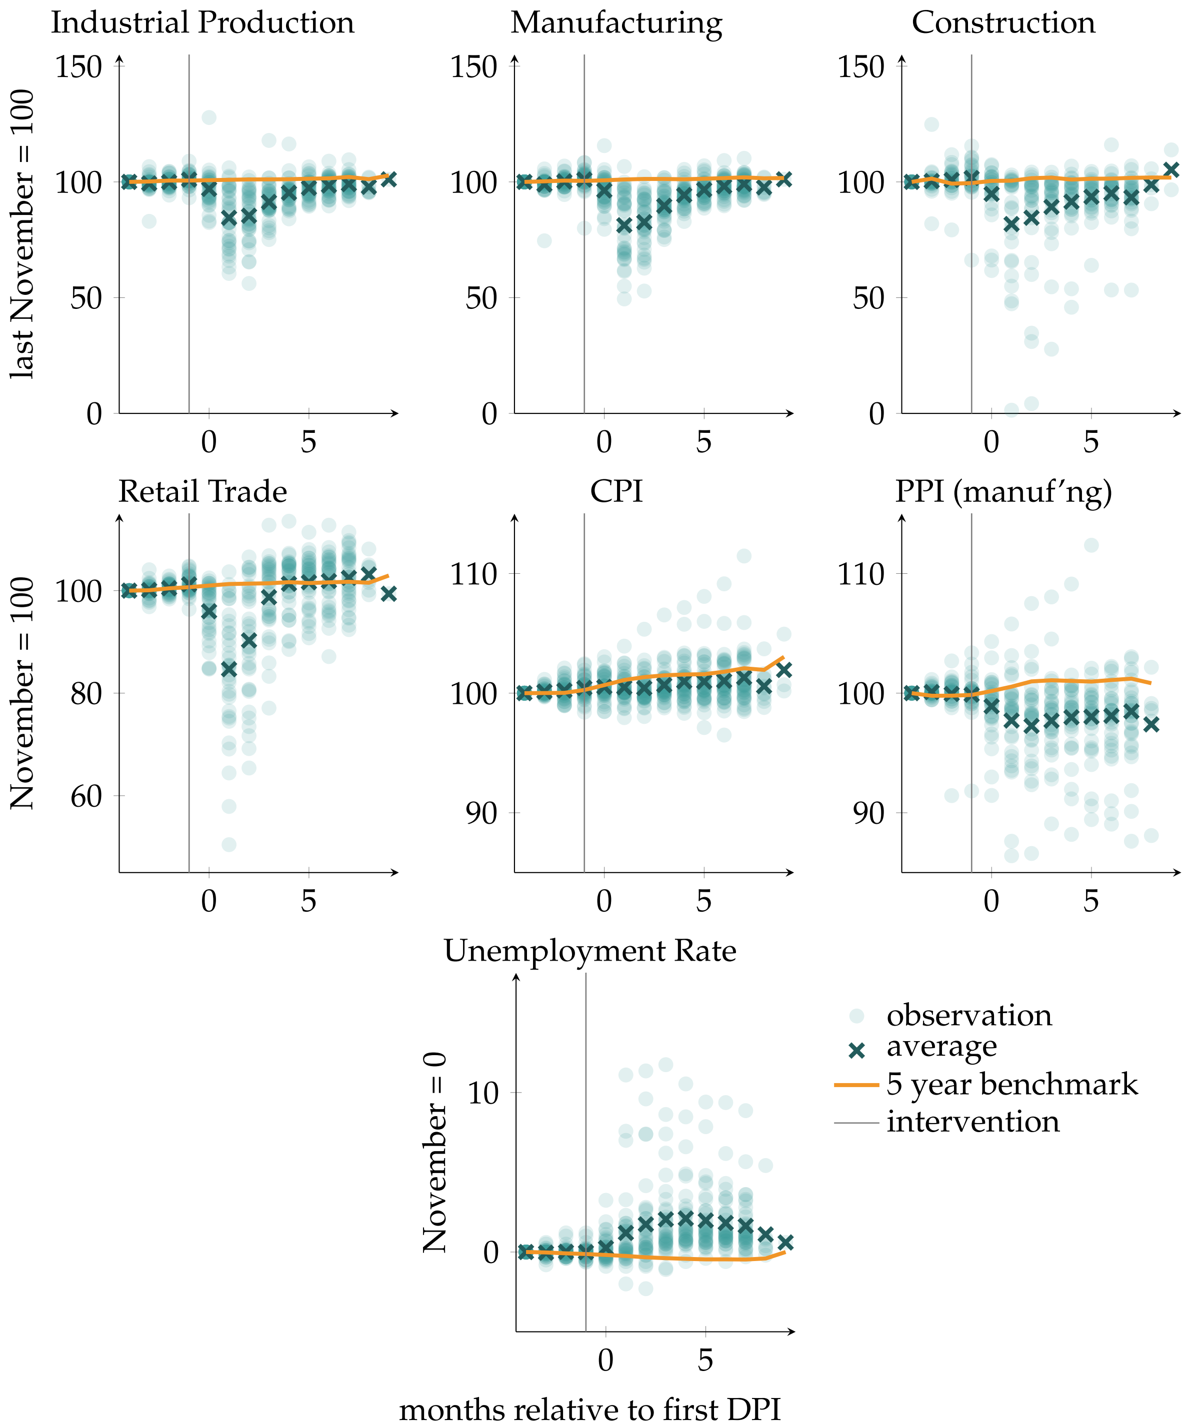

Notes: cloud: country-month observations around first DPI, darker regions show overlapping observations. X marks: cross-country averages around the first DPI. Thick orange line: averages of 5 pre-COVID years. Thin line: last month before first intervention.

Unemployment first increased after the first DPI in general, but it converged back to its benchmark levels 9 months later. Unemployment rates were on average 2 percentage points higher than their 5-year benchmarks throughout months 3-6, suggesting a substantial loss of jobs after DPIs were introduced. The question is, how much of this excess unemployment can be attributed directly to DPIs?

2.2.2 Distancing Policy Interventions (DPIs)

The primary data source for distancing policy interventions and other COVID-related interventions is Hale et al. (2020). It is a constantly updated dataset covering almost every country in the world. It reports several different COVID-related interventions on daily frequencies.

Distancing policy interventions, abbreviated as DPIs, are the main focus of this study. A DPI of type j is reported as a categorical variable $D^j_{it} \in \{0, 1, \dots k_j\}$, such that 0 signals no intervention and greater integers signal more and more stringent interventions, k_j being the most stringent type j intervention possible. One example is school closures, for which value 1 codes a recommendation, 2 a partial mandate, and 3 a mandate for all levels of education.[4]

I observe seven different DPIs:

- school closures,
- workplace closures,
- gathering limits,
- stay-at-home orders,
- within country travel restrictions, and
- cancellation of public events.

Considering the small sample size plus the fact that most countries in the sample started to intervene in the same month, in March 2020, it is very unlikely that the effects of these seven different interventions can be identified separately. Therefore, I first calculate their sum as $\overline{D}_{it} = \sum_j D^j_{it}$. This variable takes the value of 3, for example, if there was a level 1 school closure and a level 2 gathering limit in place in country i during the entire week t. It can also be non-integer when the number or the stringency of DPIs changes within a week. For example, if the level 1 school closure and the level 2 gathering limit was introduced as

[4]They also report a binary indicator for each DPI that indicates if a policy was countrywide or only local. In all my calculations presented here, I deduce 0.5 from a DPI categorical variable if it was only local, meaning a level 3 school closure gets a value of 2.5 if it was only regional.

the first interventions on a Wednesday, that would aggregate to a value of $\overline{D}_{it} = 3 \times 5/7 = 2.143$, because the aggregate value of 3 was only in place 5 days in that week.[5]

I decompose $\overline{D}_{it}$ into three distinct components, each of which is defined to be disjoint. The first component I name as the *treatment* (T_{it}). It captures the first DPIs within a country, with its magnitude remaining constant throughout the sample. T_{it} is defined as:

$$T_{it} = \begin{cases} \overline{D}_{it} = 0 & \text{if } t \leq 1 \\ \overline{D}_{it}\big|_{t=1} & \text{if } t > 1 \end{cases}, \tag{2.1}$$

that means treatment is 0 before the first ever DPI, which happens by construction. It then takes and keeps the retains of $\overline{D}_{it}$ of week 1 from the first-ever intervention throughout the sample. On week 0, T_{it} takes a value between 0 and the week 1 value of $\overline{D}_{it}$ depending on the number of days the first ever DPIs were in place on week 0. For example, if the first ever DPIs were introduced on a Monday, the week 0 and week 1 values are the same, but if they were introduced on Friday of week 0, it takes only 3/7 of its week 1 value given it was in place for only 3/7 of the week.

Figure 2.2 shows the evolution of the treatment (T_{it}) components on weekly frequencies. Squares indicate country-week observations, such that darker regions show overlapping observations. It shows a substantial heterogeneity across countries in the magnitude of their first interventions. Crosses indicate cross-country averages highlighting the general pattern across countries, which is around 8, revealing that many countries introduced their first DPIs in bundles and started some on higher than level 1. This figure also confirms the concept of this component as it is defined to be fixed throughout the sample after week 1. Any further alterations in the number or the level of restrictions are absorbed in the other two components.

The two other components I define are referred to as *extensity* (E_{it}) and *intensity* (I_{it}). Extensity captures any further changes in the number, and intensity of the level of DPIs after the first intervention. For the formal definition of E_{it} and I_{it} the definition of following two numbers is help-

[5] And the sum of all DPIs was 0 in the first two days of that week.

Figure 2.2: Treatment: T_{it}

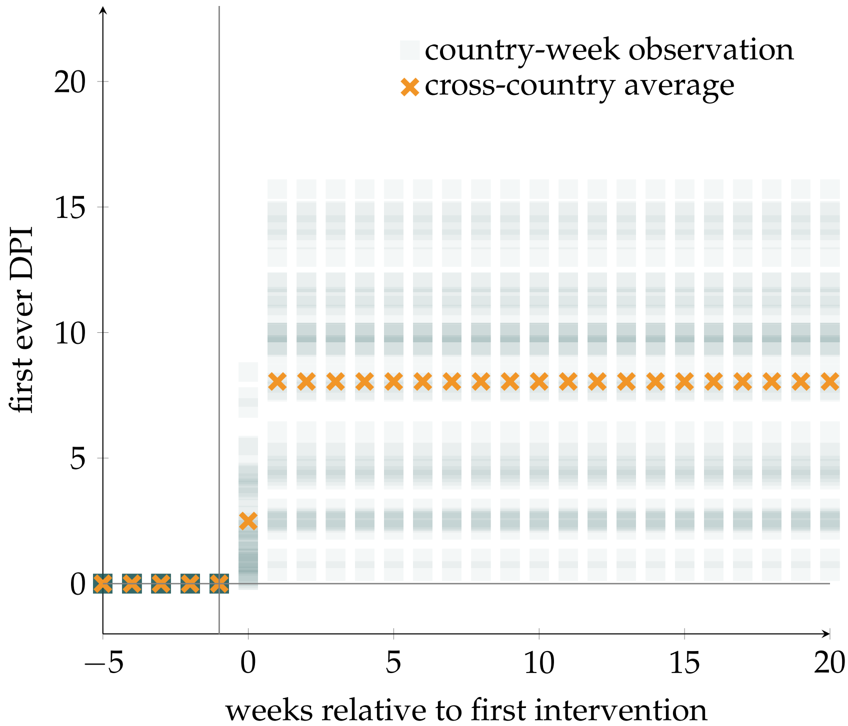

Squares: country-week observations, darker regions show overlapping observations. +: within week averages. vertical line: shows last week before the first intervention.

ful:

$$N_{it} = \sum_j \mathrm{I}\left(D^j_{it} > 0\right) \qquad\qquad L_{it} = \overline{D}_{it} - N_{it},$$

where $\mathrm{I}(\ldots)$ is an indicator function. N_{it} is the sum of the number of DPIs, while L_{it} sums the level of DPIs above 1. For example, if there were only a level 3 school closure and a level 2 gathering limit in a country i on week t, $N_{it} = 2$, because there are only 2 types of DPIs in place, and $L_{it} = 3$, because these DPIs are 3 levels above level 1 in total. E_{it} and I_{it} are formally defined as:

$$E_{it} = \begin{cases} N_{it} - N_{it}|_{t=1} & \text{if } t > 1 \\ 0 & \text{otherwise} \end{cases} \qquad I_{it} = \begin{cases} L_{it} - L_{it}|_{t=1} & \text{if } t > 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

This way $\overline{D}_{it} = T_{it} + E_{it} + I_{it}$, that is they are disjoint and contain all the information coded in D^j_{it}'s.

Figure 2.3: Treatment Dynamics: Extensity and Intensity of DPIs

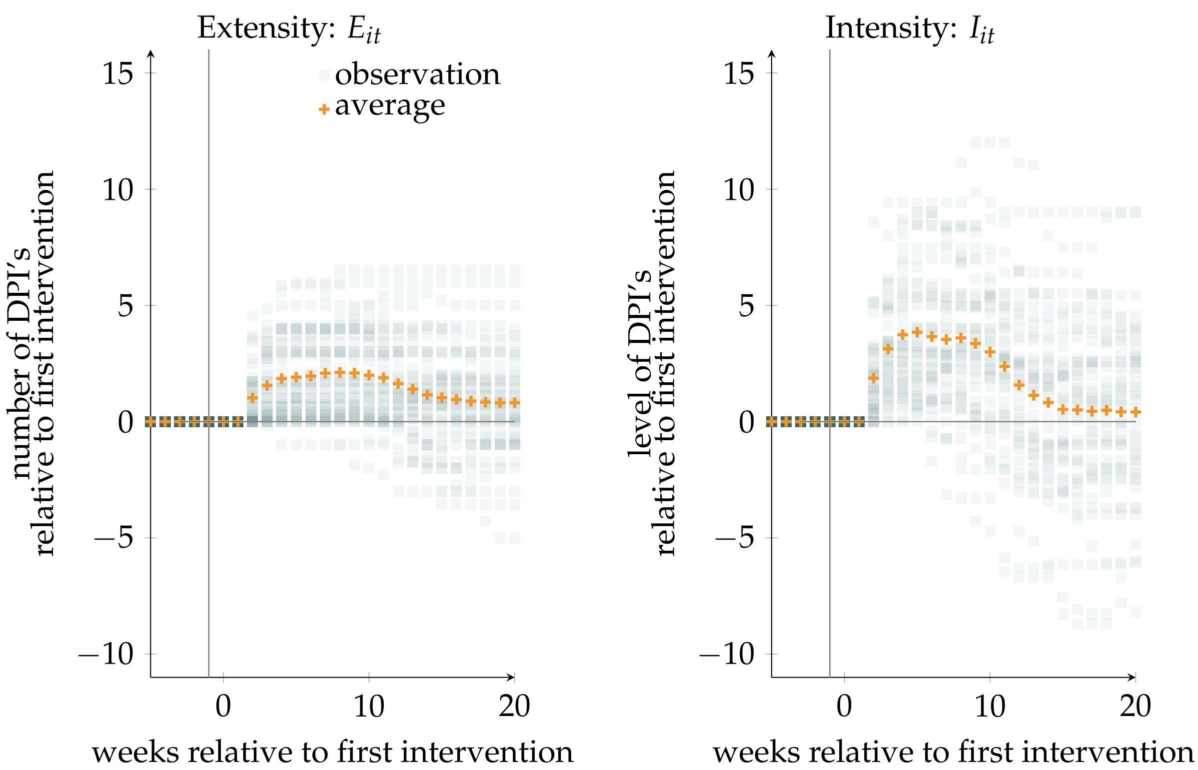

Squares: country-week observations, darker regions show overlapping observations. +: within week averages. vertical line: shows last week before the first intervention.

The two panels of Figure 2.3 show the evolution of E_{it} and I_{it}. Squares indicate country-week observations spreading out considerably in both figures showing substantial heterogeneity across countries in both the extensity and intensity of DPIs. Crosses show cross-country averages highlighting the general pattern, which is growing in the first couple of weeks and starts to decline between weeks 10 and 20. This pattern shows that after the first interventions countries tended to increase both the number and the level of DPIs in the first couple of weeks and started to slacken up restrictions only after 10 weeks.

Frequency conversion. Interventions are all reported on daily frequencies, which I convert to weekly and monthly frequencies in my estimation samples. I take the weekly (monthly) averages for all of these variables at weekly (monthly) conversions. In the case of categorical variables, this means if that categorical variable switches from category 0 to 1 on a Wednesday of a given week, the weekly conversion of that variable is 5/7=0.714 for that week. Conversion of within-month changes happens the same way. I can do that, because categorical variables are ordinal, meaning that if for example, an intervention switches from a value of 1 to 2, that means that intervention becomes more restrictive.

2.2.3 Social Mobility Index

The main contribution of this paper is the identification of the causal effects of DPIs on distancing in isolation from voluntary distancing effects. To track distancing patterns, I create a weekly index of social mobility (m_{it}) using Google's COVID-19 Aggregated Mobility Research Dataset. This dataset provides anonymized records of weekly flows of Google users[6] between NUTS3 areas. This data is available from the first week of November 2019, for every consecutive week until today. Any further details regarding this dataset can be found in the Appendix.

I calculate the social mobility index, m_{it} in two steps. First, I take only inflows into NUTS3 areas and normalize them by their average values for a 4-week period between November 3 and November 30, 2019, because this is the first available month-long period, which is as far from the time of the COVID pandemic as possible. This gives m_{it} a unit of percentage deviation from November 2019, similar to economic outcomes. In the second step, I aggregate these normalized NUTS3 level inflows country level by taking their arithmetic mean within a country-week cell. Based on this definition less social mobility (lower m_{it}) means more distancing.

Figure 2.4 shows the social mobility index m_{it} around the time of the first DPI. One circle is a single country-week observation, starker regions show overlapping observations. For this figure, I normalized m_{it} by its pre-intervention mean within each country. m_{it} fell by between 10 to 80 percent relative to the pre-intervention period within 2 weeks after the first DPI in almost every country according to the figure. This sharp decline in social mobility right after the first DPI gives a rationale for a discontinuity design in the identification of DPI-induced distancing, which is elaborated on in the next Section.

[6]Only of users who have turned on the Location History setting, which is off by default. This is similar to the data used to show how busy certain types of places are in Google Maps – helping identify when a local business tends to be the most crowded.

Figure 2.4: Social Mobility Index: m_{it}

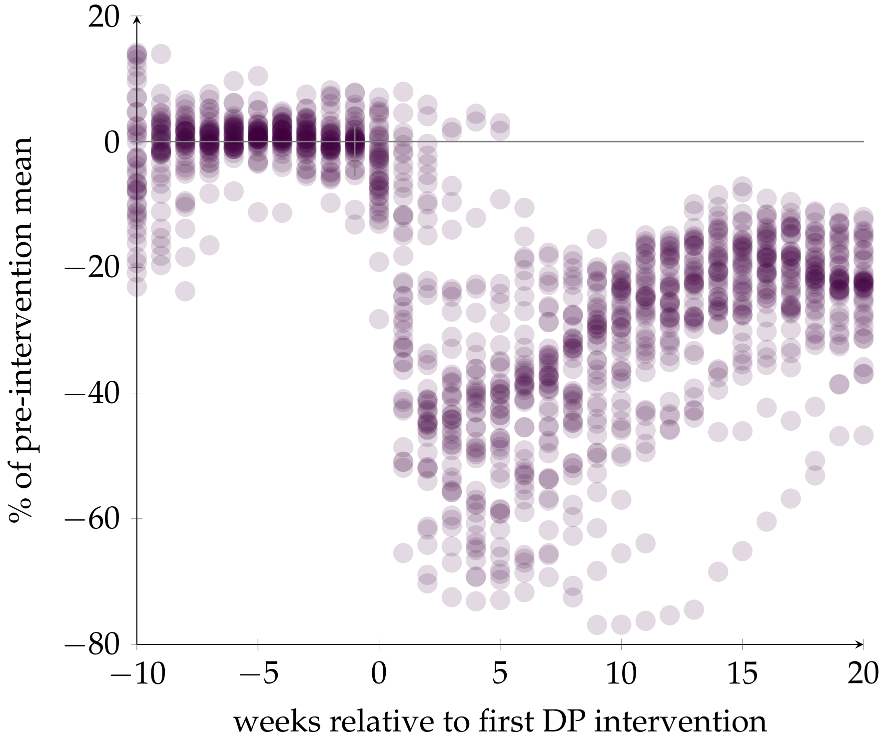

Squares: country-week observations, darker regions show overlapping observations. +: within week averages. vertical line: shows last week before the first intervention.

2.3 Empirical Strategy

The main scope of this paper is to identify the causal effects of distancing policy interventions on selected economic outcomes. An empirical strategy that leads to the identification of these effects is presented in this section. The primary empirical challenge is telling apart the economic effects of DPI-induced and voluntary distancing effects. I use a two-stage empirical strategy in which the first stage identifies voluntary distancing effects by separating social mobility m_{it} into a voluntary and a policy-induced component. I use the voluntary mobility component in the second stage as a control to be able to identify the causal effects of DPIs in isolation from voluntary distancing effects. I start with the discussion of the main identification strategy in the second stage. I then elaborate on the identification of voluntary distancing effects on social mobility in the first stage. The empirical strategy outlined in this section has been developed in an earlier paper of the author: Rácz (2022).

2.3.1 Economic Effects of DPIs

The identification of the economic effects of DPIs is based on a difference-in-differences approach. They are identified as changes in economic outcomes after the implementation of the first DPI compared to a COVID-free control period and holding other confounding factors fixed.[7] This control period is chosen to be the five years that preceded the COVID pandemic: 2015-2019, such that observations are matched by month. This strategy can be formalized by the following equation:

$$\widetilde{\Delta} y_{it} = \underbrace{\beta^T T_{it} + \beta^E E_{it} + \beta^I I_{it}}_{\text{DPI effects}} + \xi' X_{it} + \varepsilon_{it}, \tag{2.3}$$

where y_{it} is an economic outcome, and $\widetilde{\Delta}$ indicates difference from control period values. T_{it}, E_{it}, and I_{it} are capturing the first DPI intervention and further changes in the extensity and intensity of DPIs.[8]

X_{it} is a set of covariates that includes all relevant confounders that must be held constant in order to identify the effect of DPIs (*beta*). One way to find these confounders is to map out all the relevant causal links connecting distancing policy interventions to the economy.[9] Figure 2.5 shows the causality map of this paper. Each arrow represents a causal link, with thick arrows emphasizing the link to be identified. Solid lines indicate observed links; dashed lines indicate unobserved links.

The causality shows that DPIs only indirectly affect the economy through the reduction of social mobility, which has a potentially disruptive effect both on aggregate demand and supply. The effects of DPIs are conveyed by two channels: policy compliant and policy-induced voluntary distancing. People might increase their distancing after the implementation of a restriction because of compliance, but might also because they perceive it as a signal of a worsening epidemic. The primary motivation for this paper is to inform policymakers about the total effect of DPIs. These are realized through both of these channels, I do not aim to identify them separately, therefore, in this paper.

[7]The validity of this choice is supported by Figure 2.1 as it has been discussed in Section 2.2.1. The main observation there was that treatment period observations concentrate in the close neighborhood of the average of the control period values before the first DPI treatment.

[8]For details see Section 2.2.2.

[9]This technique is referred to as the DAG method by Cunningham (2021).

Figure 2.5: Causality Map

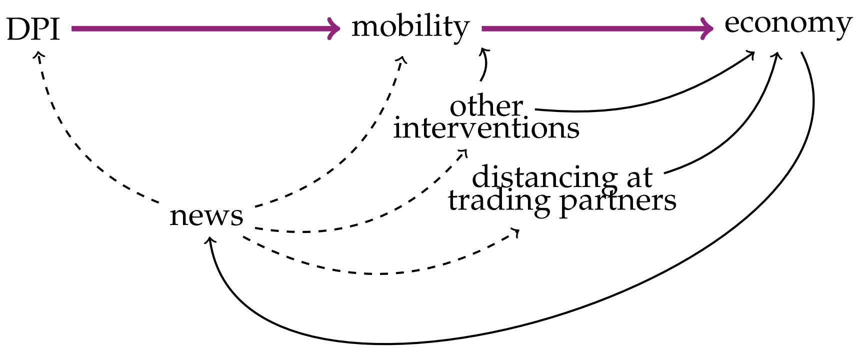

Notes: Arrows point in the direction of causality. Thick arrows: the path to be identified Solid line: observed, dashed: unobserved effect.

This map reveals three other paths through which DPIs and economic activities are also connected. First, DPIs are confounded with social mobility by news, which is a set containing any bits of information about COVID-19 that has the potential to alter government and individual distancing decisions simultaneously.[10] For example, the discovery of a large number of COVID infections raises the probability of a DPI and is also likely to discourage people from social activities. Throughout this paper, I am going to refer to this discouragement effect as *voluntary distancing*.

Second, governments implemented other COVID-related interventions that likely affected both social and economic activities. For example, income support programs aimed to prevent mass layoffs, but they might as well have encouraged people to stay home when they had COVID symptoms, decreasing social activities. Third, economies are interconnected through international trade links. Distancing in a country thus not only impacts domestic markets but can also have an influence on trading partners. For example, declining supplies increase import prices. Finally, economic shocks might also have contributed to news. For example, a negative shock to an economy could deter the government

[10]The arrow connecting news to DPI and other interventions acknowledges the fact of endogenous selection of the treatment of this study: DPIs. By closing all backdoor paths that contain this link, I simultaneously eliminate this selection bias.

from the most stringent DPIs.

The easiest way to eliminate the effects of these alternative paths would be to control for news and other interventions. Holding them fixed would identify the economic effect of DPIs. Unfortunately, this strategy is not feasible as news contains unobservable components. For example news about the risk of COVID is not observable in my data sources.

To overcome this difficulty, I instead construct a two-stage empirical design in which I estimate voluntary distancing in the first stage because it is unobserved. I do that by taking the mobility indicator m_{it}, and separating it into a voluntary and a policy-induced component using a regression discontinuity design.[11] For more details, see Section 2.3.2. I then estimate the economic effects of DPIs in the second stage, in which I control for other interventions, distancing at trading partners, and voluntary distancing. This design eliminates all alternative paths, including the reverse causality path of economic outcomes, because it is contained by the other three channels through news based on the causality map in Figure 2.5. I control for voluntary distancing by voluntary mobility predicted by the first stage. I measure distancing at trading partners by averaging social mobility index m_{it} using export and import shares.

Rácz (2022)

Second Stage: Empirical Design

This strategy is formulated by the following equation:

$$\begin{aligned}\widetilde{\Delta} y_{it} = \underbrace{\beta^T T_{it} + \beta^E E_{it} + \beta^I I_{it}}_{\text{DPI effects}} + \underbrace{\beta^V \widehat{m}_{it}^V}_{\text{voluntary mobility}} + \underbrace{\eta' P_{it}^O}_{\text{other interventions}} \\ + \underbrace{\lambda_X \textstyle\sum_j w_{ij}^X m_{j,t-1} + \lambda_M \textstyle\sum_j w_{ij}^M m_{k,t-1}}_{\text{distancing at trading partners}} + \underbrace{\xi' X_{it} + \mathrm{FE}_i}_{\text{covariates and FEs}} + \varepsilon_{it},\end{aligned} \tag{2.4}$$

where i is a country, and t is a month. T_{it}, E_{it}, and I_{it} are capturing the first DPI intervention and further changes in the extensity and intensity of DPIs. Part A. of Table 2.1 shows descriptive statistics of these indicators. $\widehat{m}_{it}^V$ is predicted voluntary mobility resulting from the first stage estimation.

[11] m_{it} is derived from Google user mobility data. For details, see Section 2.2.3.

P_{it}^{O} is a set of other COVID interventions, such as COVID-related fiscal spending, investment in vaccines and healthcare, income support programs, debt relief programs, international travel controls, and public information campaigns.[12] Part B. of Table 2.1 shows descriptive statistics of these interventions.

Information on fiscal and monetary policy interventions that are not directly COVID-related, such as tax or interest rate cuts, is not included among controls. The omission of such controls is a clear limitation of the current version of this paper because such conventional policy steps were likely to be used to mitigate inflationary and unemployment effects in many countries. Moreover, governments and central banks anticipating higher unemployment or inflationary risks were likely to intervene more strongly. This presumed correlation between conventional policy interventions and outcomes is more likely to be absorbed by COVID-related policy interventions that are controlled for but has the potential to cause omitted variable bias in the coefficients of DPI interventions.

$w_i^X j$and $w_i^M j$are export and import shares from 2019 between countries i and j, which sum to 1 across partner countries denoted by j. The terms with summations, therefore, measure the average changes in mobility at trading partners, capturing the effects of distancing at trading partners. Part D of Table 2.1 contains descriptive statistics for average social mobility at export and import partners.

[12] Source: Hale et al. (2020).

Table 2.1: Descriptive Statistics, Second Stage

Covariate	mean	st.dev.	max.	min.	unit
A. Distancing Policy Interventions (DPIs)					
Treatment	11.19	5.14	18.00	0.00	no. + lvl of DPIs
Extensity	-0.56	1.29	2.97	-4.52	no. of DPIs
Intensity	-1.71	3.12	7.18	-9.00	lvl of DPIs
B. Mobility at Trading Partners					
at Export Partners	19.50	24.41	99.98	2.87	Nov '19=100
at Import Partners	25.20	22.48	100.00	4.02	Nov '19=100
C. Other COVID Interventions					
Fiscal Spending	18.65	138.22	2151.20	0.00	billion USD
Investment in Vaccines	0.06	0.34	4.02	0.00	billion USD
Healthcare Investment	1.87	19.36	306.56	0.00	billion USD
Income Support	1.36	0.78	2.00	0.00	categorical
Debt Relief	1.15	0.79	2.00	0.00	categorical
Internat'l Travel Restr's	2.61	1.15	4.00	0.00	categorical
Information Campaigns	1.86	0.47	2.00	0.00	categorical
D. Other Covariates					
Covid Cases	178.60	335.85	2820.71	0.00	per 10^5 citizen
Covid Deaths	3.94	6.90	55.39	0.00	per 10^5 citizen
Covid Cases at Neighbors	4.64	6.70	40.80	0.00	per 10^5 citizen
Covid Deaths at Neighbors	0.13	0.19	1.21	0.00	per 10^5 citizen

Notes: 288 country-month observations of 32 countries.

X_{it} contains reported COVID cases and related deaths in population shares both domestic and from neighboring countries. These are included to address the direct and spill-over effects of COVID infections on economic activity. Part D of Table 2.1 presents descriptive statistics of these covariates. Finally, country-fixed effects are included to absorb the effects of time-invariant differences among countries, such as levels of economic development, degree of openness, or demographics, that are probably correlated with both government decisions on DPIs and changes in economic outcomes.

2.3.2 Voluntary Distancing

The first stage estimation identifies the policy-compliant component m_{it}^V of social mobility m_{it} in a regression discontinuity in time (RDiT) design.[13] The voluntary component, called voluntary mobility, is then defined as the residual of the first-stage regression.

The DPI-induced component of mobility is identified as sudden changes in m_{it} after the first DPI. The identifying assumption is that changes in social mobility due to voluntary distancing are slow, while the response to a distancing intervention is quick, at weekly frequencies. Distancing interventions prescribe a coordinated and sudden reduction in social activities after an intervention. Similarly, coordinated and sudden voluntary responses could only happen if the risk assessment of COVID news were homogeneous within countries. There is anecdotal evidence to assume that nations are much more heterogeneous in this respect, considering the simultaneous presence of virus skeptics and overly cautious people in many countries. It is more likely, therefore, that aggregate voluntary mobility responses are smooth and gradual because different fractions of society respond with different time lags and with different intensities based on their different risk assessments of the news.[14]

2.3.3 First Stage: Empirical Design

Based on these assumptions, social mobility m_{it} is modeled by the following equation:

$$m_{it} = \delta_t + \gamma^E E_{it} + \gamma^I I_{it} + \theta P_{it}^O + \zeta' Z_{it} + \mathrm{FE}_i + \nu_{it} \tag{2.5}$$

where δ_t is an event-time coefficient indicating week t after the first DPI was implemented in each country i. δ_t is included to capture the common trend in social mobility around the weeks of a type p intervention.

[13]A regular RD exploits a discontinuous change in the close neighborhood of a border separating the treated and untreated samples. RDiT is a special case when the running variable is time, which is usually a discrete variable in empirical exercises. This discreteness allows us to identify the effect by event time dummies rather than a discontinuity in a continuous polynomial like in regular RD designs. This design is related to event study designs, but it lacks a control group. For more detail see Hausman and Rapson (2018).

[14]Figure 2.4, presented in Section 2.2.3, supports this assumption, as it shows a sudden drop in social mobility at the time of the first intervention, but smooth changes in other periods.

Because they are intended to capture the effects of the intervention relative to the previous week, δ_{-1} is omitted, as δ_0^p represents week 0 of the first-ever DPIs. The treatment effects of DPIs are identified by δ_0 and δ_1, because of the main identifying assumption that the first DPIs impact mobility suddenly after their implementation. The rest of the event time dummies are, therefore, assumed to absorb the common trend in voluntary mobility changes in earlier and later weeks relative to the treatment weeks.

E_{it}, and I_{it} are capturing further changes in the extensity and intensity of DPIs after the first DPI. P^O are other interventions that may affect social mobility, such as fiscal spending, population share of vaccinated people, international travel controls, income support, and debt relief programs, public information campaigns, testing, contact tracing, mask-wearing and vaccination policies, and protection strategies for the elderly population. Parts A and B of Table 2.2 show the summary statistics of these factors.

Table 2.2: Descriptive Statistics of Covariates, First Stage

Covariate	mean	st.dev.	max.	min.	unit
A. Distancing Policy Interventions (DPIs)					
Extensity	1.04	2.03	6.50	-5.00	no. of DPIs
Intensity	1.39	3.11	12.00	-8.71	lvl of DPIs
B. Other COVID interventions					
Fiscal Spending	4.01	62.61	1957.60	0.00	billion USD
Share of Vaccinated	0.00	0.02	0.53	0.00	per citizen
Internat'l Travel Restr's	2.25	1.42	4.00	0.00	categorical
Income Support	1.09	0.87	2.00	0.00	categorical
Debt Relief	0.99	0.84	2.00	0.00	categorical
Public Info' Campaign	1.59	0.79	2.00	0.00	categorical
Testing Policy	1.63	1.03	3.00	0.00	categorical
Contact Tracing	1.24	0.80	2.00	0.00	categorical
Mask Wearing Policy	1.76	1.52	4.00	0.00	categorical
Vaccination Policy	0.32	0.83	5.00	0.00	categorical
Protection of the Elderly	1.53	1.16	3.00	0.00	categorical
C. Covariates of Voluntary Mobility					
Average Temperature	11.45	10.81	39.60	-41.89	Celsius degree
Average Humidity	70.43	16.13	97.08	12.52	percentage
Average Rainfall	14.11	18.50	209.38	0.00	mm
Average Snowfall	0.33	1.18	13.24	0.00	m
Covid Cases	0.68	1.23	9.32	0.00	per 10^5 citizen
Covid Deaths	0.01	0.02	0.20	0.00	per 10^5 citizen
Covid Cases at Neighbors	0.07	0.12	1.33	0.00	per 10^5 citizen
Covid Deaths at Neighbors	0.00	0.00	0.02	0.00	per 10^5 citizen

Z_{it} contains four weekly weather indicators, such as average temperature, humidity, snowfall, and rainfall, to absorb the effects of weather changes on social mobility. It also contains reported COVID cases and related deaths in population shares, both domestic and from neighboring countries. These are included to capture their possibly deterring effects on social mobility, which is a possible con-founder of government and individual decisions on distancing. Part D of Table 2.2 presents descriptive statistics of these covariates.

Countries differ in demographics, population density, and the quality of political and healthcare institutions, which are likely correlated with interventions, social activity, and reproduction numbers. I address these differences by including country-fixed effects, assuming the invariance of these factors on weekly frequencies.

2.4 Results

This section presents the results of the first and the second stage estimations. I start with the presentation and discussion of the first stage results. I then continue with a decomposition of social mobility into policy-induced and voluntary components based on first-stage predictions. Finally, I present and discuss the main results about the economic effects of DPIs.

2.4.1 First Stage

Figure 2.6 depicts predicted values for δ_t of equation (2.5), which measures the deviation of the social mobility index m_{it} from its final pre-intervention week values within countries. I interpret the results on this figure from left to right. Effects more than five weeks distant from the intervention are grouped, giving two coefficients: one for the distant past, and one for the distant future of interventions. Results show a slight pre-intervention adjustment in social mobility as the effects from more than five weeks before the first intervention are positive and statistically significant at a 1% level. This effect is most likely attributed to voluntary distancing motives.

In the close neighborhood of the first DPI, pre-intervention coefficients are statistically indistinguishable from zero, while post-intervention coefficients indicate strong mobility-reducing effects of the first DPI. This discontinuity in the results supports the choice of the RD strategy. Most of the post-treatment effects happen within the first two weeks, of which week 0 is a mixed week allowed to contain days both from before and after the day of the first DPI. These are the effects that are identified as the treatment effects of the first DPI. Results predict that the first DPI treatment is expected to reduce social mobility by nearly 20 percentage points in November 2019 levels. Finally, looking at effects more than five weeks after the first DPI shows a slight reversal of social mobility. This reversal is again attributed to changes in voluntary distancing motives.

Figure 2.6: Week-Fixed Effects of Social Mobility around the First DPI

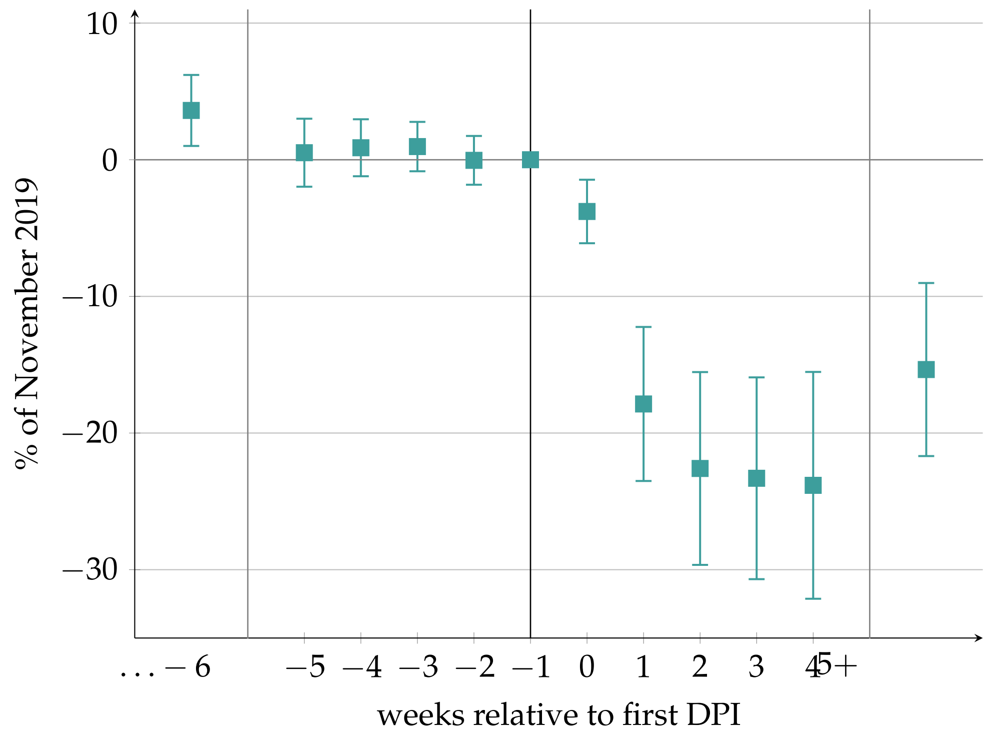

Notes: Point estimates of δ_t of equation (2.5) with 99% confidence intervals. Standard errors are allowed to cluster within weeks. Reference period: last week before the intervention. 2 870 country-week observations of 41 countries, R-squared = 0.7543.

Quantitative results for DPI effects are presented in Table 2.3.[15] Results are presented for three specifications, the first one excluding further changes in the extensity (number of) and intensity of DPIs after the first intervention. The second and third specifications include these two factors gradually. The top two rows show point estimates for δ_0 and δ_1 with their standard errors in parentheses. These two coefficients capture the effect of the first DPI. The mobility effects of the first DPI treatment were found to be robust to specifications. Based on specification 3, the first DPIs reduced social mobility index m_{it} by 17.8 percentage points measured in November 2019 levels. The magnitude of this coefficient is roughly 1.5 of the standard deviation of m_{it} in the pre-treatment sample: 12.0. Given that the average magnitude (number plus level) of the

[15] For the estimation results for other covariates see Section B.2 in the Appendix.

first DPIs was roughly 8, this result also suggests that introducing a single DPI as a first intervention reduces m_{it} by 2.2 percentage points in November 2019 levels.

Table 2.3: Effect of DPIs on Social Mobility

	(1)	(2)	(3)
Week 0	-3.285***	-3.947***	-3.786***
	(1.170)	(1.113)	(1.186)
Week 1	-14.791***	-17.013***	-17.872***
	(3.172)	(2.890)	(2.875)
Extensity		-1.615***	0.545
		(0.359)	(0.490)
Intensity			-2.098***
			(0.320)
Observations	2,870	2,870	2,870
R-squared	0.716	0.728	0.754
Country FE's	•	•	•
Countries	41	41	41
Controls	•	•	•

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within weeks. • – included, ◦ – excluded.

Specification 3 provides no statistical evidence for the effects of changes in the extensity of DPIs. However, it is found to significantly reduce mobility significantly in specification 2, where it has been included without the intensity indicator. A possible explanation for the extensity changes losing their significance when controlled for intensity changes is that these two factors are strongly correlated. The correlation between intensity and extensity is 0.79.

Based on specification 3, changing intensities of already introduced DPIs had a significant mobility-reducing effect. Increasing the intensity (total stringency level) of DPIs after the first treatment by 1 was found to decrease m_{it} by 2.1 percentage points measured in November 2019 levels. This result suggests that, although the first interventions were found to be the most effective, governments could significantly increase the distancing effects of DPIs by increasing their stringency levels. The introduction of new types of restrictions, however, was found to be ineffective in the further enhancement of social distancing.

These results suggest that the first-ever distancing interventions have on average a strong and significant effect on social mobility. This effect is possible to be fine-tuned by changes in the intensity of but not by changes in the extensity of DPIs. The desired reduction of social mobility depends on how strongly these DPI-induced reductions affected the spreading of COVID, which question is outside of the scope of this paper.[16] This paper continues toward its goal of measuring the economic effects of such mobility reductions due to DPIs.

2.4.2 Voluntary Social Mobility

The goal of the first stage estimation was to realize a prediction on voluntary distancing because that is a key control for the identification of the economic effects of DPIs. Voluntary mobility is obtained by residualizing social mobility ($m_i t$) by first-stage predictions for each policy-related covariate.

Predictions for the treatment effect of DPIs is defined as the event-time effects (δ_t) of equation (2.5) from week 0 and 1, such that it is fixed at the value of δ_1 for $t > 1$:

$$\widehat{m}_{it}^{T} = \begin{cases} \delta_t & \text{if } t \in \{0,1\} \\ \delta_1 & \text{if } t > 1 \\ 0 & \text{otherwise} \end{cases}$$

Predictions for further changes in the extensity ($\widehat{m}_{it}^{E}$), and the intensity ($\widehat{m}_{it}^{I}$) of DPIs, and other interventions (P_{it}^{O})are simply the product of these variables with their coefficients:

$$\widehat{m}_{it}^{E} = \gamma^{E} E_{it}, \qquad \widehat{m}_{it}^{I} = \gamma^{I} I_{it}, \qquad \widehat{m}_{it}^{O} = \theta P_{it}^{O}$$

[16]See Perra (2021) for a summary of the related literature.

A prediction for voluntary mobility ($\widehat{m}_{it}^{V}$) is then obtained as the following residual:

$$\widehat{m}_{it}^{V} = \text{mobility}_{it} - \widehat{m}_{it}^{T} - \widehat{m}_{it}^{E} - \widehat{m}_{it}^{I} - \widehat{m}_{it}^{O}.$$

Figure 2.7 depict cross-country averages for the predicted voluntary mobility component ($\widehat{m}_{it}^{V}$), the effect of other interventions ($\widehat{m}_{it}^{O}$) and the sum of the DPI related components: $\widehat{m}_{it}^{T}$, $\widehat{m}_{it}^{E}$, and $\widehat{m}_{it}^{I}$ in calendar time. This is a stacked column graph, therefore the sum of the columns tracks the social mobility indicator m_{it}. Social mobility declines around mid-March as most countries in the sample intervened for the first time in March 2020. The dominant factor in this decline is found to be the effect of DPIs, although voluntary mobility had a substantial contribution as well. Figure 2.8 breaks down the effect of DPIs of Figure 2.7 into its three components: $\widehat{m}_{it}^{T}$, $\widehat{m}_{it}^{E}$, and $\widehat{m}_{it}^{I}$. This figure reveals that the effect of DPIs was predominantly due to the first DPIs. This graph also gives visual support for the conclusion that further changes in the intensity of DPIs had a significant effect on mobility, while further changes in the number of DPIs did not.

The aim of the first stage estimation is to create a proxy for voluntary distancing, which is an important control in the estimation of the economic effects of DPIs in the second stage. It is the predicted voluntary mobility component $\widehat{m}_{it}^{V}$. It is obtained as a residual, and therefore, it is important to investigate which factors drive the variance of this variable. The contribution of different covariates, country-fixed effects, and the error term to the total variance of voluntary mobility, $widehatmV_it$, is shown in table ref: tab: variance. Covariates account for 28 percent to the total variance of $\widehat{m}_{it}^{V}$, with fixed effects accounting for another 12 percent. The unexplained component accounts for the remaining 60% of its variance.

2.4.3 Economic Effects of Distancing Policy Interventions

In this subsection, I present the estimation results for equation (2.4) for seven different economic outcomes. Economic outcomes are measured

Figure 2.7: Historical Decomposition of Social Mobility Index in 42 OECD Economies

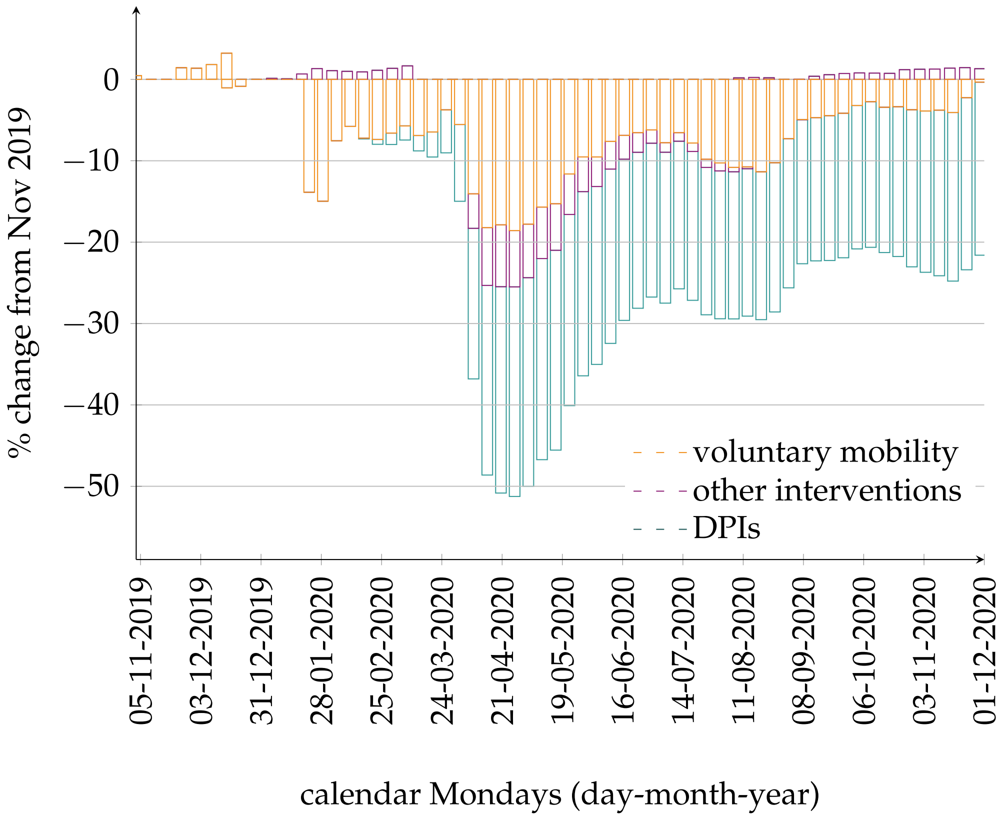

Notes: cross country averages. Predicted by specification (3) of Table 2.3

as a percentage of their most recent November values. The average values of 2015–2019, a COVID-free control period, are subtracted from each indicator. I start with the presentation of three different specifications that include voluntary mobility, other interventions, and mobility at trading partners one-by-one to investigate omitted variable biases caused by these factors. I present these specifications using industrial production as the outcome. After that, I present results for the other six economic outcomes using only the most complete specifications. I first continue with the four sector outputs: industrial, and manufacturing production, construction output, and retail trade. I then continue with the two price indicators: CPI, and PPI in manufacturing. I conclude the analysis with unemployment effects.

Figure 2.8: Historical Decomposition of Predicted DPI effects in 42 OECD Economies

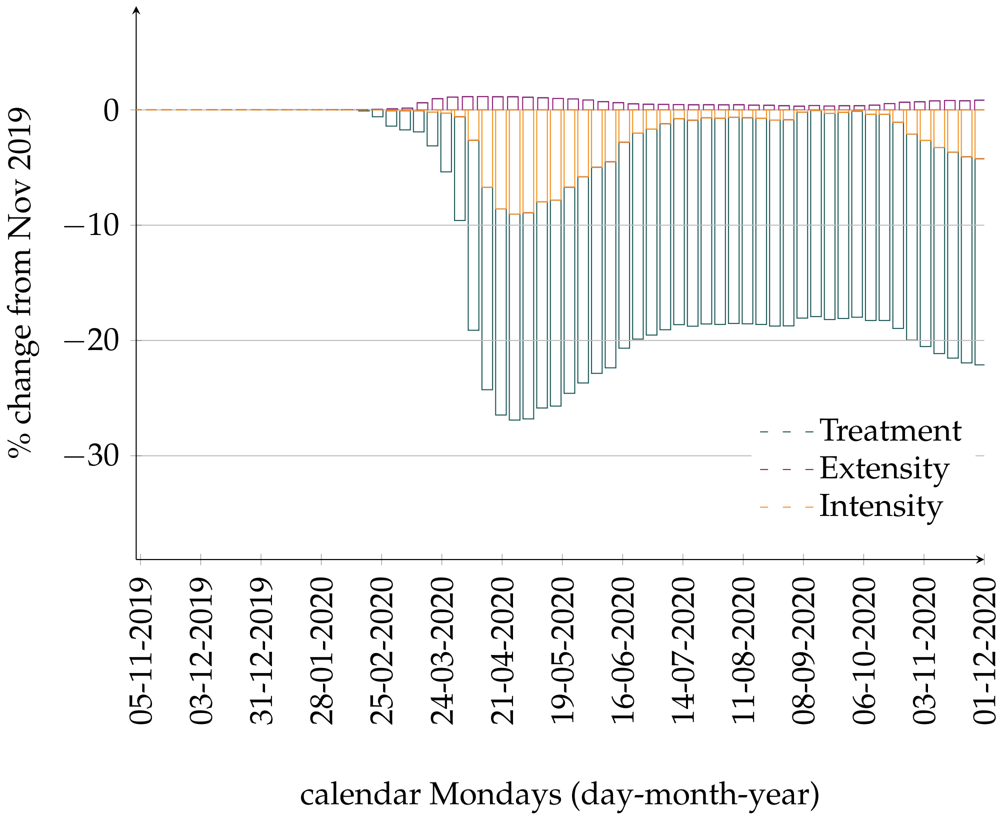

Notes: cross country averages. Predicted by specification (3) of Table 2.3

Second Stage Results: Industrial Production

Results for industrial production are presented in Table 2.5.[17] The first specification includes only DPI factors along with country-fixed effects. Treatment and intensity effects are already significant and strong in this simple specification. The second specification includes voluntary social mobility, which is found to be significant and positively correlated with industrial production. The positive sign of this coefficient is in line with the intuition that less mobility implies lower rates of economic activity. The third specification includes other interventions. These turn out to be important control factors, as their inclusion significantly decreases the coefficients of DPI factors. The fourth specification includes mobility at export and import partners. The inclusion of these two indicators decreases slightly further the coefficients of DPIs, revealing a modest omitted variable bias in previous specifications due to international spillovers. The coefficients of these two factors are statistically insignifi-

[17] For the estimation results for other covariates see Section B.3 in the Appendix.

Table 2.4: Variance Decomposition of Voluntary Mobility

Covariate	Variance	Proportion (%)
Average Temperature	2.19	2.21
Average Humidity	1.01	1.03
Average Rainfall	0.14	0.14
Average Snowfall	0.00	0.00
Covid Cases $_{t-1}$	0.03	0.03
Covid Cases $_{t-2}$	0.05	0.05
Covid Deaths $_{t-1}$	6.73	6.80
Covid Deaths $_{t-2}$	2.65	2.67
Covid Cases $_{t-1}$ at Neighbors	0.58	0.59
Covid Cases $_{t-2}$ at Neighbors	3.96	4.00
Covid Deaths $_{t-1}$ at Neighbors	9.58	9.68
Covid Deaths $_{t-2}$ at Neighbors	0.42	0.42
Total of Covariates	27.34	27.62
FE_i	12.21	12.33
Residual	59.43	60.05
Total	98.98	100.00

Notes: Using on results from specification (3) of Table 2.3. Covid cases and deaths are measured in population shares both for domestic and neighboring countries.

cant, however.

The most complete specification shows that the first DPIs and further changes in their intensity had a significant effect on industrial production, while I found no evidence for the effects of changes in the extensity of DPIs. When compared to its 2015-2019 averages in November values, a single level 1 DPI reduces industrial production by 0.8 percentage points on average. A change in the intensity of DPIs reduces industrial production by 1.15 percent. A one percentage point deviation of voluntary mobility from its November 2019 levels is found to decrease industrial production by .4 percent.

Table 2.5: Effect of DPIs on Industrial Production

	(1)	(2)	(3)	(4)
Treatment	-1.230***	-1.124***	-0.871***	-0.793***
	(0.203)	(0.132)	(0.244)	(0.231)
Extensity	-0.687	-0.723	-0.628	-0.679
	(0.524)	(0.625)	(0.533)	(0.559)
Intensity	-1.809***	-1.559***	-1.223***	-1.157***
	(0.447)	(0.282)	(0.242)	(0.217)
Voluntary Mobility		0.335**	0.397**	0.400**
		(0.130)	(0.131)	(0.137)
Mobility $_{t-1}$ at Import Partners				0.411
				(0.702)
Mobility $_{t-1}$ at Export Partners				-0.285
				(0.662)
Observations	288	288	288	288
R-squared	0.584	0.623	0.666	0.669
Country FE	•	•	•	•
Countries	32	32	32	32
Other Interventions	○	○	•	•

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within months.
• – included, ○ – excluded.

Output Losses

Table 2.6 shows the results of specification 4 for the four different sector outputs.[18] Column 1 simply repeats the results for industrial production in column 4 of table 2.5 for comparison. Column 2 is manufacturing, which is a sub-sector of the wider industry sector, and as a consequence, the results are very similar in the first two columns. The first DPI treatment and further changes in DPI intensities had a significant negative effect on manufacturing production, but there was no effect from the changes in extensity. Voluntary mobility had a similar impact on manufacturing as it did on the entire industry sector, and no evidence of

[18]For the estimation results for other covariates see Section B.3 in the Appendix.

spillover effects from mobility changes at trading partners was found.

Table 2.6: Effect of DPIs on Sector Outputs

	(1) Industrial Production	(2) Manuf'ing Production	(3) Constr' Output	(4) Retail Trade
Treatment	-0.794***	-0.957***	-2.141**	-1.234***
	(0.233)	(0.202)	(0.753)	(0.223)
Extensity	-0.670	-0.302	3.039**	0.938
	(0.557)	(0.589)	(0.941)	(0.695)
Intensity	-1.145***	-1.413***	-2.682***	-2.447***
	(0.218)	(0.202)	(0.436)	(0.473)
Voluntary Mobility	0.402**	0.454**	0.301	0.531***
	(0.136)	(0.143)	(0.208)	(0.112)
Mobility $_{t-1}$ at Import Partners	0.405	0.525	0.126	0.515
	(0.702)	(0.758)	(0.882)	(0.416)
Mobility $_{t-1}$ at Export Partners	-0.278	-0.388	-0.222	-0.552
	(0.663)	(0.723)	(0.947)	(0.424)
Observations	288	288	189	270
R-squared	0.670	0.687	0.683	0.784
Country FE	•	•	•	•
Countries	32	32	21	30
Other Interventions	•	•	•	•

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within months. • – included, ○ – excluded.

Results for construction output are presented in column 3. This indicator is only available for a substantially smaller set of countries; therefore, its results are not quantitatively comparable with other columns. The introduction of a single level 1 DPI treatment decreases construction output by 2.1 percentage points from November 2019 levels. A further change in DPI intensity is found to decrease it by another 2.7 percentage points. A unit increase in DPI extensity, which is the introduction of a new type of DPI, had, however, a positive, albeit only marginally significant effect on construction output. Without further investigation, a possible explanation could be that when DPI restrictions extend to more and more building types, such as schools, office buildings, or concert halls, that gives way to more and more reconstructions. A one percentage point decline in mobility was found to decrease construction output by 0.3 percentage points. No evidence was found for the spillover effects of mobility changes across trading partners.

Column 4 shows results for retail trade. Retail trade responded significantly to the first DPI treatment, changes in DPI intensity, and voluntary mobility. I found no evidence of significant responses to DPI extensity and mobility spillovers from trade partners. A single level 1 DPI introduced as a first treatment decreases retail trade by 1.2 percentage points, while a unit change in the level of DPIs decreases retail trade by 2.4 percentage points from November 2019 levels. In November 2019 values, a 1% decrease in voluntary mobility reduces retail trade by 0.5 percentage points.

Sector outputs were found to respond strongly to first DPI treatments, changes in DPI intensities, and voluntary mobility. On the other hand, I found no evidence of significant responses to DPI extensity and mobility spillovers from trade partners, except in construction. These findings altogether suggest that distancing behaviors that were either voluntary or DPI-compliant generated substantial output losses.

It is crucial to compare the consequences of voluntary and DPI-induced distancing when forming policy conclusions about DPI efficiency. Voluntary mobility and DPI components are measured in different units, so Table 2.6 coefficients are not directly comparable between rows. One possible way to address this issue would be to use the estimates for DPIs of equation (2.5), for example, $\widehat{\gamma}^I I_{it}$ for intensity changes, directly on the right-hand side of equation (2.4), instead of the policy variables. $\widehat{\gamma}^I I_{it}$ contains the same information as the policy variable, I_{it}, as they differ only in a constant multiplier $\widehat{\gamma}^I$. But this multiplier translates the unit of the policy variable into the unit of voluntary mobility changes, making these two factors comparable. Although this strategy appears simple and straightforward, it is impossible to implement because the most important policy variable, treatment (T_{it}), is not included in the first-stage equation. The reason it is not included is that the effect of the DPI treatment is captured by the RDiT design that builds on the key identifying assumption of sudden responses to policy changes. Giving this design

Table 2.7: Predicted Distancing Effects by Sectors in Month 2 of the First DPI

	Industrial Production	Manuf'ing Production	Constr' Output	Retail Trade
Treatment	-10.17	-12.25	-27.41	-15.80
	(2.99)	(2.59)	(9.64)	(2.85)
Extensity	-0.22	-0.10	1.01	0.31
	(0.18)	(0.19)	(0.31)	(0.23)
Intensity	0.16	0.20	0.38	0.35
	(-0.03)	(-0.03)	(-0.06)	(-0.07)
Voluntary Mobility	-1.44	-1.62	-1.08	-1.90
	(-0.48)	(-0.51)	(-0.74)	(-0.40)
Total Change	-16.15	-19.20	-19.52	-11.68
Explained by Distancing	-11.67	-13.77	-27.10	-17.04
percent	72.26	71.72	138.83	145.89
Unexplained by Distancing	-4.48	-5.43	7.58	5.36
percent	27.74	28.28	-38.83	-45.89

Notes: Predicted effects. Calculated as changes in cross-country averages between month -1 and month 2, and multiplied by the coefficients of column 4 of Table 2.6. Standard errors in parenthesis are calculated similarly, using the s.e. of the corresponding coefficient.

up is considered to be a greater cost than the gain of the comparison that would emerge from a different design would provide.

I work around this problem by picking a different strategy to make the effects of DPIs and voluntary distancing comparable. It is a decomposition of the changes in sector outputs around the months of the first DPI interventions. I calculated predicted values of DPI and voluntary distancing effects by multiplying the changes of these factors from month -1 to month 2 for all factors with their coefficients.[19] Table 2.7 shows these predicted effects for all four sector outputs averaged across countries. The bottom of the table contains the change of the explained sector outcome and summary calculations about what fraction of this total change could be explained by the predicted distancing effects. Figures show that although voluntary distancing caused significant losses to sector

[19] I did the same multiplication with the standard errors.

outputs, its effect was an order of magnitude smaller than that of DPIs in the short run. For example, the first DPI treatment explains about 10 percentage points of the total industrial output loss relative to the last month before the first DPI.[20] Voluntary distancing, on the other hand, explains only 1.4 percentage points.

The largest negative effect of the first DPI treatment was identified in the construction sector, -27.4 percentage points. The effect of DPI extensity was found to be significant only in the case of construction output, where it contributed 1 percentage point, offsetting slightly the overall 19,5 percentage point decline observed in the sector. In month 2, the effect of DPI intensity changes was found to contribute the least to total changes. Voluntary mobility was found to decrease retail trade the most, by almost 2 percentage points.

Only 70% of the total losses in industry and manufacturing are explained by distancing factors, implying that output losses in these sectors were caused by other factors, such as other COVID-related interventions. In construction and retail trade, on the other hand, distancing factors altogether predicted more losses than was observed. This finding suggests that other factors, such as fiscal and monetary support programs, could mitigate the short-term costs of distancing in these two sectors.

Inflationary Effects

Table 2.8 contains results for consumer prices and producer prices of the manufacturing industry.[21] I found no evidence of inflationary effects of DPIs except for DPI extensity. Results show that extending the set of DPIs by a new intervention decreases consumer prices by 0.1 percent. Column 6 shows results for producer prices in manufacturing, providing no evidence of either voluntary or DPI-induced distancing effects from domestic markets. Distancing in export markets, on the other hand, is marginally significant, with a one-point increase in social mobility in export markets lowering domestic manufacturing prices by 0.27 percent. The sign of this effect is in contrast with the economic intuition that falling demand reduces prices.

[20]As a comparison Deb et al. (2021) find that losses in industrial production were about 10 percent over 30 days following the implementation of containment measures.

[21]For the estimation results for other covariates see Section B.3 in the Appendix.

In summary, evidence for the inflationary effects of any kind of distancing could not be identified by this study. This suggests that neither DPIs nor voluntary distancing brings on little to no inflationary costs. One possible explanation for insignificant inflationary effects is the omission of conventional monetary policy interventions, such as rate cuts. Countries anticipating stronger inflationary risks due to their specific mix of DPIs might have cut their rates more strongly, mitigating the inflationary effects of DPIs. Another possible explanation for insignificant inflationary effects is that the typical shock response of prices tends to have a several-quarter time lag. That might suggest that inflationary effects of DPIs emerge on time horizons, for example, a year later, that are unable to be captured with the current design.

Table 2.8: Effect of DPIs on Prices

	(1)	(2)
	CPI	PPI manuf'ing
Treatment	-0.009	-0.049
	(0.009)	(0.049)
Extensity	-0.090**	0.178
	(0.028)	(0.157)
Intensity	0.016	-0.077
	(0.015)	(0.058)
Voluntary Mobility	-0.005	-0.008
	(0.003)	(0.022)
Mobility $_{t-1}$ at Export Partners	-0.010	-0.275*
	(0.023)	(0.133)
Mobility $_{t-1}$ at Import Partners	0.024	0.231
	(0.025)	(0.139)
Observations	288	252
R-squared	0.887	0.838
Country FE	•	•
Countries	32	28
Other Policies	•	•

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within months. • – included, ◦ – excluded.

Unemployment Effects

Table 2.9 presents second-stage results for the unemployment rate in the same four specifications as Table 2.5 for industrial production.[22] The first specification reveals strong positive unemployment responses to the first DPI treatment.[23] When voluntary mobility is introduced as a control, this strong response is maintained. However, when other COVID interventions have been introduced the coefficient of DPIs collapsed and lost its significance. This observation about coefficients is maintained when mobility at trading partners is included in the final specification. This finding suggests that the observed hike in unemployment on Figure 2.1 of Section 2.2 after the first DPI interventions are explained by other COVID related interventions.[24]

There might be other explanations for the lack of unemployment effects as well. For example, conventional or not COVID-focused fiscal policy interventions are not controlled for in the current version of this paper. For example, governments anticipating higher unemployment risks might choose to relax taxes more intensively compared to other governments, offsetting the unemployment effects of their DPIs. Another possible explanation could be based on anecdotal evidence that labor adjusted on the intensive margins first in the early months of the COVID restrictions, such as lowering working hours, the extension of sick leaves, or enforced holidays. Employers aimed to keep their employees as they expected the restrictions-induced production halt to be temporary and considered the cost of rehiring to be higher than the cost of labor intensity adjustments.

2.5 Conclusion

This paper identifies the causal effects of distancing policy interventions (DPIs) on seven short-term economic indicators: industrial and manufacturing production, construction output; retail trade; CPI; PPI in manufacturing; and the unemployment rate. Effects are identified

[22]For the estimation results of other covariates from specification 4 see Section B.3 in the Appendix.
[23]And a slightly negative response to DPI intensity changes.
[24]For a similar result, see Kong and Prinz (2020).

Table 2.9: Effect of DPIs on the Unemployment Rate

	(1)	(2)	(3)	(4)
Treatment	0.163***	0.171***	0.058	0.026
	(0.031)	(0.029)	(0.069)	(0.070)
Extensity	-0.133*	-0.133*	-0.076	-0.043
	(0.065)	(0.061)	(0.074)	(0.060)
Intensity	0.059	0.073	0.019	-0.006
	(0.100)	(0.095)	(0.090)	(0.084)
Voluntary Mobility		0.017	-0.007	-0.013
		(0.011)	(0.009)	(0.009)
Mobility $_{t-1}$ at Export Partners				-0.108*
				(0.057)
Mobility $_{t-1}$ at Import Partners				0.008
				(0.065)
Observations	279	279	279	279
R-squared	0.745	0.747	0.788	0.810
Country FE	•	•	•	•
Countries	31	31	31	31
Other Interventions	○	○	•	•

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within months.
• – included, ○ – excluded.

from within-country changes of these indicators from before and after the first ever DPI treatment relative to the averages of a COVID-free control period: 2015–2019. Causal effects are identified by controlling for three important confounding factors: voluntary distancing, other COVID-related interventions, and distancing at trading partners.

Among these confounders, voluntary distancing is an unobserved factor. Voluntary distancing is therefore estimated in a regression discontinuity framework using mobility data. It is realized as a residual after the identification of DPI-induced distancing effects as sudden changes in mobility after the first DPI intervention. Results suggest that the first-ever distancing intervention had, on average, a strong and significant effect on social mobility. This effect was fine-tuned by changes in the intensity of DPIs. Voluntary motives were also found to contribute to a significant portion of mobility patterns.

I found significant output losses due to DPIs, but no evidence for inflationary and unemployment effects. Findings suggest that DPIs caused substantial output losses. Results also show that although voluntary distancing caused significant losses to sector outputs, its effect was an order of magnitude smaller than that of DPIs. Only 70% of total losses in industry and manufacturing are explained by either voluntary or DPI-induced distancing, implying that other factors, such as other COVID-related interventions contributed substantially to output losses in these sectors. In construction and retail trade, on the other hand, distancing factors altogether predicted more losses than was observed. This finding suggests that other factors, such as fiscal and monetary support programs, could mitigate the short-term costs of distancing in these two sectors.

This study did not identify any evidence for the inflationary effects of any kind of distancing. This suggests that neither DPIs nor voluntary distancing brings on little to no inflationary costs. Although a significant hike in unemployment rates can be observed after DPI interventions took place, no evidence was found in support when controlling for voluntary distancing, other COVID interventions, and distancing at trading partners. Findings suggest that the observed hike in unemployment is related to other COVID-related interventions.

These findings provide evidence of the economic cost of DPIs to consider for governments that are planning to implement such interventions during an epidemic. These findings also contribute to a more complete cost-benefit analysis of distancing policy interventions on the cost side. The costs identified here are mainly output losses, while no evidence was found for inflationary costs or unemployment responses.

Chapter 3

Network Origins of Aggregate Dynamics

3.1 Introduction

There is growing evidence of producer-level idiosyncratic shocks affecting macroeconomic aggregates as they are propagated and amplified by input-output linkages, i.e., the production network. For example, an oil shortage affects oil refineries, the production of fertilizers, the energy sector, plastic manufacturing, chemical production, and all other sectors that directly or indirectly rely on crude oil as an input. In this paper, I measure the speed of this propagation process by pinning down the average time it takes for the effect of a shock to propagate from an input supplier to its direct customers.

Quantifying the average propagation time of an economy is the main contribution of this paper. Most papers in the production networks literature, kicked off by Acemoglu et al. (2012), assume that this shock propagation process is instantaneous, or at least goes through entirely[1] within a time period. The seminal model of this literature, Long and Plosser (1983), however, was able to generate dynamic propagation. They assumed that producers could only transmit the effect of a shock with a one-period delay. Although their model is dynamic, they did not specify the length of a time period. This paper applies their model predictions to annual time series and input-output data of 66 private industries in the

[1] Including all direct and indirect customers of any distance.

US economy from the past two decades to quantify the average propagation time, which is equivalent to the delay of shock transmission in Long and Plosser (1983).

Once quantitative results for the average propagation time are established, it is possible to investigate how the staggered propagation of producer-level shocks affects aggregate dynamics. Applying these new results to the model, I predict a series of aggregate GDP and calculate its auto-correlation function for the third degree. I then determine what proportion of aggregate dynamics can be traced back to network propagation by comparing the predicted auto-correlations to their observed counterparts. This secondary result speaks to the network origins of aggregate dynamics, which is another important contribution of this paper.

The rest of the paper is organized in the following way: Section 3.2 introduces the model of Long and Plosser (1983) and presents the main predictions that are employed in the empirical analysis. This model is a real business cycle model with mostly standard assumptions: markets are perfectly competitive, prices are fully flexible, and the demand side is represented by a single utility-maximizing household. It has two alterations to the text-book type RBC model: It assumes (i) multi-sector production with (ii) time-to-build friction.

The multi-sector nature of the model means the following: Each sector is represented by a single profit-maximizing producer that uses labor and intermediate goods as inputs to produce a single good. Each sector's output can be used either as a final consumption good or as an intermediate input. The web of intermediate input-output linkages defines the production network.

The time-to-build friction is responsible for the dynamics of the model. It assumes that inputs have to be purchased one period earlier than their products can be sold on output markets.[2] The only shocks the model includes are TFP shocks. The model assumes no savings, capital accumulation, or any means of consumer-side dynamics. Therefore, dynamics are solely driven by the network propagation of TFP shocks, which is paced by the time-to-build friction.

[2]Intuition would suggest that this lagged sourcing of inputs is more plausible at higher time frequencies, such as monthly or weekly compared to annual. Instead of inferring the validity of this assumption from intuition, this paper pins down the average time delay of shock propagation, which is the consequence of the time-to-build assumption in this model.

The most important feature of the model is the staggered propagation of TFP shocks through input-output linkages. This propagation process unfolds the following way: Each industry responds to its own TFP shocks by the adjustment of its output price and quantity within the time period of the shocks. This way, the effect of the TFP shock is transmitted to customer industries as they have to adjust their input demands to these new input prices. They do not pass the effects of the original shocks by adjusting their output prices in this period, however, because of the time-to-build friction: inputs purchased in this period produce output only a period later. Therefore, the effect of the original TFP shocks appears in the prices of first-order customer industries one period later. And the propagation does not stop at direct customers, as now in the second period, the buyers of direct customers are faced with adjusted input prices. They also have to adjust their own input demands and wait one period before passing through the effects to their customers because of time-to-build. For example, a negative TFP shock, for example, workers' strikes, in coal mining reduces the output of this sector, which makes coal products more expensive. Industries that use coal as an input, including the energy sector, respond by decreasing their demand for coal, which reduces their own output and thus increases their output price as well. This process spreads the effect of the coal mining TFP shock to energy users, thus indirect customers of coal mining. This process then goes on until it reaches the most distant indirect customers of coal mining, thus diffusing to the whole economy through the production network. But most importantly, the propagation process takes one step in one time period when it walks down the production network of input-output links.

The model summarizes this shock propagation process by a single equation that connects industry value-added growth rates to TFP shocks. This equation predicts that value-added growth rates are the linear combination of contemporaneous and lagged TFP shocks combined by the production network. This combination captures the network propagation of shocks that have been described above: contemporaneous TFP shocks have an effect on their own industries only. TFP shocks lagged by one period are affecting direct customers of the shocked industries. Similarly, the kth lag of TFP shocks is affecting the kth order customers.

The central problem of this paper is that the average propagation time,

10.14754/CEU.2022.06

i.e. the time frequency of the model, is unknown. It is pinned down using annual data, which is described in Section 3.3. This study obtains annual value-added, TFP, and input-output data from the U.S. Bureau of Economic Analysis (BEA). The estimation sample is a panel of the 66 private sector breakdown of the US economy. The sample spans each year between 1997 and 2020.

The empirical strategy for the identification of the average propagation time is presented in Section 3.4. The average propagation time is the time unit of the model's time frequency, which is assumed to be at most one year and noted by δ. The main empirical challenge of the paper is the identification of a time interval that is at most as long as the unit of the frequency of the observations. I tackle this challenge with the following strategy: I evaluate the model at sub-annual frequencies. I assume that one of these sub-annual time frequencies is the true frequency of the model. The primary goal of the empirical strategy is to find that frequency. I predict the observed annual value-added growth rates using TFP shocks and input-output coefficients at each of these possible sub-annual frequencies. The sub-annual frequency that minimizes the sum of the squared errors between model-predicted and observed annual value-added growth rates is identified as the most likely time frequency of the model. The average propagation time δ is identified by the unit time length of this error-minimizing time frequency.

This empirical strategy is evaluated in two different specifications. The first one uses the model predictions faithfully. That means this specification includes both contemporaneous and lagged TFP shocks when it evaluates the model. Contemporaneous TFP shocks – predictions themselves – are likely to be strongly correlated with value-added growth rates of their own sectors, raising a threat of endogeneity in the first specification. In a second specification, I use only lagged TFP shocks therefore in the identification of δ.

The second part of this section discusses the empirical strategy for determining what proportion of aggregate dynamics can be traced back to network propagation. I take model predictions for GDP growth evaluated at the previously pinned down value of propagation time δ. I then compare the first, second, and third-order auto-correlation coefficients of the predicted and actually observed GDP growth rates.

Results are presented in Section 3.5 in three parts. The first part shows results for the average propagation time, δ. I find that δ was most likely between 4 and 8 months in the US in the past two decades.

This propagation time is sufficiently slow to consider the network propagation of sector-level shocks a dynamic process even at annual frequencies. In other words, it is likely that industry-level shocks exert their impact on the economy much beyond the time horizon of a year. Consequently, the network propagation of these socks themselves might generate persistent macroeconomic aggregates. Thus, the second section predicts annual GDP growth rates, assuming the average propagation time to be 6 months. I then compare the predicted and observed GDP growth rate series using their first three auto-correlation coefficients. I do this to calculate the proportion of GDP dynamics that can be traced back to dynamic network propagation. I find that dynamic network propagation can account for 82 percent of the first order auto-correlation of annual GDP growth. This result provides evidence for the network origins of aggregate dynamics.

This finding is reinforced when compared to predicted GDP auto-correlations under the assumption of instantaneous network propagation, ($\delta = 0$). Under that assumption, GDP growth rates are predicted to be only weakly auto-correlated, if at all. This finding provides further evidence for the network origins of GDP dynamics. In the third part of this section I estimate δ separately for pre- and post-2008 sub-samples. I find that the network propagation of TFP shocks has likely accelerated after the Great Recession. I conclude my findings in the final Section.

Literature

This paper primarily contributes to the literature on production networks. Following the study of Acemoglu et al. (2012), there is a renewed interest in this research topic. Since then, many papers have investigated the relevance and consequences of the network propagation of micro-level shocks. The paper of Acemoglu et al. (2012) shows that producer level shocks are amplified and propagated into macroeconomic aggregates, explaining a significant portion of aggregate volatility. This happens because of the particularly unbalanced structure of the production network, in which there exist a few disproportionately influential input supplier sectors, such as the oil sector, that are important direct or indirect suppliers to most of the producers in an economy. Shocks to these sectors therefore have a large effect on the economy proportional to the relative size of these sectors in terms of output.

The model of Acemoglu et al. (2012) is generalized to imperfect competition by baqaee2018cascading, allowing for firm entry and exit. Acemoglu et al. (2010) generalizes the model of Acemoglu et al. (2012) beyond first-order approximations, finding that idiosyncratic producer-level shocks might trigger a cascade of failures of businesses. According to Baqaee and Farhi (2019), producer level shocks account for much higher levels of aggregate volatility than was found with a linear approximation and account for a 20% reduction in TFP growth rates in the second half of the last century.

Carvalho and Gabaix (2013) shows that the Great Moderation period between 1984 and 2007 that was named after the characteristically low levels of aggregate volatility could be explained by the diversification of production technologies that led to the shrinkage of the influence of the previously most influential suppliers. They also show that the buildup in the influence of the financial sector contributed substantially to the high volatility time period that started with the Great Recession in 2008. This is an incomplete list of the theoretical research done in the field of production networks, which has been extended by many others in recent years. A good summary of further studies can be found in Carvalho (2014). This paper investigates empirically if this network propagation process has a time dimension to it. Allowing for dynamic propagation lets me address whether aggregate dynamics have network origins.

Empirical evidence for the relevance of production networks was provided by several papers. For example, Carvalho et al. (2021) finds significant network propagation effects using the natural experiment of an earthquake in Japan in 2011. Acemoglu et al. (2016) shows that supply-side shocks, such as TFP shocks, propagate downstream through the production network, whereas demand-side shocks, such as external demand shocks, propagate upstream through the production network. This finding is in line with the predictions of the seminal model. Barrot and Sauvagnat (2016) shows that input specificity is "an important determinant of the propagation of idiosyncratic shocks in the economy." The model in this paper retains all the features of the original timeless model that explain these findings because the propagation process is based on the same principles. It, however, stretches this propagation process out in the time dimension, making the propagation staggered. This paper contributes to these empirical results with a quantitative measurement of the time step of this network propagation process.

There are other papers that study the dynamic effects of network propagation of producer-level shocks. Carvalho and Grassi (2019) model an economy in which large firm dynamics propagate through input-output links, generating aggregate dynamics. A major difference from their approach is that in this paper, the sole source of dynamics is the result of the staggered transmission of shocks through the links of input supply. In this setup, even shocks with zero auto-correlations generate aggregate dynamics. Another example is the recent work of Liu and Tsyvinski (2020). Their model contains a continuous version of Long and Plosser (1983), which model is employed in this paper. Consequently, their model can also predict similar results to those that this paper uses in the empirical analysis. This paper contributes to their findings by taking the first step towards empirical evidence for dynamic network propagation and the network origins of aggregate dynamics.

This paper quantifies the average propagation time of the US economy, which is the average time it takes for the effect of a cost side shock to appear in the output prices of a sector. Therefore, this paper

This is related to the literature on the frequency of price changes and on the pricing responses to shocks. Papers find that most firms change their prices at least once a year. These figures are found to be similar in most advanced economies, based mostly on survey-level evidence. See, for example Kwapil et al. (2010) for Austria, Hoeberichts and Stokman (2010) for the Netherlands, Martins (2010) for Portugal, or Fabiani et al. (2005) or Druant et al. (2012) for the EU. Kwapil et al. (2010) also finds that Austrian firms respond to cost-side shocks with an average time lag of 3 to 6 months. These findings are in the same ballpark as the main finding of this paper on the average propagation time being 6 months in the US. It also has to be mentioned that this propagation time could be dependent on the level of aggregation, as it might take several firm-to-firm propagation steps before the effect of a shock crosses the boundaries of a sector.

Replicating persistent macroeconomic series is an ongoing challenge for macroeconomic models, as it was emphasized by Cogley and Nason (1995), who showed that the main reason for insufficient persistence is the weak internal propagation of these models. Chari et al. (2000) shows that assuming sticky prices does not solve this persistence puzzle. This paper shows that the staggered propagation of shocks in the production network, as suggested by Long and Plosser (1983) can account for a large portion of the observed persistence in macroeconomic series with an average propagation time of 6 months.

Finally, my findings provide the basis for new opportunities in economic policy. Grassi and Sauvagnat (2019) shows how to use input-output data and production network models to draw predictions for different domains of policy, including industrial or fiscal policies. The primary policy gain from these models is the opportunity for better-targeted interventions. This paper provides evidence for the dynamic propagation of sector-level shocks. This finding extends the targeting opportunities of policy interventions in the time dimension. This extension is particularly important in the prevention of welfare-discounting consequences of economic shocks.

3.2 Model

The first part of this Section presents the model setup. In the second part, I present two model predictions. The first one relates industry output to sector-level TFP shocks. This equation summarizes shock propagation and is therefore utilized in the empirical identification of the average propagation time. The second prediction shows how GDP aggregates from the industry-level output. This enables the model to predict the GDP series, which can be used to measure the GDP auto-correlation resulting from the dynamic propagation of sector-level TFP shocks. These auto-correlations can then be compared to actual GDP auto-correlations to determine the empirical significance of dynamic network propagation.

In the model notation, I use τ for time indexes to save t for annual observations. The unknown parameter in this paper is the time unit of the discrete-time scale indexed by τ. I assume that τ is a discrete divisor of a year, for example, quarterly to retain a tractable correspondence between τ and t. This is a strong assumption that constrains the empirical investigation into the domain of the following possibilities: annual, semi-annual, three periods a year, quarterly, and so on until, for example, the 365th of a year, i.e., daily. This domain has a low resolution at its early end, which is a limitation of the current version of this paper. This limitation is especially problematic considering the result is expected to be in the ballpark of a few months. As a result, in a future version of the paper, this restriction will be relaxed by allowing the unknown time scale τ to have any possible frequency.

3.2.1 Setup

This paper builds on the model of Long and Plosser (1983), which is presented in the following. Consider an economy of an infinitely lived representative consumer and n industries. Each industry is represented by a producer of a single good. These products are sold in competitive markets with flexible prices. Each product is allowed to be used two ways, either as a final consumption good or as an intermediate input by any producers.[3]

A production side time-to-build friction constrains producers to source their inputs one period ahead of production. This friction is responsible for the dynamic propagation of sector-level shocks. Because when an input supplier is hit by a shock, its customer industries need to adjust their output and prices to this shock when their production is ready, which happens one period later then their input purchases. A crucial unknown factor in this process is how long is one period in this model? This is the quantitative question this paper is focused on.

[3]Including the producer of this input itself.

Preferences

The representative consumer derives its period utility from the consumption of a variety of n products in each time period τ. It has Cobb-Douglas preferences over this variety of goods:

$$C_\tau = \prod_{i=1}^{n} c_{i\tau}^{\gamma_i} \tag{3.1}$$

where $c_{i\tau}$ is consumption of product i in time-period τ. $\gamma_i \in [0, 1)$ is the corresponding preference parameter, such that $\sum_i \gamma_i = 1$, which implies that γ_i equals the consumption share of product i. It is assumed that the consumer has no means of savings or investments, which rules out any sort of consumer-side dynamics. This way, the only source of dynamics is the result of the dynamic propagation of supply-side shocks. Without savings or investment possibilities, the consumer consumes the entirety of their income each time period. The consumer earns its income by providing its unit labor endowment to producers.

Production

Each product is produced by a separate industry. Because each industry is represented by a single producer i, the terms "industry," "sector," and "producer" are used interchangeably throughout this text. Each product is either purchased by the consumer or used as an intermediate input by another producer, or the producer itself. This web of input-output linkages defines the production network of the economy. Each producer i combines its labor and intermediate inputs using the following Cobb-Douglas technology:

$$y_{i\tau} = e^{z_{i\tau}} l_{i\tau-1}^{\lambda_i} \prod_{j=1}^{n} x_{ij,\tau-1}^{a_{ij}} \tag{3.2}$$

where $y_{i\tau}$ is the total output of sector i, $e^{z_{i\tau}}$ is a Hicks-neutral total factor productivity shock, $l_{i\tau-1}$ is the labor input, and $x_{ij,\tau}$ is intermediate input produced by sector j. Assuming constant returns-to-scale, the coefficients λ_i and a_{ij} represent expenditure shares of the respective input, and $\lambda_i + \Sigma_j = 1na_ij = 1$. The collection of a_ij coefficients represents the production network of this economy, which is defined in the following definition:

Definition 3.1. *Production Network. The production network of the economy is defined as the directed graph of links a_{ij} connecting each industry j to i. The strength of an edge a_{ij} represents the share of intermediate input j in the total expenditure of industry i. The production network is fully characterized by its adjacency matrix A, which is the matrix of a_{ij} coefficients:*

$$A = [a_{ij}]_{ij}. \tag{3.3}$$

Producer i sells its output either as final consumption $c_{i\tau}$ or as an intermediate input to another producer j as $x_{ij\tau}$.[4] Profits are realized from these sources of revenues after the deduction of input costs. Producers maximize these profits.

TFP shocks

Producers are prone to TFP shocks, which they realize and adjust to before going to the output markets. TFP shocks $e^{z_{i\tau}}$ are assumed to be random walks in logarithms, therefore TFP growth rates $\Delta z_{i\tau}$ are assumed to be mean zero i.i.d. processes for each industry i. These shocks are the only source of stochastics in this model.

Time-to-Build Friction

A crucial feature of the production technology is that the output $y_{i\tau}$ of time period τ is produced by inputs purchased a period earlier: in $\tau - 1$. This assumption covers all inputs, including labor.[5] This friction is referred to as the time-to-build assumption following Kydland and Prescott (1982). This friction is stated formally in the following assumption:

[4] or to itself as $j = i$ is allowed.

[5] Changing the model to contemporaneous labor sourcing would have no effect on the predictions used in this paper. In this model, there is a single labor market for all producers each period. The chosen numeraire is the period real wage w_{t-1}, which is set to 1 and thus drops from log forms. With contemporaneous sourcing of labor, $w_t = 1$ could be set, resulting in the same equilibrium conditions.

Assumption 3.1. *Time-to-build. Inputs purchased in time period t produce output in $t+1$. Formally, $y_{i\tau+1}$ is produced by $l_{i\tau}$ and $x_{ij\tau}$ for $\forall\, i = 1 \ldots n$.*

Labor input is sourced with a one-period lag similarly to intermediate inputs.[6] This is assumed in Long and Plosser (1983) and followed here. The intuitive interpretation of this production structure is that once each input is purchased (including the hiring of labor) the production process is ready to start, but it takes time before it is finished and ready for sales.

Equilibrium

There are n markets for goods and one for labor per time period τ. Prices are flexible; competition is perfect in all markets. General equilibrium is defined by the set of prices, final consumption, intermediate input purchases, wages, and employment that clear each and every market.

The shock propagation process unfolds in equilibrium in the following way. Each industry adjusts its output price and quantity in response to its own TFP shocks. It does not adjust its input demand because the time-to-build friction means its inputs were already purchased in the preceding period.[7] Thus, the effects of TFP shocks appear only on output markets; hence, they propagate only downstream. The propagation unfolds in consecutive periods. Input users of each industry i adjust their input demand by the new price of good i that was adjusted for the TFP shock in i. The production of these industries is thus affected by the shock of i. This effect appears in output markets in the next period, when these customer industries are ready to sell their output. This is the time when the second-order customers of i adjust their input demands to the shock of i from two periods before. And this step-by-step adjustment to shock i propagates through all indirect customers of i taking one step down the supply chain each period. This dynamic propagation process is formalized and analyzed in detail in the next section.

[6] This assumption could be modified to contemporaneous labor sourcing without doing any harm to the main predictions used in the empirical part of this paper. The reason is that the equilibrium wage rate w_{t-1} is set to be the numeraire in each period τ. Setting the numeraire to be w_τ would offset the effect of switching from lagged sourcing of labor to contemporaneous sourcing of labor.

[7] It is true even without the time-to-build friction, because of the constant returns to scale in the Cobb-Douglas production function. With such a production technology, the optimal output price and quantity are set such that the product, therefore the nominal revenues, stay constant. And because inputs are purchased as constant fractions of total revenues, their demand is unchanged by shocks.

The model provides a definition for GDP. $c_{i\tau}$ denotes consumption from good i in period τ. Because consumption is the only type of final use, aggregate consumption C_τ, defined by equation (3.1), is equivalent to GDP in this economy. Mostly, for this reason, $c_{i\tau}$ is proxied by sector value-added in the empirical part of this paper, even though value-added is not exactly equivalent to $c_{i\tau}$ in this economy. This choice, therefore, is a limitation of the current version of this paper.[8]

3.2.2 Dynamic Propagation

In general equilibrium, industry value-added is a linear combination of contemporaneous and lagged TFP shocks in logarithms combined by the production network A. This prediction is formalized by equation (3.4):

$$\mathbf{c}_\tau = \sum_{h=0}^{\infty} A^h \mathbf{z}_{\tau-h} + \kappa, \tag{3.4}$$

where $\mathbf{c}_\tau = [\ln c_{i\tau}]_i$ and $\mathbf{z}_\tau = [z_{i\tau}]_i$ are n-long vectors of real value-added and TFP shocks.[9] κ is a vector of constants consisting only of model parameters. This formula reveals that a TFP shock of any sector can affect any other sectors through the production network, represented by the matrix A. This happens because the vector of the shocks $\mathbf{z}_\tau$ is multiplied by matrix A. Moreover higher order lags of $\mathbf{z}_\tau$ (older vintages of TFP shocks) are multiplied by higher powers of A. Because A is the adjacency matrix of the production network, A^k is a matrix of walk counts of k-long walks from industry j to industry i. In this context, that means any a_{ij} element of matrix A represents the direct supplier relations ships from j to i, and any a^2_{ij} element of A^2 represents a two-step indirect supplier relationship from j to i. For example, A contains the link of coal mining $\rightarrow$ steel production, while A^2 contains coal mining $\rightarrow$ energy production $\rightarrow$ steel production. Similarly at higher powers, A^k traces kth step indirect supplier relationships, such as coal mining $\rightarrow$ ...($k-2$ intermediate sectors)...$\rightarrow$ housing for example.

[8]The difference between $c_{i\tau}$ and real value-added can be seen from the market clearing condition: $\mathbf{y}_\tau = \mathbf{c}_\tau + A\mathbf{x}_\tau$, where $\mathbf{y}_\tau$ is real gross output by industry, and $\mathbf{x}_\tau$ is real intermediate use by industry, and the definition of real value-added: $\text{RVA}_\tau = \mathbf{y}_\tau - A\mathbf{x}_{\tau-1}$, which is real gross output less real intermediate use that produced the given output.

[9]Derivation of this prediction is included in the Appendix.

The property that higher powers of matrix A are combining older shocks captures the idea of dynamic propagation. For example, a shock to coal mining in τ affects only the value-added of coal mining itself in τ, because it appears only in coal prices in τ. Although customers of coal, such as power plants have to accommodate this higher coal price into their coal demands, they do not pass this cost shock forward into their output prices and supplied quantities until the next time period $\tau + 1$, because of the time-to-build friction. This friction ties inputs purchased in τ to output sold in $\tau + 1$. Therefore the effect of this shock is transmitted to customers one step at a time period. This staggered propagation of TFP shocks is referred to as "dynamic propagation" in this paper.

The propagation described here is the same downstream propagation that happens in other production network models, such as in Acemoglu et al. (2012). The main difference is that in this economy the propagation is stretched out in time, taking only one step at a time period. This happens because of the time-to-build friction. Papers following up on Acemoglu et al. (2012), which are most papers in the production networks literature, are omitting the time-to-build friction, therefore predicting an instantaneous shock propagation.

This paper aims to measure the length of this time delay, for which I introduce the notation δ. δ represents "average propagation time." It is called propagation time because it measures the time delay with which an industry passes through the effects of a TFP shock to its direct customers. It is average because it is plausible to assume that there is some heterogeneity across sectors in propagation times. It is the limitation of the model, that it assumes a uniform propagation time for each sector.

δ is bounded to be non-negative. One extreme case is, when $\delta = 0$. In this case, propagation is instantaneous. Consequently, network propagation loses its time dimension, and the model becomes equivalent to the model of Acemoglu et al. (2012). In the other extreme case, if $1/\delta = 0$ there is no propagation at all, and equation (3.4) reduces to $\mathbf{c}_\tau = \mathbf{z}_\tau$. Thus, if δ is sufficiently large, that provides evidence for the dynamic propagation of TFP shocks in the production network on annual time frequencies. On the other hand, if δ is found to be a small number, that would mean most of the propagation happens within a year. Therefore, the network propagation of sector-level shocks has little relevance to annual frequencies. Therefore, if δ is found to be sufficiently large, it provides evidence for the dynamic propagation of TFP shocks in the production network on annual time frequencies. On the other hand, if δ is found to be a small number, that would mean most of the propagation happen within a year, therefore the network propagation of sector-level shocks has little relevance on annual frequencies.

3.2.3 Relevance of Production Network Dynamics

The secondary objective of this paper is to quantify the fraction of aggregate dynamics originating from network propagation of sector shocks. The macroeconomic aggregate this paper focuses on is GDP, which is equivalent to total consumption C_τ in the model. It is an aggregate of industry value-added series defined by equation (3.1): $\ln C_\tau = \gamma' \mathbf{c}_\tau$. Different industries respond to the same TFP shock with different time delays depending on how far they are situated in the supply chains of the shocked sector. In general, the more distant an industry is from the source of the shock, the later it responds.[10] These delayed responses make the aggregate effect of a sector-level shock fold out over time generating auto-correlation in GDP. These GDP dynamics can be captured by the substitution of equation (3.4) into (3.1) that results in the following equation:

$$\ln C_\tau = \gamma' \sum_{h=0}^{\infty} A^h \mathbf{z}_{\tau-h} + \gamma' \kappa. \tag{3.5}$$

This equation shows that GDP is also the function of TFP shocks from different time periods, which implies serial dependence of GDP observations from different time periods. If the propagation time δ was found to be comparable in magnitude to the length of a year that would predict that GDP is correlated with TFP shocks from multiple years. This prediction of GDP auto-correlations is the implication of the dynamic propagation of shocks in the production network. Calculating the fraction

[10]An industry can be a customer of another industry at multiple orders. For example, the manufacturing of chemical products might use oil as a direct input, and also as an indirect input through energy production, or through even longer chains, when it sources plastic materials. This multiplicity, however, does not harm conclusions, because all these possibly complex effects of TFP shocks arriving through different branches of the production networks are all summarized by input prices at each stage of propagation.

of observed GDP auto-correlation that can be explained by this network propagation effect determines the relevance of this propagation process.

3.3 Data

I use three data sets from the U.S. Bureau of Economic Analysis (BEA) in my empirical analysis. Industry value-added and industry GDP shares are obtained from the annual GDP by Industry dataset. Industry TFP data is obtained from the Integrated BEA GDP-BLS[11] Productivity Account. Production network coefficients are derived from Input-Output Accounts. The estimation sample is annual. It covers the 66 private industries out of the 71 sector breakdown of the US economy.[12] The sample includes years from 1997 to 2020. Sector TFPs are observable from 1987. These extra 10 years turn out to be helpful as identification leverages on lagged values of TFP shocks.[13]

TFP series are observed as unitless volume indexes. Therefore I transform both value-added and TFP series to growth rates. That transformation is consistent with the first differences of equations (3.4) and (3.5):

$$\Delta \mathbf{c}_\tau = \sum_{h=0}^{\infty} A^h \Delta \mathbf{z}_{\tau-h}, \tag{3.6}$$

$$\Delta \ln C_\tau = \gamma' \Delta \mathbf{c}_\tau \tag{3.7}$$

Input-output coefficients of matrix A are calculated from the industry-by-industry indirect requirements tables of the annual input-output tables of the BEA. The 66×66 matrix of indirect requirements correspond to the Leontief-inverse of an economy, which is defined as $\Gamma_t = (I - A_t)^{-1}$, such that I is a 66×66 identity matrix. A_t is referred to as the industry-by-industry direct requirements table, which can be obtained by expressing A_t as a function of the reported Γ_t matrix using this formula: $A_t = \Gamma_t^{-1} - I$. Finally, I define matrix A as the full sample average of the annual A_t matrices, such that: $A = 1/T \sum_{t=1}^{T} A_t$.

[11] U.S. Bureau of Labor Statistics.
[12] A list of the 66 sectors is included in Appendix C.1.
[13] See Appendix C.2 for details on TFP.

I need to calculate GDP shares (γ_i) to be able to empirically evaluate equation (3.1). I identify γ_i as the long-term average of the GDP shares of every industry i. I obtain nominal GDP shares from the GDP by Industry dataset of the BEA.

3.4 Empirical Strategy

The empirical strategy for determining average propagation time δ is presented in the first part of this section. The second part discusses the empirical strategy of the determination of what fraction of the annual auto-correlation of GDP can be attributed to the network propagation of sector shocks.

I define δ as the time unit of an unknown time-frequency that divides a year without a remainder.[14] I define a set of sub-annual time periods by splitting up a year into 1, 2, ... 365 equal parts, which would define annual, semi-annual, ... daily frequencies. These sub-annual frequencies define a set of possible values for δ as $12 \times \{1, 1/2, 1/3 \dots 1/365\}$. The multiplier 12 sets the unit of δ in months.

The main empirical challenge of the paper is the identification of a time interval that is shorter than the unit of the frequency of the observations. I tackle this challenge with the following strategy: I evaluate the model at each of these frequencies. I assume that one of these sub-annual time frequencies is the true frequency of the model, and the primary goal of the paper is to find that frequency. I predict the observed annual value-added growth rates using TFP shocks and input-output coefficients at each of these possible sub-annual frequencies. The sub-annual frequency that minimizes the sum of the squared errors between model-predicted and observed annual value-added growth rates is identified as the most likely time frequency of the model. The unit time length of this error minimizing time-frequency identifies δ, the average propagation time.

The key difference between model evaluations at different sub-annual frequencies is the number of steps the shock propagation process takes

[14]The reason for the strong assumption that I consider only whole divisors of a year is that the propagation process is discrete, and it is difficult to handle if a propagation step is split by the end of a year.

within a year. In the semi-annual setup, it takes two steps, while in the monthly setup, it makes twelve steps in a year. At each step, the effect of a shock propagates from suppliers to their direct customers. A two-step propagation, therefore, spreads the effects of a shock to first and second-order customers, while a twelve-step propagation reaches first, second, ... twelfth-order customers within a year. The variation caused by the different evaluations, therefore, appears in the predicted annual value-added growth rates. In a semi-annual setup, value-added growth rates of year t are affected by year t TFP shocks from first and second-order suppliers. In a monthly setup however value-added growth rates of year t are affected by year t TFP shocks from first, second ... twelfth-order suppliers. This difference is reflected in the formulas as different powers of matrix A, as the ijth element of A^k represents the link from supplier j to customer i as its kth order customer.

The following paragraphs formalize this idea.

3.4.1 Average Propagation Time

Although δ is the parameter in focus, and it has a meaningful interpretation as average propagation time, working with δ in the formulas of this section is inefficient and needlessly complicated at particular points. The formulas in this section are derived from model equation (3.6) and use the powers of the network matrix A to track the steps of the dynamic propagation process. It is more useful to switch from the average propagation time to the number of steps the propagation makes in a year because that number directly enters the power of A.

Let me define, therefore, *average propagation frequency* $\varphi = 12/\delta$, which is in a reciprocal relationship with δ. φ is counting the number of sub-annual periods (indexed by τ) within a year.[15] For example, if the average propagation time $\delta = 2$ months, the propagation process takes 4 steps within a year, therefore its average propagation frequency $\varphi = 4$.

The following aggregation rules are defined to connect theoretical quan-

[15] φ is not simply $1/\delta$, because δ is defined to be measured in months, while the unit of φ is 1/year.

tities labeled by asterisks to annual observations indexed by t:

$$\Delta \mathbf{c}_t = \sum_{s=0}^{\varphi-1} \Delta \mathbf{c}^*_{\varphi t+s} \tag{3.8}$$

$$\Delta \mathbf{z}_{t-\lceil h/\varphi \rceil} = \sum_{s=0}^{\varphi-1} \Delta \mathbf{z}^*_{\varphi t+s-h}, \tag{3.9}$$

where $\lceil \ldots \rceil$ is the ceiling function.[16]

The first rule connects observed value-added growth rates $\mathbf{c}_t$ to their unobserved theoretical counterparts $\mathbf{c}*_\tau$. $\mathbf{c}*_\tau$ has an unknown time frequency, which can be any whole divisor of a year, as $\varphi \in \{1, 2, \ldots 365\}$. By this formula, I assume that sub-annual growth rates sum up to the annually observed growth rates. This assumption is captured by the summation on the right-hand side. This summation has exactly φ terms, thus if the model time is quarterly, the summation index $s \in \{0, 1, 2, 3\}$. Each sub-annual growth rate has a time index $\varphi t + s$, because the sub-annual time indices have to cover φT numbers. For example, if the model time is quarterly, and the sample ends at year $T = 2$, the largest index for the quarterly series is 7, given the two-time scales meet at 0.

The second rule assumes that the annual TFP shocks are representative, therefore they are imputed for each sub-annual time period within a year. TFP shocks are defined as the growth rates of observed annual TFP series.[17] The left-hand side of equation (3.9) is a definition of the annual lags of the TFP shocks using model frequency lag noted by h. For example, if the model is quarterly, the unit of h is a quarter. This translation of sub-annual lags into annual series is formalized by the ceiling function. The left-hand side states that a TFP shock $\Delta \mathbf{z}_t$ with a sub-annual frequency time-lag h is equivalent to its annual time-lag from the year that would contain the sub-annual lag h. A time-lag $h = 3$, for example, points 3 quarters back, for which period the annual observation of $\Delta \mathbf{z}_t - 1$ is imputed. The right-hand side defines the annual TFP shocks as the sum of their unobserved sub-annual period components in the same way as equation (3.8) does for value-added growth rates.

The additivity of growth rates is an approximation that is more accu-

[16] Rounds a number to the smallest integer that is greater than the number. For example $\lceil 1.2 \rceil = 2$.

[17] The observed TFP process is assumed to be a random walk in logarithms, therefore its growth rates series is the shock process itself.

rate for single-digit growth rates. Figure 3.1 shows the value-added growth rates for the 66 private industries in the US between 1998 and 2020. Each line represents a single industry. Figure 3.2 is the same figure for TFP growth rates. These figures demonstrate that the majority of sector growth rates stayed within the single-digit boundaries (shaded area).

Figure 3.1: Annual Growth Rates of Real value-added by Private Industries.

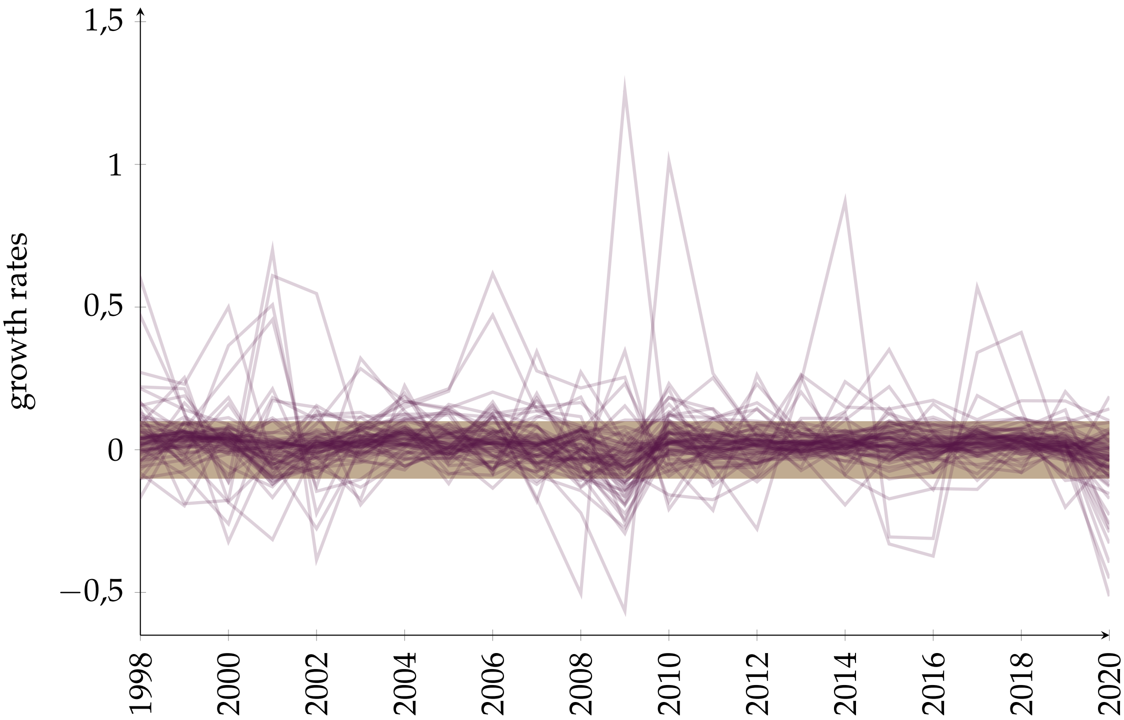

Notes: Each line represents one of the 66 industries. Single digit interval [-0.1,0.1] is highlighted by the shaded area.

Using aggregation rules (3.8), and (3.9) the following proposition converts equation 3.6 to an annual frequency empirical model:

Proposition 3.1. *empirical specification 1.*

$$\Delta \mathbf{c}_t = \sum_{k=0}^{\infty} B_k(A,\delta)\Delta \mathbf{z}_{t-k} + \mathbf{u}_t, \tag{3.10}$$

such that $B_k(A,\delta) = A^{k\varphi}(I - A^{\varphi})(I - A)^{-1}$, *and* $\delta = 12/\varphi$.

Proof. See Appendix C.4 ■

Here $B_k(A,\delta)$ is a series of $n \times n$ matrices by each annual time-lag k.

Figure 3.2: Annual Growth Rates of Total Factor Productivity by Private Industries

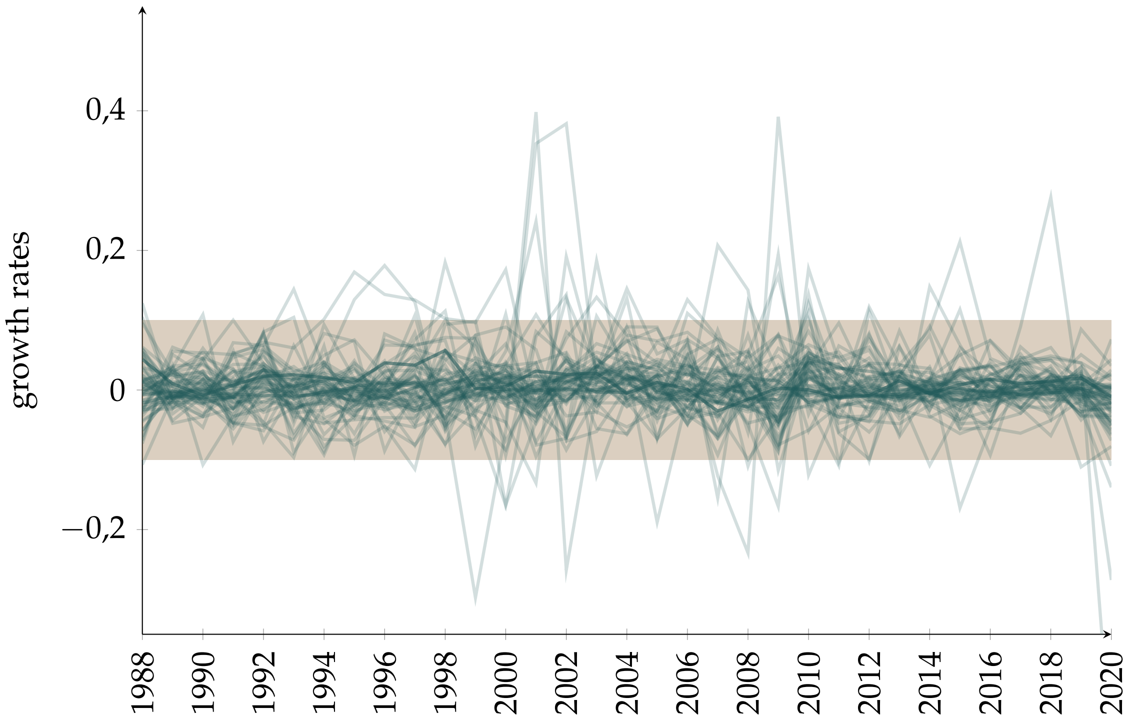

Notes: Each line represents one of the 66 industries. Single digit interval [-0.1,0.1] is highlighted by the shaded area.

Each of these matrices is a function of the production network A and the average propagation time δ. I is an $n \times n$ identity matrix.

$\mathbf{u}_t$ is the vector of the estimation errors in year t. It may consist of measurement errors in sector value-added $\mathbf{c}_t$, $\mathbf{z}_{t-k}$ TFP shocks, or the input-output coefficients contained in A. It may absorb the possible sector heterogeneity of propagation time. It can also reflect misspecifications of the theoretical model. Each of these might be a potential threat to the identification.

The most concerning is the possible correlation between contemporaneous TFP shocks and the error term as a result of measurement error in TFP. To address this threat, I propose a second empirical specification:

Proposition 3.2. *empirical specification 2.*

$$\Delta \mathbf{c}_t = \beta \mathbf{z}_t + \sum_{k=1}^{\infty} B_k(A, \delta) \Delta \mathbf{z}_{t-k} + \mathbf{u}_t. \tag{3.11}$$

This specification assumes that time lags of TFP shocks are uncorrelated with the error term. This identifying assumption is formalized by the following equation:

$$\text{Cov}\left[\Delta \mathbf{z}_{t-k}, \mathbf{u}_t | k > 0\right] = 0 \tag{3.12}$$

Identification

δ is identified by non-linear least squares (NLLS) in both of these specifications. It is carried out by pooling the whole sample. These empirical models are only defined for discrete values of δ putting a constraint on NLLS. The constrained NLLS estimator of δ is summarized by the following equation:

$$\widehat{\delta}_{NLLS} = \text{argmin} \frac{1}{NT} \sum_{i,t=1}^{N,T} u_{it}^2 \tag{3.13}$$
$$\text{such that } \delta \in [12/1, 12/2, 12/3, \ldots]$$

There are two sources of identifying variation in these empirical designs. One is the different powers of A at the different sub-annual frequencies. The other is that at lower frequencies, such as semi-annual, older vintages of annual TFP shocks have more influence on current value-added because they are multiplied by lower powers of A compared to a higher frequency setup. Lower powers of matrix A have larger coefficients because any element $a_i j$ of matrix A has a maximum value of 1 and is non-negative.

3.4.2 Predicted Auto-Correlation of GDP

To determine the aggregate relevance of dynamic network propagation I calculate predict GDP growth series by the following application of equation (3.7):

$$\Delta \ln \widehat{C}_t = \gamma' \Delta \widehat{\mathbf{c}}_t(\widehat{\delta}_{NLLS}), \tag{3.14}$$

where $\widehat{\mathbf{c}}_t(\widehat{\delta})$ are predictions of equation (3.11) evaluated at $\widehat{\delta}$. I then calculate the first three auto-correlation coefficients of $\Delta \ln \widehat{C}_t$ and compare them to the actually observed auto-correlations of annual GDP growth within the same time frame as the estimation sample.

As a benchmark, I also predict GDP growth $\Delta \ln \widehat{C}_t$ assuming instantaneous propagation, $\delta = 0$. Auto-correlation of GDP in this benchmark is not necessarily zero, because it is also based on actual annual TFP shocks. Thus, it retains any possible auto-correlation of aggregate TFP shocks.

3.5 Results

In this section, I present the results in three parts. The first part presents the results for average propagation time δ. The second part uses the results for δ to predict annual GDP growth from sector TFP shocks, input-output coefficients, and industry GDP shares. I compare the first three auto-correlation coefficients of this predicted GDP growth series to the same moments of observed GDP growth rates from the same time period. I do this to determine the relevance of the dynamic propagation of sector shocks that is predicted by the model of Long and Plosser (1983). In the third part, I split the sample into two sub-samples in 2008 and estimate the average propagation time for both sub-samples. I do that to investigate if the average propagation time has changed after the Great Recession of 2008.

3.5.1 Propagation Time

Table 3.1 shows the results for the propagation time parameter δ in two specifications. The first specification is based on equation (3.10). It identifies δ by utilizing both contemporaneous TFP shocks and its lags. This specification identifies the average propagation time to be 6 months. This value is the result of a constrained NLLS estimation that allows only integer numbers as the number of occurrences of propagation periods within a year. It is difficult to generate meaningful confidence intervals under this integer constraint. As a result, the unconstrained NLLS results for δ are in brackets, and the standard errors are in parenthesis. The unconstrained NLLS finds a δ of 6.1 months, which is almost equivalent to the constrained NLLS result. With a 2 standard error interval around the constrained result of 6 month, the average propagation time is between 4.6 and 7.4 months.

The second specification is based on equation (3.11). Only the lagged values of the TFP shocks are used to identify δ in this model. The contemporaneous TFP shock is included with a second parameter, β, which is identified by OLS. Similarly to the first specification, this specification identifies δ as 6. The unconstrained NLLS, however, finds it to be roughly one month longer: 7.6. This difference between the two specifications might suggest that the bias caused by the inclusion of contemporaneous TFP shocks in the first specification introduces a considerable bias. Considering the constrained result of 6 months, the two-standard deviation interval for the average propagation time is found to be between 4.6 and 7.4 months. This is a slight slacker interval, suggesting a trade-off between the bias of specification 1 and the efficiency cost of specification 2.

Table 3.1: Propagation Frequency

	(1)	(2)
δ	6***	6***
	[6.114]	[7.607]
	(0.702)	(0.964)
β		1.728***
		(0.109)
VA Observations	1518	1518
TFP Observations	2178	2178
R-squared	0.478	0.580

Notes: *** p<0.01, ** p<0.05, * p<0.1. NLLS point estimates for δ constrained to discrete values of $12 \times \{1, 1/2 \ldots 1/365\}$. Unconstrained NLLS estimates are shown in brackets, and standard errors estimated by 1000 bootstrap repetitions are shown in parenthesis. OLS estimates for β with robust standard errors in parenthesis.

These findings suggest that the average propagation time was roughly half a year in the US economy in the past two and a half decades. This propagation time is sufficiently slow to consider the network propagation of sector-level shocks a dynamic process even at annual frequencies. In other words, it is likely that industry-level shocks exert their impact on the economy much beyond the time horizon of a year. Consequently, the network propagation of these socks themselves might generate persistent macroeconomic aggregates. The next section quantifies the effect of this aggregate persistence resulting from network propagation.

3.5.2 Network Origins of Aggregate Dynamics

In this section, I take the result of $\widehat{\delta} = 6$ and evaluate equation (3.7) to predict GDP growth. I then calculate its first three auto-correlation coefficients to compare them to the same statistics of actually observed GDP growth rates for the same sample period.

Table 3.2 shows the results for these auto-correlations. The first two columns show coefficients for actual GDP growth data. The first column shows results for total GDP, which covers the entirety of the economy, including all private and public sectors. The first-order auto-correlation is found to be about 0.32. The second-order coefficient is about -0.03, which I consider being negligible. The third coefficient, that is the correlation of GDP growth in 3 years' distance, is -0.167. The small sample size of only 23 time-series observations between 1998 and 2020 does not allow for testing if these coefficients are significantly different from 0, as their standard errors are about 0.2. This is a clear limitation of the current version of this paper.[18]

Because the model considers only private sectors, I calculate a private sector GDP series as a more fitting benchmark. I define private GDP by subtracting the value-added by public sectors from total GDP. Column (2) shows auto-correlations of private GDP growth rates. Figures show that private and total GDP dynamics are very similar.

Columns (3) and (4) show model predictions for $\delta = 6$ and $\delta = 0$. The model evaluated at $\delta = 6$ replicates auto-correlation patterns of annual GDP growth rather well. I find that dynamic network propagation with an average propagation time of 6 months predicts a first-order auto-correlation of annual GDP growth rates of 0.27. This figure is roughly 80This result provides evidence for the network origins of aggregate dynamics.

This conclusion is strengthened by the numbers in column (4). The final column shows the auto-correlation coefficients for a model predic-

[18]For higher order auto-correlations the sample is even smaller because two components of the correlation coefficients have shrinking overlaps.

tion that assumes $\delta = 0$, which is the instantaneous propagation case. The only source of GDP dynamics in this model is auto-correlated TFP shocks. Results show that this model variation produces very few first and second-order GDP auto-correlations at annual frequencies. This result confirms that the strong first-order auto-correlation of column (3) is the result of dynamic network propagation.

The model has, on the other hand, a hard time matching higher-order auto-correlations. Although the second-order auto-correlations seem to be indistinguishable from zero, their signs are consistently negative for observations while positive for model predictions. The third-order auto-correlations are found to be negative for both of the models, consistent with observations. Their magnitude, however, is better matched by the instantaneous propagation model than the dynamics. These higher-order auto-correlations of annual time series connect two and three-year distant observations. These might be relevant for business cycle frequencies, as they are usually defined as 2–10 year cycles. It is difficult to find references to studies of such low-frequency auto-correlations.

Table 3.2: Actual and Model Predicted GDP Auto-Correlations

	Actual GDP		Model Prediction	
	(1)	(2)	(3)	(4)
Time-lag (Years)	Total	Private Sectors	Dynamic Propagation	Instantaneous Propagation
1	0.319	0.332	0.273	0.009
	(0.212)	(0.211)	(0.215)	(0.224)
2	-0.027	-0.008	0.050	0.014
	(0.224)	(0.224)	(0.224)	(0.224)
3	-0.167	-0.150	-0.052	-0.207
	(0.236)	(0.236)	(0.235)	(0.236)

Notes: Annual auto-correlation coefficients. *** p<0.01, ** p<0.05, * p<0.1, standard errors in parenthesis. Number of observations = 23.

3.5.3 The Great Recession

The Great Recession of 2008–2009 was the "most devastating global economic crisis since the Great Depression," according to Grusky et al. (2011). This brings up the question of whether the Great Recession altered the average propagation frequency of the US economy. 2008 sits almost exactly in the middle of my sample. This makes me able to split my sample into two sub-samples that are comparable in size and estimate propagation frequency separately on both samples.

Results for this comparative exercise are presented in Table 3.3. The first column shows results for the pre-crisis period of 1998–2007. The average propagation time is found to be 6 months. This result is equivalent to the full sample result. The unconstrained NLLS point estimate for δ is 8.4 with a standard error of 1.2. These figures suggest that the average propagation time before the Great Recession was between 3.6 and 8.4 months.

The second column shows results for the post-2008 period. In this subsample, I find the average propagation time to be 4 months. Column 3 shows the difference between the point estimates of the two sub-samples. This difference is found to be statistically significant based on the standard error of this difference. This result suggests that the network propagation of TFP shocks has accelerated after the Great Recession.

Table 3.3: Propagation Frequency Before and After 2008

	1998-2007	2008-2019	Difference
δ	6***	4***	-2***
	[8.513]	[5.931]	[-2.582]
	(1.354)	(1.055)	(0.249)
Observations	726	792	1518
R-squared	0.570	0.608	

Notes: *** p<0.01, ** p<0.05, * p<0.1. NLLS point estimates for δ constrained to discrete values of $12 \times \{1, 1/2 \ldots 1/365\}$. Unconstrained NLLS estimates are shown in brackets, standard errors estimated by 1000 bootstrap repetitions are shown in parenthesis.

3.6 Conclusion

In this paper, I estimate the average propagation time of the US economy. The average propagation time is the average time it takes for the effects of a producer-level shock to get transmitted from one industry to the other through input-output links. Using annual value-added and TFP data from a sample of 66 private sectors between 1998 and 2020, I find that the average propagation time was between 4 and 8 months. When I split the sample in 2008, the time of the Great Recession, I find that the average propagation time accelerated to between 2.4 and 5.6 months. These results suggest that the average propagation time was about a one-half year in the US economy in the past two and a half decades. This propagation time is sufficiently slow to consider the network propagation of sector-level shocks a dynamic process even at annual frequencies. In other words, TFP shocks exert their effect on the economy beyond the time horizon of one year. I find that network propagation produces persistent macroeconomic aggregates, accounting for 82% of the first-order auto-correlation of annual GDP growth.

These results provide evidence for the network origins of aggregate dynamics. That suggests there is more to be found in the direction of dynamic network propagation models within this literature. One possible extension could be to relax the assumption of a uniform propagation time. Another possible extension could be the investigation of continuous time models or other types of supplier link level frictions that might predict similar dynamic patterns as the theory that provides the basis for this paper.

This finding also suggests that there is more to be explored empirically about dynamic network propagation. One might be interested in estimating the average propagation time in other economies as well. Or to extend the exploration to other aggregates or dynamic moments as well.

My findings provide the basis for new opportunities in economic policy. Grassi and Sauvagnat (2019) shows how to use input-output data and production network models to draw predictions for different domains of policy including industrial or fiscal policies. The primary policy gain from these models is the opportunity for better-targeted interventions. This paper provides evidence for the dynamic propagation of sector-level shocks. This finding extends the targeting opportunities of policy interventions in the time dimension. This extension is prominently important in the prevention of welfare-diminishing consequences of economic shocks.

Bibliography

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Cascades in networks and aggregate volatility," Technical Report, National Bureau of Economic Research 2010.

_ , Ufuk Akcigit, and William Kerr, "Networks and the macroeconomy: An empirical exploration," *NBER Macroeconomics Annual*, 2016, *30* (1), 273–335.

_ , Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The network origins of aggregate fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Arnon, Alexander, John Ricco, and Kent Smetters, "Epidemiological and economic effects of lockdown," *Brookings Papers on Economic Activity*, 2020, *2020* (3), 61–108.

Askitas, Nikolaos, Konstantinos Tatsiramos, and Bertrand Verheyden, "Estimating worldwide effects of non-pharmaceutical interventions on COVID-19 incidence and population mobility patterns using a multiple-event study," *Scientific reports*, 2021, *11* (1), 1–13.

Baqaee, David Rezza and Emmanuel Farhi, "The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Barrot, Jean-Noël and Julien Sauvagnat, "Input specificity and the propagation of idiosyncratic shocks in production networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Barrot, Jean-Noel, Basile Grassi, and Julien Sauvagnat, "Sectoral effects of social distancing," in "AEA Papers and Proceedings," Vol. 111 2021, pp. 277–81.

Berry, Christopher R, Anthony Fowler, Tamara Glazer, Samantha Handel-Meyer, and Alec MacMillen, "Evaluating the effects of shelter-in-place policies during the COVID-19 pandemic," *Proceedings of the National Academy of Sciences*, 2021, *118* (15).

Bodenstein, Martin, Giancarlo Corsetti, and Luca Guerrieri, "Economic and epidemiological effects of mandated and spontaneous social distancing," *Working Paper*, 2021.

Boranova, Vizhdan, Raju Huidrom, Ezgi Ozturk, Ara Stepanyan, Petia Topalova, and Shihangyin Frank Zhang, "Cars in Europe: Supply Chains and Spillovers during COVID-19 Times," Technical Report, International Monetary Fund 2022.

Carvalho, Vasco and Xavier Gabaix, "The Great Diversification and its Undoing," *THE AMERICAN ECONOMIC REVIEW*, 2013, *103* (5), 1697–1727.

Carvalho, Vasco M, "From micro to macro via production networks," *The Journal of Economic Perspectives*, 2014, pp. 23–47.

_ **and Basile Grassi**, "Large firm dynamics and the business cycle," *American Economic Review*, 2019, *109* (4), 1375–1425.

_ **, Juan R Garcia, Stephen Hansen, Álvaro Ortiz, Tomasa Rodrigo, José V Rodríguez Mora, and Pep Ruiz**, "Tracking the COVID-19 crisis with high-resolution transaction data," *Royal Society Open Science*, 2020, *8* (8), 210218.

_ **, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi**, "Supply chain disruptions: Evidence from the great east japan earthquake," *The Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.

Castex, Gonzalo, Evgenia Dechter, and Miguel Lorca, "COVID-19: The impact of social distancing policies, cross-country analysis," *Economics of disasters and climate change*, 2021, *5* (1), 135–159.

Chari, Varadarajan V, Patrick J Kehoe, and Ellen R McGrattan, "Sticky price models of the business cycle: can the contract multiplier solve the persistence problem?," *Econometrica*, 2000, *68* (5), 1151–1179.

Chen, Sophia, Deniz O Igan, Nicola Pierri, Andrea F Presbitero, Maria Soledad, and Martinez Peria, "Tracking the economic impact of COVID-19 and mitigation policies in Europe and the United States," *IMF Working Papers*, 2020, *2020* (125).

Chernozhukov, Victor, Hiroyuki Kasahara, and Paul Schrimpf, "Causal impact of masks, policies, behavior on early Covid-19 pandemic in the US," *Journal of Econometrics*, 2021, *220* (1), 23–62.

Cogley, Timothy and James M Nason, "Output dynamics in real-business-cycle models," *The American Economic Review*, 1995, pp. 492–511.

Cori, Anne, Neil M Ferguson, Christophe Fraser, and Simon Cauchemez, "A new framework and software to estimate time-varying reproduction numbers during epidemics," *American journal of epidemiology*, 2013, *178* (9), 1505–1512.

Cox, Nicholas J., "CSIPOLATE: Stata module to perform cubic spline interpolation," Statistical Software Components, Boston College Department of Economics February 2009.

Cunningham, Scott, *Causal Inference*, Yale University Press, 2021.

Deb, Pragyan, Davide Furceri, Jonathan D Ostry, and Nour Tawk, "The economic effects of Covid-19 containment measures," *Open Economies Review*, 2021, pp. 1–32.

Druant, Martine, Silvia Fabiani, Gabor Kezdi, Ana Lamo, Fernando Martins, and Roberto Sabbatini, "Firms' price and wage adjustment in Europe: Survey evidence on nominal stickiness," *Labour Economics*, 2012, *19* (5), 772–782. Special Section on: Price, Wage and Employment Adjustments in 2007-2008 and Some Inferences for the Current European Crisis.

Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Bettina Landau, Claire Loupias, Fernando Martins, Thomas Mathä, Roberto Sabbatini, Harald Stahl et al., "THE PRICING BEHAVIOUR OF FIRMS IN THE EURO AREA: NEW SURVEY EVIDENCE," *Working Paper*, 2005.

Goolsbee, Austan and Chad Syverson, "Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020," *Journal of public economics*, 2021, *193*, 104311.

Grassi, Basile and Julien Sauvagnat, "Production networks and economic policy," *Oxford Review of Economic Policy*, 12 2019, *35* (4), 638–677.

Grusky, David B, Bruce Western, and Christopher Wimer, *The great recession*, Russell Sage Foundation, 2011.

Gupta, Sumedha, Kosali I Simon, and Coady Wing, "Mandated and voluntary social distancing during the covid-19 epidemic: A review," 2020.

_, Thuy D Nguyen, Felipe Lozano Rojas, Shyam Raman, Byungkyu Lee, Ana Bento, Kosali I Simon, and Coady Wing, "Tracking Public and Private Response to the COVID-19 Epidemic: Evidence from State and Local Government Actions," Technical Report, National Bureau of Economic Research 2020.

Hale, Thomas, Sam Webster, Anna Petherick, Toby Phillips, and Beatriz Kira, "Oxford COVID-19 Government Response Tracker," 2020.

Haug, Nils, Lukas Geyrhofer, Alessandro Londei, Elma Dervic, Amélie Desvars-Larrive, Vittorio Loreto, Beate Pinior, Stefan Thurner, and Peter Klimek, "Ranking the effectiveness of worldwide COVID-19 government interventions," *Nature human behaviour*, 2020, *4* (12), 1303–1312.

Hausman, Catherine and David S Rapson, "Regression discontinuity in time: Considerations for empirical applications," *Annual Review of Resource Economics*, 2018, *10*, 533–552.

Heckman, James J and Guilherme Sedlacek, "Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market," *Journal of political Economy*, 1985, *93* (6), 1077–1125.

Hoeberichts, Marco and Ad Stokman, "Price setting behaviour in the Netherlands: Results of a survey," *Managerial and Decision Economics*, 2010, *31* (2-3), 135–149.

Islam, Nazrul, Stephen J Sharp, Gerardo Chowell, Sharmin Shabnam, Ichiro Kawachi, Ben Lacey, Joseph M Massaro, Ralph B D'Agostino, and Martin White, "Physical distancing interventions and incidence of coronavirus disease 2019: natural experiment in 149 countries," *British Medical Journal*, 2020, *370.*

Koh, Wee Chian, Lin Naing, and Justin Wong, "Estimating the impact of physical distancing measures in containing COVID-19: an empirical analysis," *International Journal of Infectious Diseases*, 2020, *100*, 42–49.

Kong, Edward and Daniel Prinz, "Disentangling policy effects using proxy data: Which shutdown policies affected unemployment during the COVID-19 pandemic?," *Journal of Public Economics*, 2020, *189*, 104257.

Kwapil, Claudia, Johann Scharler, and Josef Baumgartner, "How are prices adjusted in response to shocks? Survey evidence from Austrian firms," *Managerial and Decision Economics*, 2010, *31* (2-3), 151–160.

Kydland, Finn E and Edward C Prescott, "Time to build and aggregate fluctuations," *Econometrica: Journal of the Econometric Society*, 1982, pp. 1345–1370.

Liu, Ernest and Aleh Tsyvinski, "Dynamical structure and spectral properties of input-output networks," Technical Report, National Bureau of Economic Research 2020.

Liu, Ying, Albert A Gayle, Annelies Wilder-Smith, and Joacim Rocklöv, "The reproductive number of COVID-19 is higher compared to SARS coronavirus," *Journal of travel medicine*, 2020.

Long, John B and Charles I Plosser, "Real business cycles," *The Journal of Political Economy*, 1983, pp. 39–69.

Maloney, William F and Temel Taskin, "Determinants of social distancing and economic activity during COVID-19: A global view," *World Bank Policy Research Working Paper*, 2020, (9242).

Martins, Fernando, "Price stickiness in Portugal evidence from survey data," *Managerial and Decision Economics*, 2010, *31* (2-3), 123–134.

Perra, Nicola, "Non-pharmaceutical interventions during the COVID-19 pandemic: A review," *Physics Reports*, 2021.

Petersen, Eskild, Marion Koopmans, Unyeong Go, Davidson H Hamer, Nicola Petrosillo, Francesco Castelli, Merete Storgaard, Sulien Al Khalili, and Lone Simonsen, "Comparing SARS-CoV-2 with SARS-CoV and influenza pandemics," *The Lancet infectious diseases*, 2020.

Picard, Robert, "GEOINPOLY: Stata module to match geographic locations to shapefile polygons," Statistical Software Components, Boston College Department of Economics May 2015.

Rácz, Olivér M., "The Effect of Distancing Policies on the Reproduction Number of COVID-19," *Workingpaper*, 2022.

Seres, Gyula, Anna Helen Balleyer, Nicola Cerutti, Anastasia Danilov, Jana Friedrichsen, Yiming Liu, and Müge Süer, "Face masks increase compliance with physical distancing recommendations during the COVID-19 pandemic," *Journal of the Economic Science Association*, 2021, pp. 1–20.

Silva, JMC Santos and Silvana Tenreyro, "The log of gravity," *The Review of Economics and statistics*, 2006, *88* (4), 641–658.

Ullah, Akbar and Olubunmi A Ajala, "Do Lockdown and Testing Help in Curbing Covid-19 Transmission?," *Covid Economics*, 2020, *13*, 138–185.

Wahltinez, O. et al., "COVID-19 Open-Data: curating a fine-grained, global-scale data repository for SARS-CoV-2," 2020. Work in progress.

Appendix A

Appendix for Chapter 1

A.1 Policy Pairs

Table A.1: Percent of Countries Implementing a Policy Pair within 3 or Less Days.

	Place Restriction				Mobility Restriction		
	School	Event	Gather	Work	Stay H	Move	Transp't
Schools		43.12	44.44	38.53	23.15	24.07	10.00
Events	43.12		53.70	26.61	15.74	21.30	10.00
Gatherings	44.44	53.70		38.89	29.91	27.10	24.00
Workplaces	38.53	26.61	38.89		40.74	37.96	40.00
Stay Home	23.15	15.74	29.91	40.74		48.60	42.42
Movement	24.07	21.30	27.10	37.96	48.60		43.43
Transport	10.00	10.00	24.00	40.00	42.42	43.43	

Coloring: Dark yellow $\geq$ 66.7%, *light yellow* $\geq$ 50%.

Table A.2: Percent of Countries Implementing a Policy Pair within 5 or Less Days.

	Place Restriction				Mobility Restriction		
	School	Event	Gather	Work	Stay H	Move	Transp't
Schools		66.97	58.33	57.80	39.81	41.67	26.00
Events	66.97		64.81	44.95	34.26	34.26	23.00
Gatherings	58.33	64.81		52.78	44.86	44.86	35.00
Workplaces	57.80	44.95	52.78		52.78	50.00	48.00
Stay Home	39.81	34.26	44.86	52.78		60.75	51.52
Movement	41.67	34.26	44.86	50.00	60.75		56.57
Transport	26.00	23.00	35.00	48.00	51.52	56.57	

Coloring: Dark yellow $\geq$ 66.7%, *light yellow* $\geq$ 50%.

Table A.3: Percent of Countries Implementing a Policy Pair within 7 or Less Days.

	Place Restriction				Mobility Restriction		
	School	Event	Gather	Work	Stay H	Move	Transp't
School Closure		76.15	67.59	71.56	50.00	53.70	39.00
Events Cancelled	76.15		70.37	55.96	44.44	49.07	29.00
Gathering Limit	67.59	70.37		65.74	57.94	58.88	44.00
Workplace Closure	71.56	55.96	65.74		62.96	60.19	54.00
Stay Home Order	50.00	44.44	57.94	62.96		67.29	58.59
Movement Restricted	53.70	49.07	58.88	60.19	67.29		64.65
Public Transport Closed	39.00	29.00	44.00	54.00	58.59	64.65	

Coloring: Dark yellow $\geq$ 66.7%, *light yellow* $\geq$ 50%.

Table A.4: Percent of Countries Implementing a Policy Pair within 9 or Less Days.

	Place Restriction				Mobility Restriction		
	School	Event	Gather	Work	Stay H	Move	Transp't
Schools		85.32	75.93	77.98	55.56	60.19	49.00
Events	85.32		74.07	66.06	52.78	56.48	41.00
Gatherings	75.93	74.07		72.22	65.42	63.55	53.00
Workplaces	77.98	66.06	72.22		70.37	68.52	65.00
Stay Home	55.56	52.78	65.42	70.37		70.09	64.65
Movement	60.19	56.48	63.55	68.52	70.09		70.71
Transport	49.00	41.00	53.00	65.00	64.65	70.71	

Coloring: Dark yellow $\geq$ 66.7%, *light yellow* $\geq$ 50%.

Table A.5: Number of Countries Implementing both Policies of a Policy Pair.

	School	Event	Gather	Work	Transp't	Stay H	Move
Schools		109	108	109	100	108	108
Events	109		108	109	100	108	108
Gatherings	108	108		108	100	107	107
Workplaces	109	109	108		100	108	108
Transport	100	100	100	100		99	99
Stay Home	108	108	107	108	99		107
Movement	108	108	107	108	99	107	

Notes

A.2 Timing of Distancing Policy Interventions by Country

Figures A.1 and A.2 show the time of the first distancing interventions relative to the day of the first reported COVID-19 case. These figures demonstrate that there is a sufficiently large variation in the adoption times of distancing interventions to make their effects feasible to identify with panel econometric methods. It is also apparent from the figures that many countries implemented their first distancing interventions before they even had a confirmed COVID case within their borders.

Figure A.1: First Place Restrictions by Countries

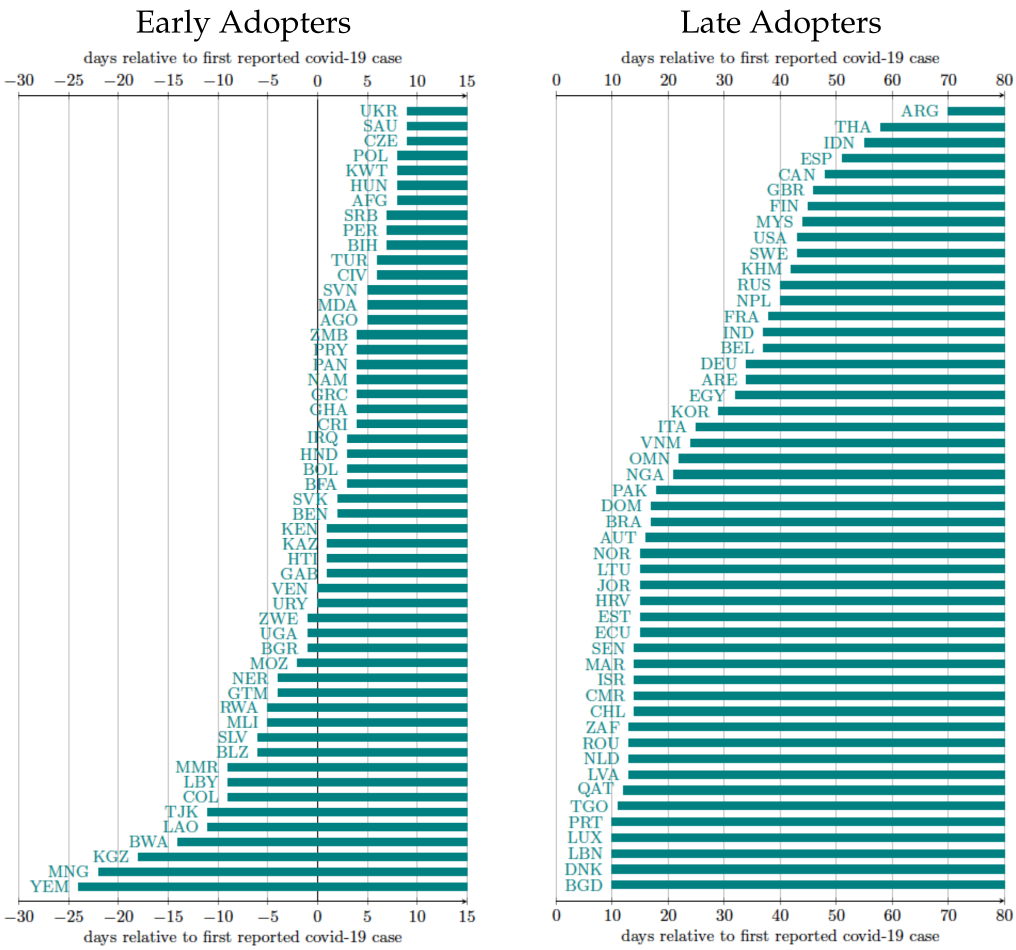

10.14754/CEU.2022.06

Figure A.2: First Mobility Restrictions by Countries

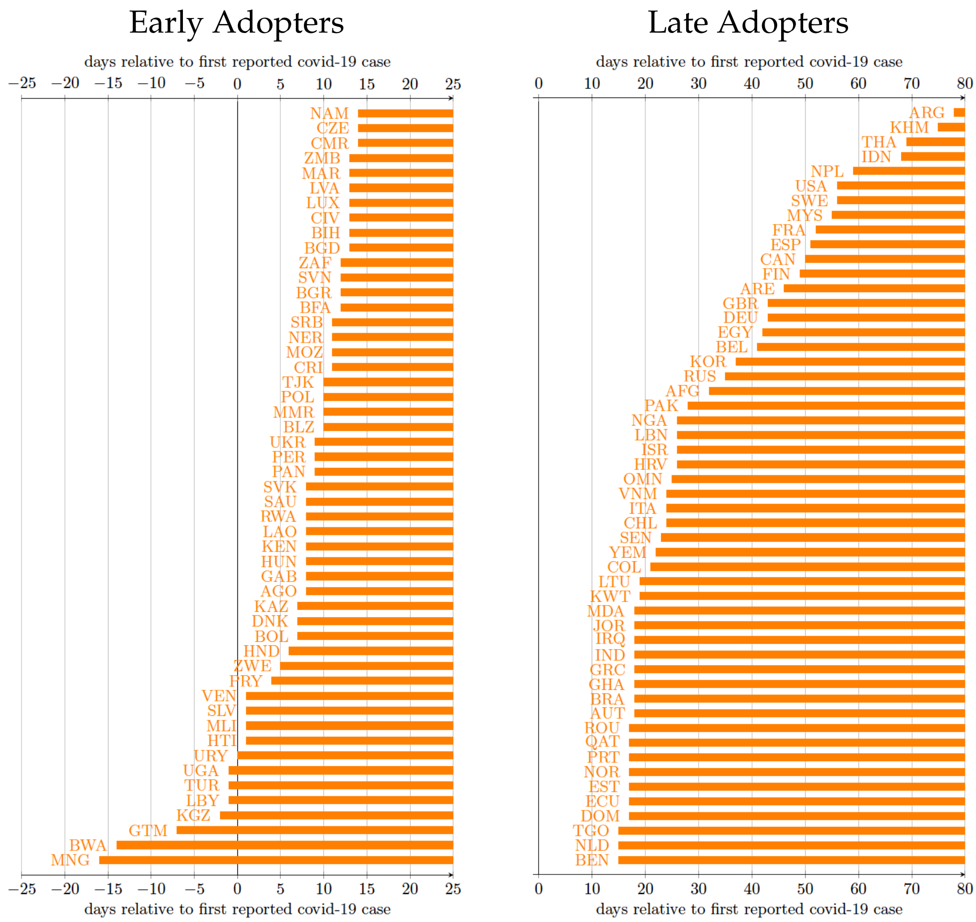

A.3 Validation of the Calculation of Reproduction Numbers

Here I compare my definition for R_t to an estimation of R_t using the methodology of Cori et al. (2013). It is a parametric calculation for which I use the following parameters: mean SI = 6, standard deviation of SI = 3, aimed posterior CV = .3, length of time-steps = 7, number of steps estimated = 1, posterior mean=5, posterior st.d. = 5. I input new case incidence data for each country from Wahltinez et al. (2020). For my definition of R_t I calculate R_t^I first, than normalize it by its within country mean and multiply it by the within country mean of Cori et al. (2013).[1]

[1]This renormalization does not harm my conclusions as it is based on $R_t \propto R_t^I$.

Figure A.3: Validation of R_t by Cori et al. (2013) – Germany

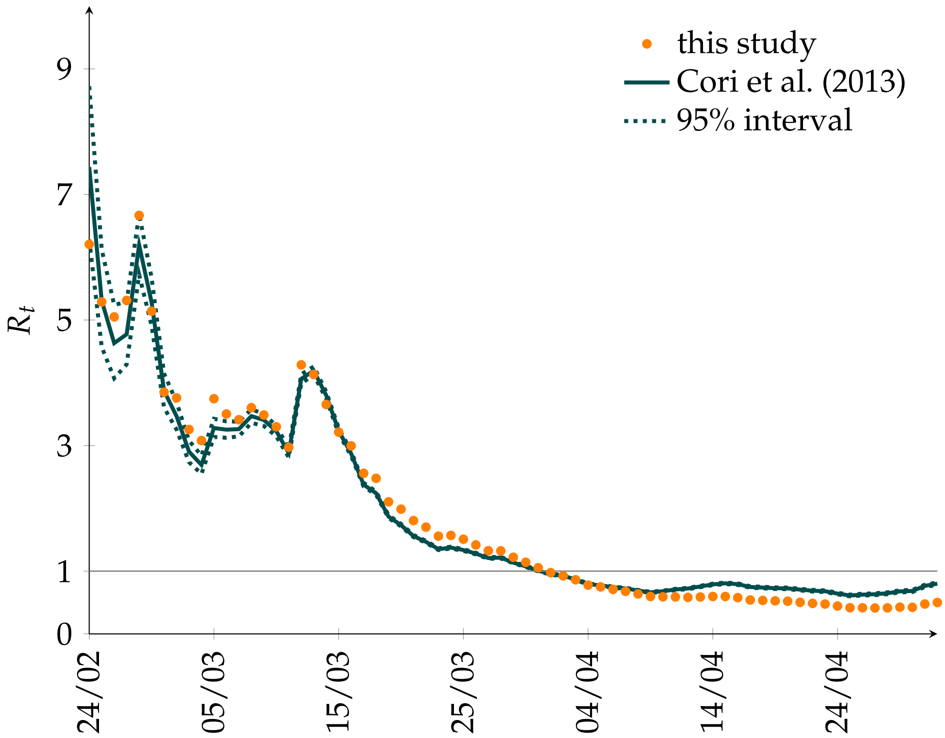

Figure A.4: Validation of R_t by Cori et al. (2013) – Italy

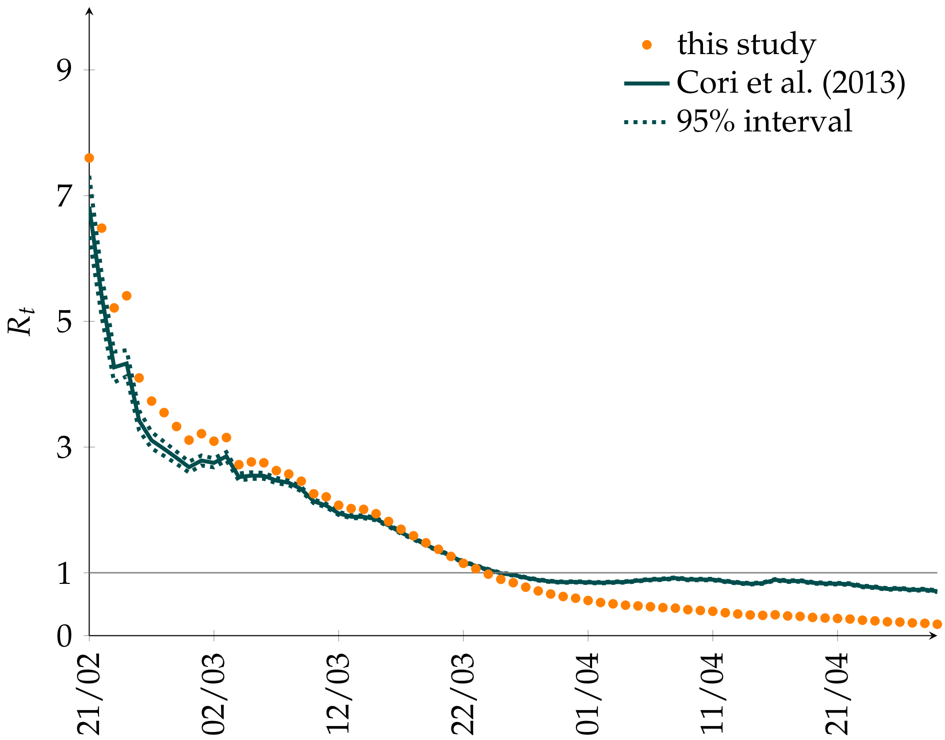

Figure A.5: Validation of R_t by Cori et al. (2013) – France

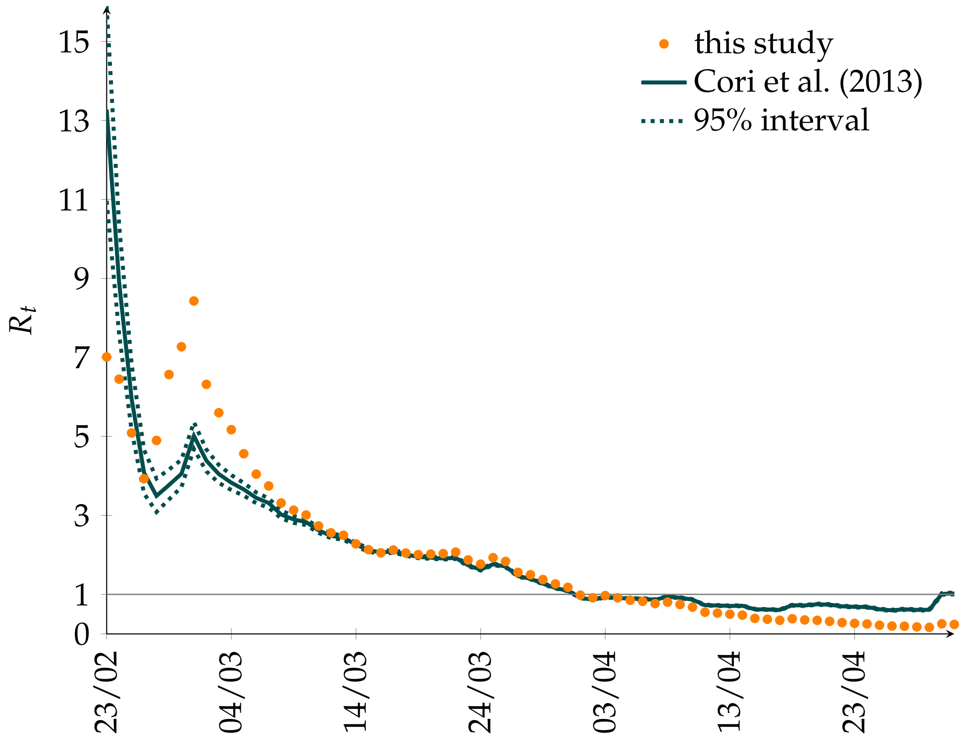

Figure A.6: Validation of R_t by Cori et al. (2013) – Spain

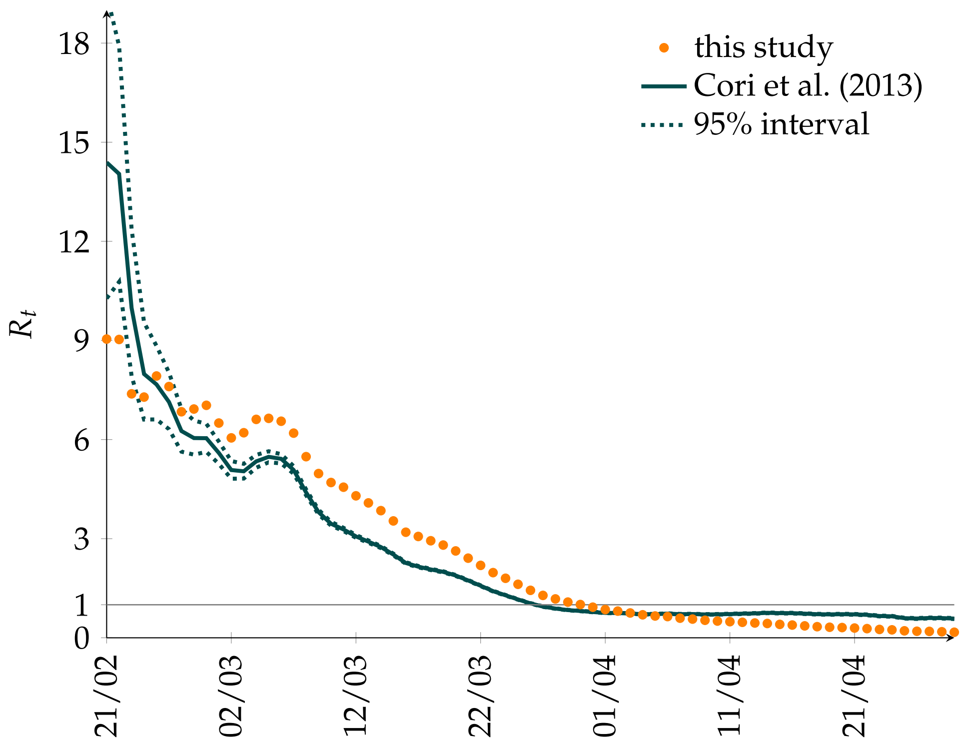

Figure A.7: Validation of R_t by Cori et al. (2013) – UK

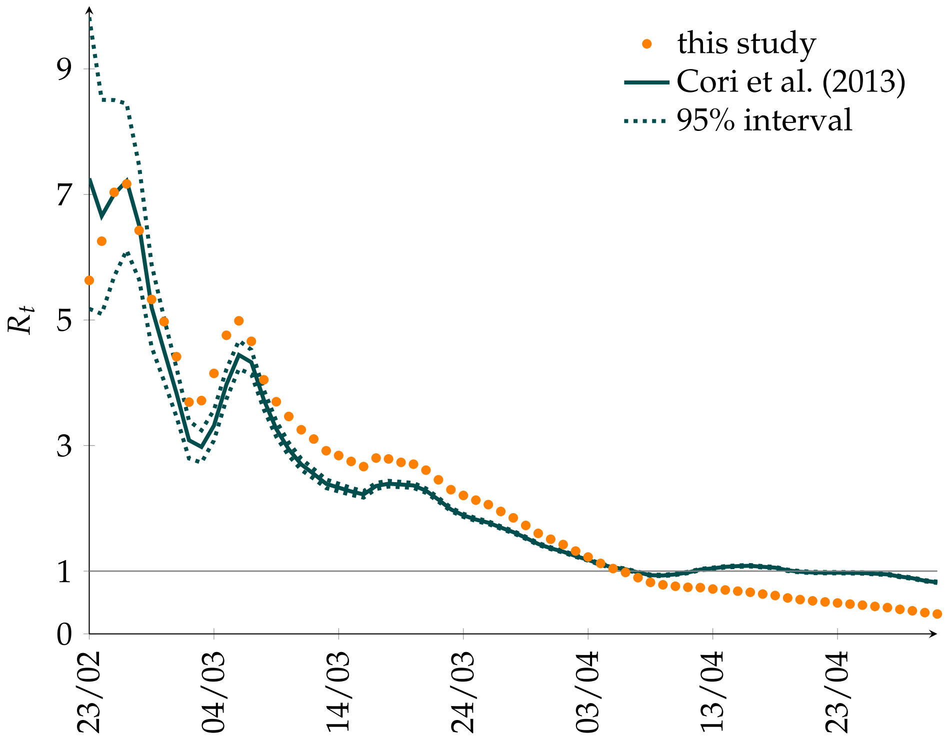

A.4 COVID-19 Aggregated Mobility Research Dataset

Description The Google COVID-19 Aggregated Mobility Research Dataset contains anonymized mobility flows aggregated over users who have turned on the Location History setting, which is off by default. This is similar to the data used to show how busy certain types of places are in Google Maps — helping identify when a local business tends to be the most crowded. The dataset aggregates flows of people from region to region, which is here further aggregated at the level of NUTS3 areas, weekly.

To produce this dataset, machine learning is applied to logs data to automatically segment it into semantic trips https://www.nature.com/articles/s41467-019-12809-y. To provide strong privacy guarantees, all trips were anonymized and aggregated using a differentially private mechanism https://research.google/p to aggregate flows over time (see https://policies.google.com/technologies/anony This research is done on the resulting heavily aggregated and differentially private data. No individual user data was ever manually inspected, only heavily aggregated flows of large populations were handled.

All anonymized trips are processed in aggregate to extract their origin and destination location and time. For example, if users traveled from location a to location b within time interval t, the corresponding cell (a, b, t) in the tensor would be $n \pm err$, where err is Laplacian noise. The automated Laplace mechanism adds random noise drawn from a zero mean Laplace distribution and yields (ϵ, δ)-differential privacy guarantee of $\epsilon = 0.66$ and $\delta = 2.1 \times 10 - 29$ per metric. Specifically, for each week W and each location pair (A, B), we compute the number of unique users who took a trip from location A to location B during week W. To each of these metrics, we add Laplace noise from a zero-mean distribution of scale 1/0.66. We then remove all metrics for which the noisy number of users is lower than 100, following the process described in https://research.google/pubs/pub48778/, and publish the rest. This yields that each metric we publish satisfies (ϵ, δ)-differential privacy with values defined above. The parameter ϵ controls the noise intensity in terms of its variance, while δ represents the deviation from pure ϵ-privacy. The closer they are to zero, the stronger the privacy guarantees.

Limitations These results should be interpreted in light of several important limitations. First, the Google mobility data is limited to smartphone users who have opted in to Google's Location History feature, which is off by default. These data may not be representative of the population as whole, and furthermore their representativeness may vary by location. Importantly, these limited data are only viewed through the lens of differential privacy algorithms, specifically designed to protect user anonymity and obscure fine detail. Moreover, comparisons across rather than within locations are only descriptive since these regions can differ in substantial ways.

Data Availability The Google COVID-19 Aggregated Mobility Research Dataset used for this study is available with permission from Google LLC.

A.5 Covariates of the First Stage

Table A.6: Effect of News Components in the First Stage

Cases $_{t-1}$	-0.005	Neighbors' Deaths $_{t-1}$	-0.895*
	(0.003)		(0.457)
$\sum_{s=2}^{7}$Cases $_{t-s}$	-0.070***	$\sum_{s=2}^{7}$N's' Deaths $_{t-s}$	-3.400**
	(0.016)		(1.441)
$\sum_{s=8}^{14}$Cases $_{t-s}$	-0.067***	$\sum_{s=8}^{14}$N's' Deaths $_{t-s}$	-3.117*
	(0.023)		(1.844)
Deaths $_{t-1}$	-0.266	Neighbors' Place R's $_{t-1}$	0.343***
	(0.183)		(0.102)
$\sum_{s=2}^{7}$Deaths $_{t-s}$	-2.000**	$\sum_{s=2}^{7}$N's' Place R's $_{t-s}$	-0.278
	(0.811)		(0.199)
$\sum_{s=8}^{14}$Deaths $_{t-s}$	-3.029***	$\sum_{s=8}^{14}$N's' Place R's $_{t-s}$	-0.046
	(0.885)		(0.241)
Neighbors' Cases $_{t-1}$	0.015	Neighbors' Mobility R's $_{t-1}$	-0.130
	(0.016)		(0.083)
$\sum_{s=2}^{7}$N's' Cases $_{t-s}$	-0.055*	$\sum_{s=2}^{7}$N's' Mobility R's $_{t-s}$	-0.598***
	(0.032)		(0.108)
$\sum_{s=8}^{14}$N's' Cases $_{t-s}$	0.180***	$\sum_{s=8}^{14}$N's' Mobility R's $_{t-s}$	0.435**
	(0.054)		(0.167)
Observations			50,036
R-squared			0.685

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors clustered at country level in parenthesis. Dependent variable is changes in activity relative to a five week benchmark period from before the epidemic. Homeland Cases and Deaths are reports from day $t-1$ published on day t in country i, Neighbors' Cases and Deaths are sum of reports from countries sharing a land border with i. All reports measured in case per 10 000 citizens.

10.14754/CEU.2022.06

A.6 Second Stage Daily Event Study

Figure A.8: Effects of Distancing Policies on the Reproduction Number on the Same Day

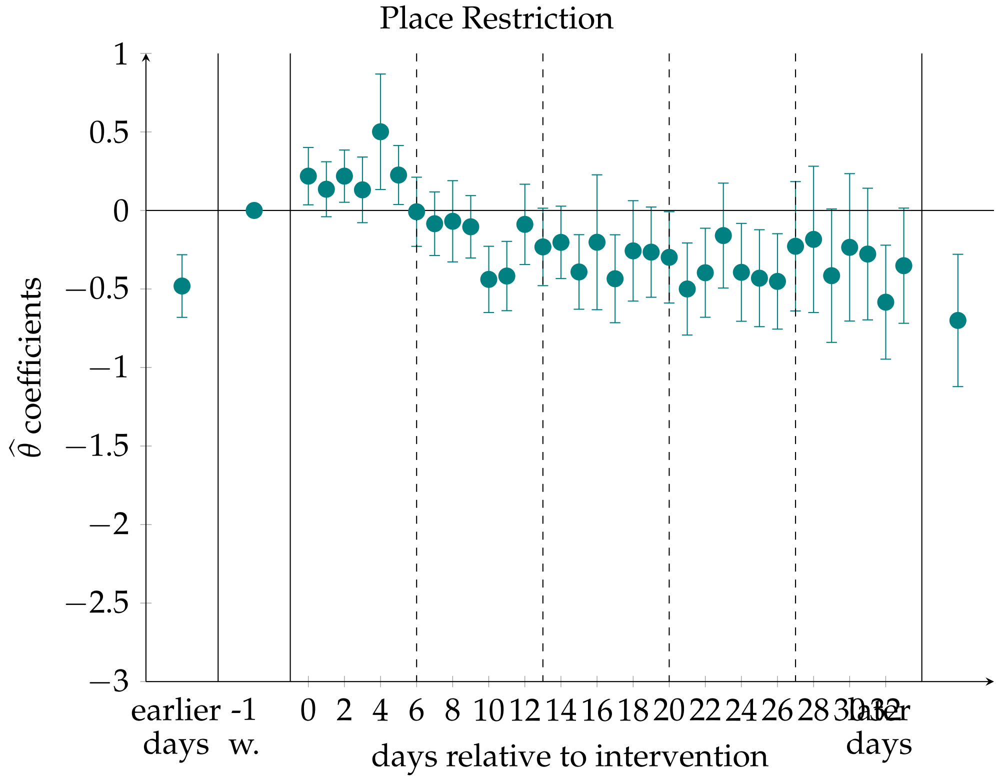

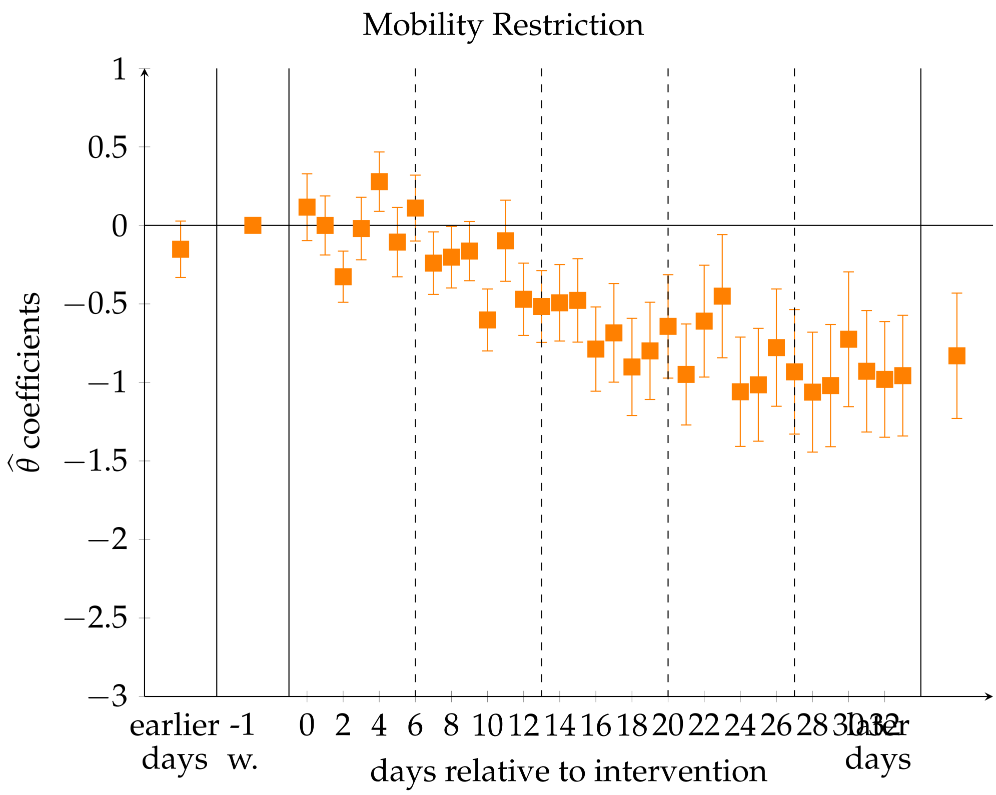

10.14754/CEU.2022.06

A.7 Sensitivity to latency parameter

Table A.7: Effect of Distancing Policies on Reproduction Number l days later.

	(1)	(2)	(3)	(4)
latency parameter l	7	9	11	13
Place Restrictions $_{t-4}$	-0.274*	-0.315**	-0.309**	-0.271**
	(0.149)	(0.144)	(0.149)	(0.135)
Mobility Restriction $_t$	-0.568***	-0.556***	-0.552***	-0.580***
	(0.128)	(0.121)	(0.127)	(0.133)
Voluntary Activity $_t$	0.138***	0.143***	0.142***	0.144***
	(0.042)	(0.043)	(0.047)	(0.048)
Observations	26,151	26,151	26,151	26,151
Countries	109	109	109	109
Preventive Policies	•	•	•	•
Country and Day FE's	•	•	•	•

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parenthesis allowing for country level clustering. Dependent variable is instantaneous reproduction number $R^I_{i,t+l}$. • = included ○ = excluded. Controlled for daily weather conditions and weekly seasonality.

Appendix B

Appendix for Chapter 2

B.1 COVID-19 Aggregated Mobility Research Dataset

Description The Google COVID-19 Aggregated Mobility Research Dataset contains anonymized mobility flows aggregated over users who have turned on the Location History setting, which is off by default. This is similar to the data used to show how busy certain types of places are in Google Maps — helping identify when a local business tends to be the most crowded. The dataset aggregates flows of people from region to region, which is here further aggregated at the level of NUTS3 areas, weekly.

To produce this dataset, machine learning is applied to logs data to automatically segment it into semantic trips

https://www.nature.com/articles/s41467-019-12809-y. To provide strong privacy guarantees, all trips were anonymized and aggregated using a differentially private mechanism https://research.google/pubs/pub48778/ to aggregate flows over time (see

https://policies.google.com/technologies/anonymization). This research is done on the resulting heavily aggregated and differentially private data. No individual user data was ever manually inspected, only heavily aggregated flows of large populations were handled.

All anonymized trips are processed in aggregate to extract their origin and destination location and time. For example, if users traveled from location a to location b within time interval t, the corresponding cell

(a, b, t) in the tensor would be $n \pm err$, where err is Laplacian noise. The automated Laplace mechanism adds random noise drawn from a zero mean Laplace distribution and yields (ϵ, δ)-differential privacy guarantee of $\epsilon = 0.66$ and $\delta = 2.1 \times 10 - 29$ per metric. Specifically, for each week W and each location pair (A, B), we compute the number of unique users who took a trip from location A to location B during week W. To each of these metrics, we add Laplace noise from a zero-mean distribution of scale 1/0.66. We then remove all metrics for which the noisy number of users is lower than 100, following the process described in https://research.google/pubs/pub48778/, and publish the rest. This yields that each metric we publish satisfies (ϵ, δ)-differential privacy with values defined above. The parameter ϵ controls the noise intensity in terms of its variance, while δ represents the deviation from pure ϵ-privacy. The closer they are to zero, the stronger the privacy guarantees.

Limitations These results should be interpreted in light of several important limitations. First, the Google mobility data is limited to smartphone users who have opted in to Google's Location History feature, which is off by default. These data may not be representative of the population as whole, and furthermore their representativeness may vary by location. Importantly, these limited data are only viewed through the lens of differential privacy algorithms, specifically designed to protect user anonymity and obscure fine detail. Moreover, comparisons across rather than within locations are only descriptive since these regions can differ in substantial ways.

Data Availability The Google COVID-19 Aggregated Mobility Research Dataset used for this study is available with permission from Google LLC.

B.2 First Stage Results for Covariates

Table B.1: First Stage Results for Covariates – 1

VARIABLES	(1)	(2)	(3)
Cases $_{t-1}$	-0.119	-0.069	-0.224
	(0.636)	(0.596)	(0.512)
Cases $_{t-2}$	-0.416	-0.354	-0.051
	(0.753)	(0.717)	(0.658)
Deaths $_{t-1}$	-122.590**	-125.633***	-108.987***
	(45.508)	(40.947)	(37.613)
Deaths $_{t-2}$	54.836	59.764*	61.981**
	(34.975)	(33.427)	(27.370)
Cases $_{t-1}$ at	1.966	2.044	5.517
Neighbors	(4.821)	(4.207)	(4.573)
Cases $_{t-2}$ at	25.366***	24.222***	17.958***
Neighbors	(6.270)	(6.260)	(6.075)
Deaths $_{t-1}$ at	-1,625.307***	-1,538.692***	-1,175.645***
Neighbors	(356.618)	(346.804)	(294.015)
Deaths $_{t-2}$ at	203.146	201.314	214.978
Neighbors	(315.422)	(308.279)	(269.606)
Fiscal spending	-0.004*	-0.003	-0.004
	(0.002)	(0.002)	(0.002)
Share of vaccinated $_{t-2}$	3.798	7.963	4.707
	(14.339)	(13.505)	(14.963)
Travel Cont's: Screening	6.730**	7.250***	7.329***
	(2.510)	(2.299)	(2.347)
Quarantine	0.005	0.901	0.914
	(2.532)	(2.392)	(2.353)
Trageted Ban	-0.957	0.139	0.611
	(2.956)	(2.785)	(2.729)
Total Ban	-9.701***	-7.600**	-5.435*
	(3.585)	(3.210)	(2.976)
Observations	2,870	2,870	2,870
R-squared	0.716	0.728	0.754
Country FE's	•	•	•
Extensity	○	•	•
Intensity	○	○	•
Countries	41	41	41

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within weeks. • – included, ○ – excluded.

Table B.2: First Stage Results for Covariates – 2

VARIABLES	(1)	(2)	(3)
Income Support (≤50%)	-3.882	-3.449	-1.857
	(3.023)	(3.237)	(2.846)
Income Support (¿50%)	-2.302	-1.987	-0.869
	(2.464)	(2.582)	(2.359)
Debt Reief: Narrow	-1.628	-1.430	-2.346
	(1.643)	(1.691)	(1.482)
Broad	-2.197	-2.078	-2.544
	(2.136)	(2.133)	(1.844)
Info' Camp'n: Urging	1.371	1.243	0.942
	(1.709)	(1.580)	(1.618)
Coordinated	-0.435	1.318	1.719
	(2.229)	(1.893)	(1.862)
Testing: Symptoms + else	-0.200	0.056	-0.629
	(1.275)	(1.220)	(1.281)
w/ Symptoms	3.681*	2.714	2.369
	(1.896)	(1.733)	(1.864)
Open for All	7.147***	6.017***	5.739**
	(2.169)	(2.005)	(2.169)
Contact Tracing: Limited	-1.306	-1.038	-0.913
	(1.598)	(1.368)	(1.201)
Comprehensive	-0.649	-0.606	-1.731
	(1.427)	(1.240)	(1.214)
Masks: recommended	1.094	1.531	1.843
	(2.333)	(2.430)	(2.059)
specific places	3.231*	3.078*	3.197*
	(1.759)	(1.818)	(1.588)
public places	4.680**	4.615**	4.840***
	(1.953)	(1.977)	(1.744)
everywhere	3.760	5.029**	5.805**
	(2.287)	(2.451)	(2.237)
Observations	2,870	2,870	2,870
R-squared	0.716	0.728	0.754
Country FE's	•	•	•
Extensity	○	•	•
Intensity	○	○	•
Countries	41	41	41

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within weeks. • – included, ○ – excluded.

Table B.3: First Stage Results for Covariates – 3

VARIABLES	(1)	(2)	(3)
Vaccination: 1 group	-1.403	-1.466	-1.944
	(1.838)	(1.823)	(1.518)
2 groups	1.216	1.561	1.725
	(1.387)	(1.295)	(1.305)
3 groups	1.359	1.097	2.827*
	(1.324)	(1.297)	(1.505)
3+ groups	5.635	3.814	7.857**
	(3.833)	(3.262)	(3.810)
universal	7.104	4.828	5.007
	(6.505)	(6.063)	(6.852)
Elderly Protection: Recomm'	-1.674	-1.403	-2.794
	(1.842)	(1.884)	(1.730)
Narrow	-6.189***	-4.772***	-4.987***
	(1.961)	(1.763)	(1.495)
Extensive	-8.027***	-5.987***	-4.297**
	(2.183)	(2.145)	(1.699)
Mean Temperature	-0.098	-0.106	-0.124**
	(0.075)	(0.065)	(0.060)
Mean Humidity	0.089**	0.084**	0.079**
	(0.035)	(0.035)	(0.038)
Total Rainfall	-0.029	-0.028	-0.021
	(0.018)	(0.018)	(0.017)
Total Snowfall	-0.000	0.000	-0.000
	(0.000)	(0.000)	(0.000)
Observations	2,870	2,870	2,870
R-squared	0.716	0.728	0.754
Country FE's	•	•	•
Extensity	○	•	•
Intensity	○	○	•
Countries	41	41	41

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within weeks. • – included, ○ – excluded.

10.14754/CEU.2022.06

B.3 Second Stage Results for Covariates

Table B.4: Second Stage Results for Covariates – 1

	(1) industrial production	(2) manuf'ing production	(3) cons-truction	(4) retail trade	(5) CPI	(6) PPI manuf'ing	(7) unemployment rate
Cases $_{t-1}$	0.007***	0.007***	0.003	0.002	-0.000***	0.002**	0.001
	(0.002)	(0.002)	(0.005)	(0.003)	(0.000)	(0.001)	(0.001)
Deaths $_{t-1}$	-0.029	-0.057	-0.083	0.121*	0.002	-0.072***	-0.019*
	(0.072)	(0.086)	(0.133)	(0.060)	(0.004)	(0.010)	(0.009)
Cases $_{t-1}$ at Neighbors	0.283**	0.402***	0.075	0.259**	-0.022***	0.006	0.046
	(0.089)	(0.085)	(0.258)	(0.102)	(0.006)	(0.032)	(0.034)
Deaths $_{t-1}$ at Neighbors	-5.005	-6.687	-3.368	-5.061	-0.098	0.792*	-1.160*
	(3.371)	(3.711)	(6.812)	(2.930)	(0.199)	(0.377)	(0.507)
Fiscal Spending	-0.000	-0.008	-0.050***	-0.001	0.000	0.000	-0.005***
	(0.000)	(0.005)	(0.012)	(0.004)	(0.000)	(0.001)	(0.001)
Investment in Vaccines	0.350	0.346	2.256	1.663	-0.010	-0.078	0.241
	(0.614)	(0.457)	(1.825)	(1.101)	(0.053)	(0.123)	(0.242)
Investment in Healthcare	0.076**	0.066**	0.332***	-0.016	-0.002	-0.002	0.025**
	(0.026)	(0.026)	(0.081)	(0.032)	(0.003)	(0.008)	(0.011)
International Travel Controls							
Screening	6.693**	6.738**	4.381	11.743**	-0.107	-1.945**	-1.073***
	(2.179)	(2.446)	(5.329)	(4.865)	(0.137)	(0.685)	(0.211)
Quarantine	1.076	0.842	7.213	5.269	-0.246**	-2.479***	-0.343
	(1.738)	(2.097)	(5.703)	(2.938)	(0.086)	(0.735)	(0.189)
Targeted Ban	-1.226	-1.890	14.420*	5.681*	-0.399***	-3.411***	0.302
	(1.262)	(1.436)	(7.238)	(2.922)	(0.102)	(0.751)	(0.338)
Total Ban	-4.882**	-5.719**	19.837**	3.211	-0.375**	-2.955***	0.444
	(2.027)	(2.204)	(7.352)	(3.311)	(0.132)	(0.786)	(0.564)

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within months. • – included, ○ – excluded.

10.14754/CEU.2022.06

Table B.5: Second Stage Results for Covariates – 2

	(1) industrial production	(2) manuf'ing production	(3) cons- truction	(4) retail trade	(5) CPI	(6) PPI manuf'ing	(7) unemployment rate
Income Support Programs							
$\leq$ 50%	0.486	0.635	-0.849	0.965	-0.285**	0.045	0.449**
	(3.033)	(2.691)	(2.862)	(2.485)	(0.086)	(0.200)	(0.168)
$<$ 50%	-0.091	0.277	-0.340	4.432	-0.283**	-0.344	0.522
	(2.413)	(2.032)	(2.544)	(2.432)	(0.106)	(0.392)	(0.295)
Debt Relief							
Narrow	-0.253	0.271	1.810	-0.047	0.116	0.240	0.590*
	(0.571)	(0.729)	(2.188)	(1.324)	(0.103)	(0.427)	(0.280)
Broad	-1.510	-1.278	6.137	0.131	-0.075	0.208	0.753*
	(0.993)	(1.217)	(3.614)	(1.890)	(0.147)	(0.383)	(0.336)
Public Information Campaigns							
Officials Urging	-2.068*	-2.190**	1.547	-0.175	-0.079	1.286**	0.069
	(1.012)	(0.828)	(2.563)	(0.831)	(0.123)	(0.543)	(0.476)
Coordinated	1.401	2.213	-2.126	-0.098	-0.109	0.728	-0.032
	(2.196)	(2.220)	(3.782)	(0.966)	(0.096)	(0.728)	(0.403)
Observations	288	288	189	270	288	252	279
R-squared	0.670	0.687	0.683	0.784	0.887	0.838	0.810
Country FE	•	•	•	•	•	•	•
Countries	32	32	21	30	32	28	31

Notes: *** p<0.01, ** p<0.05, * p<0.1, standard errors in parentheses allowed to cluster within months. • – included, ∘ – excluded.

10.14754/CEU.2022.06

B.4 Historical Decompositions

Figure B.1: Historical Decomposition of Industrial Production

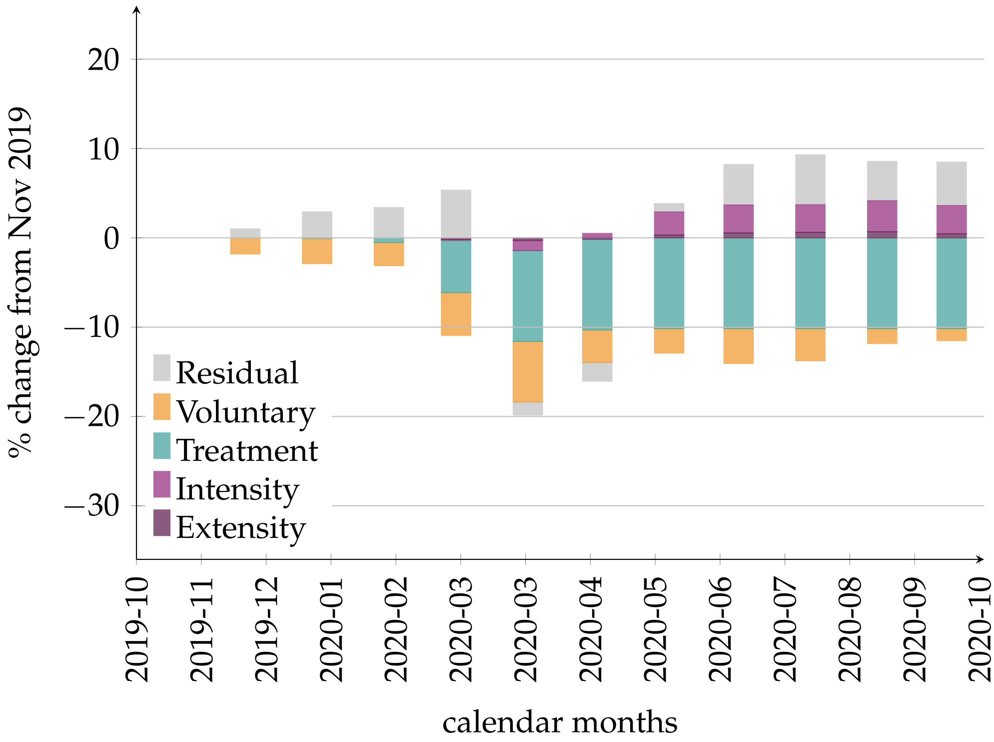

Figure B.2: Historical Decomposition of Manufacturing Production

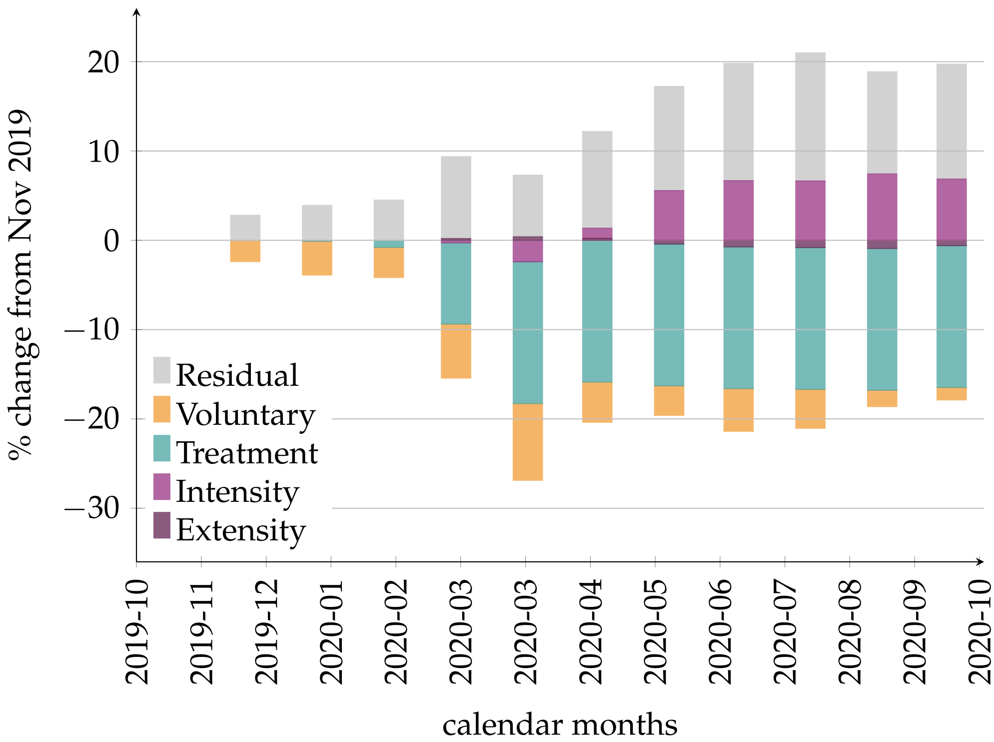

Figure B.3: Historical Decomposition of Construction Output

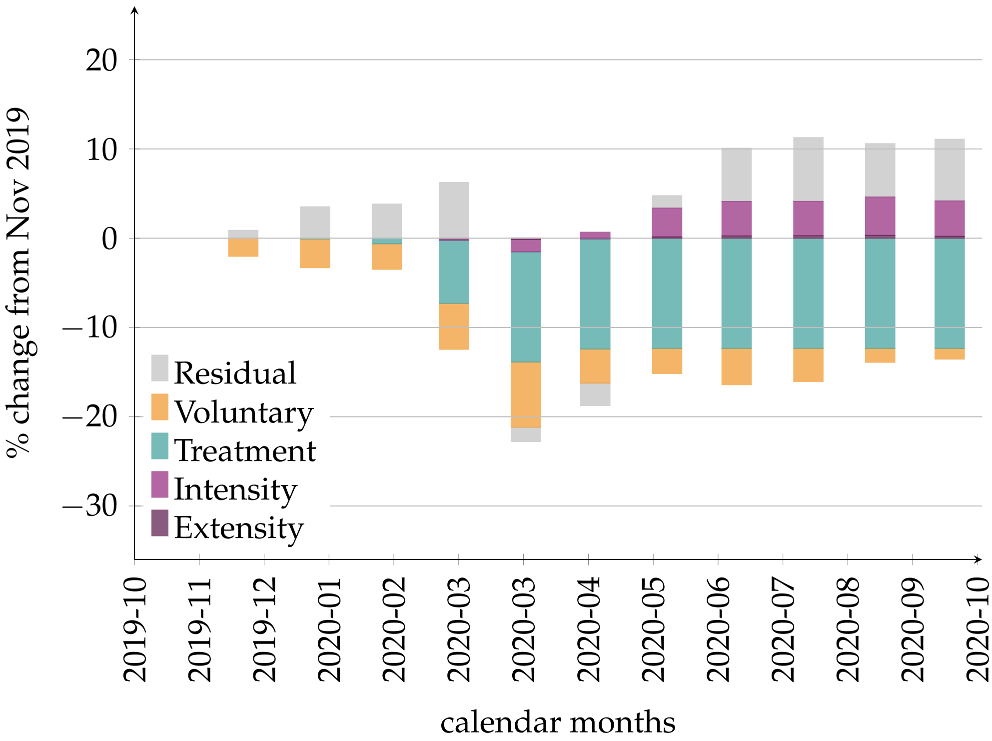

Figure B.4: Historical Decomposition of Retail Trade

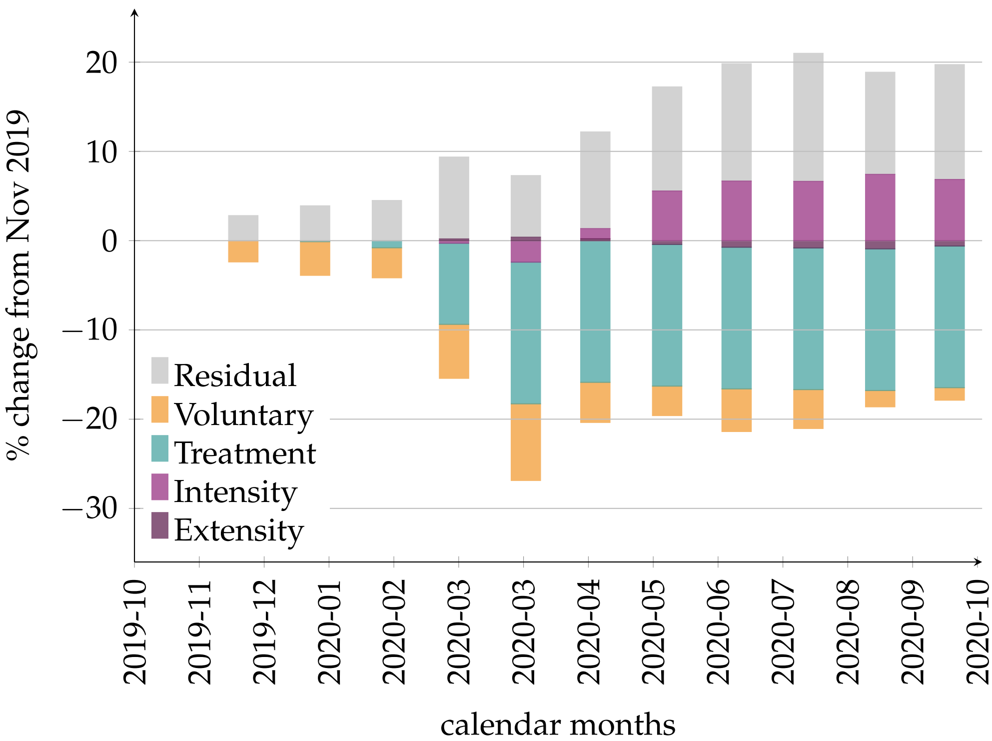

Appendix C

Appendix for Chapter 3

C.1 List of Industries

1. Farms
2. Forestry, fishing, and related activities
3. Oil and gas extraction
4. Mining, except oil and gas
5. Support activities for mining
6. Utilities
7. Construction
8. Wood products
9. Nonmetallic mineral products
10. Primary metals
11. Fabricated metal products
12. Machinery
13. Computer and electronic products
14. Electrical equipment, appliances, and components
15. Motor vehicles, bodies and trailers, and parts
16. Other transportation equipment
17. Furniture and related products
18. Miscellaneous manufacturing
19. Food and beverage and tobacco products
20. Textile mills and textile product mills
21. Apparel and leather and allied products
22. Paper products
23. Printing and related support activities
24. Petroleum and coal products
25. Chemical products
26. Plastics and rubber products
27. Wholesale trade
28. Motor vehicle and parts dealers
29. Food and beverage stores
30. General merchandise stores
31. Other retail
32. Air transportation
33. Rail transportation
34. Water transportation
35. Truck transportation
36. Transit and ground passenger transportation
37. Pipeline transportation
38. Other transportation and support activities
39. Warehousing and storage
40. Publishing industries, except internet (includes software)
41. Motion picture and sound recording industries
42. Broadcasting and telecommunications
43. Data processing, internet publishing, and other information services
44. Federal Reserve banks, credit intermediation, and related activities
45. Securities, commodity contracts, and investments

46. Insurance carriers and related activities
47. Funds, trusts, and other financial vehicles
48. Housing
49. Other real estate
50. Rental and leasing services and lessors of intangible assets
51. Legal services
52. Computer systems design and related services
53. Miscellaneous professional, scientific, and technical services
54. Management of companies and enterprises
55. Administrative and support services
56. Waste management and remediation services
57. Educational services
58. Ambulatory health care services
59. Hospitals
60. Nursing and residential care facilities
61. Social assistance
62. Performing arts, spectator sports, museums, and related activities
63. Amusements, gambling, and recreation industries
64. Accommodation
65. Food services and drinking places
66. Other services, except government

C.2 TFP Data

TFP series are available for a slightly less refined breakdown of 61 private sectors. For example, I observe TFP shocks for the real estate sector, while value-added is observed separately for housing and other real estate. I assume that the TFP shocks in the parent industry represent the TFP shock of its sub-industries. That means I impute the TFP of the higher-level aggregates to all of its breakdowns. The following table lists these industries.

Table C.1: Correspondence of 61 and 66 Industry Breakdowns

61 industries	66 industries
Retail trade	Motor vehicle and parts dealers Food and beverage stores General merchandise stores Other retail
Real estate	Housing Other real estate
Hospitals, and nursing,... ...and residential care facilities	Hospitals Nursing and residential care facilities

C.3 Data Description

The U.S. Bureau of Labor Statistics produces the official industry-level total factor productivity (TFP) measures for the United States (www.bls.gov/productivity). The official TFP measures differ from the integrated TFP measures due to differences in output concept and coverage. The official TFP indexes are constructed using a sectoral output concept rather than gross output. Sectoral output measures are calculated by subtracting from gross output those transactions that occur between establishments within the same industry, creating a measure of the value of goods and services consumed outside the industry. The official TFP indexes are not fully consistent with official GDP statistics; however, by using the sectoral output approach they eliminate double counting which reduces the influence of vertical integration on the industry estimates of TFP growth. The integrated TFP measures, on the other hand, are fully consistent with official GDP statistics; however, they are constructed using a gross rather than a sectoral output concept. Additionally, GDP statistics that conform to international guidelines include measures of economic activities for which there are no market prices. Such activities are excluded from official productivity statistics because output is generally not measured independently from inputs. There is an inherent tradeoff between having a set of accounts that add up to the best estimate of GDP and a set of accounts that generate the best productivity measures by industry. For more detail on the relative strengths and weaknesses of the official and integrated statistics, see Fleck, S., Rosenthal, S., Russell, M., Strassner, E., & Usher, L. (2012). A Prototype BEA/BLS Industry-Level Production Account for the United States, which also appears in D. W. Jorgenson, J. S. Landefeld, & P. Schreyer, Measuring Economic Sustainability and Progress.

Complete information of the methods and data underlying these measures can be found be at the Productivity Handbook of Methods. Information is also available on the BLS Productivity Website (http://www.bls.gov/productivity). For further information, contact the Division of Major Sector Productivity. Source: Bureau of Labor Statistics, Office of Productivity and Technology.

C.4 Proof of Proposition 3.1

Proof. Applying aggregation rules (3.8) and (3.9) to (3.6) yields:

$$\Delta \mathbf{c}_t = \sum_{k=0}^{\infty} \left[\sum_{h=k/\varphi}^{(k+1)\varphi-1} A^h \right] \Delta \mathbf{z}_{t-k}$$

This can be simplified by the introduction of $B_k(A, \delta)$, such that $\delta = 12/\varphi$:

$$\begin{aligned} B_k(A,\delta) &= \sum_{h=k/\varphi}^{(k+1)\varphi-1} A^h = \left(A^{k\varphi} - A^{(k+1)\varphi}\right)(I-A)^{-1} \\ &= A^{k\varphi}(I - A^{\varphi})(I-A)^{-1}. \end{aligned}$$

■

C.5 Derivation of Model Predcitions

Each firm i equipped by production technology of equation (3.2) has the following profit maximization exercise to solve in each period τ:

$$\max_{l_{i\tau}, x_{ij\tau}} \text{Profit}_{i\tau} = \frac{p_{i,\tau+1}}{\rho} y_{i,\tau+1} - w_\tau l_{i\tau} - \sum_{j=1}^{n} p_{j\tau} x_{ij\tau}, \tag{C.1}$$

where ρ is the period discount factor. Long and Plosser (1983) show that this discount factor is a constant.

Assuming rational expectations, and $\mathbb{E}[e^{z_{i,\tau+1}}] = 1$, the following input demand functions maximize the profit of producer i in τ:

$$l_{i\tau} = p_{i,\tau+1} \frac{\lambda_i y_{i,\tau+1}}{\rho w_\tau}, \tag{C.2}$$

$$x_{ij\tau} = p_{i,\tau+1} \frac{a_{ij} y_{i,\tau+1}}{\rho p_{j\tau}}. \tag{C.3}$$

Plugging (C.2) and (C.3) into (3.2) yields the following supply equation in logarithms after rearrangements:

$$\ln p_{i,\tau+1} = -z_{i,\tau+1} + \lambda_i \ln w_\tau \sum_{j=1}^{n} a_{ij} \ln p_{j\tau} + \ln K_i^{-1} \tag{C.4}$$

Such that $K_i \equiv \lambda_i^{\lambda_i} \prod_{j=1}^{n} a_{ij}^{a_{ij}} / \rho$.

The representative consumer maximizes its utility for each period independently because it has no means of wealth accumulation. Therefore it spends all of its income earned as wages. Therefore its demand for each good i is the following:

$$c_{i\tau} = \gamma_i \frac{w_\tau}{p_{i\tau}}, \tag{C.5}$$

given the labor endowment of the consumer is fixed, it is normalized to 1.

Taking logarithms of equation (C.5) and rearrangement to $\ln p_{i\tau}$ yields:

$$\ln p_{i\tau} = \ln \gamma_i + \ln w_\tau - \ln c_{i\tau} \tag{C.6}$$

Market clearing sets prices of demand and supply equal, allowing the substitution for $\ln p_{i\tau}$ in the price equation (C.4) using equation (C.6). This substitution yields the following equation in vector notation after rearrangements and shifted backward one period:

$$\mathbf{c}_\tau = A\mathbf{c}_{\tau-1} + \mathbf{z}_\tau + \mathbf{k}, \tag{C.7}$$

where $\mathbf{k}$ is a function of model parameters, therefore a constant. This equation shows that the vector of value-added series follows an AR(1) process with the TFP process standing in as innovations. Using the lag operator (L) equation (C.7) can be rewritten the following MA(∞) model:

$$\mathbf{c}_\tau = (1 - \mathrm{L})^{-1}\mathbf{z}_\tau + \kappa, \tag{C.8}$$

which is equivalent to equation (3.4) such that $\kappa = (1 - \mathrm{L})^{-1}\mathbf{k}$. ■