

Spatial-SIR with Network Structure and Behavior: Lockdown Rules and the Lucas Critique[‡]

Alberto Bisin Andrea Moro

April 12, 2022

First version: August 4, 2020

Abstract

We introduce a model of the diffusion of an epidemic with demographically heterogeneous agents interacting socially on a spatially structured network. Contagion-risk averse agents respond behaviorally to the diffusion of the infections by limiting their social interactions. Schools and workplaces also respond by allowing students and employees to attend and work remotely. The spatial structure induces local herd immunities along socio-demographic dimensions, which significantly affect the dynamics of infections. We study several non-pharmaceutical interventions; e.g., i) lockdown rules, which set thresholds on the spread of the infection for the closing and reopening of economic activities; ii) neighborhood lockdowns, leveraging granular (neighborhood-level) information to improve the effectiveness public health policies; iii) selective lockdowns, which restrict social interactions by location (in the network) and by the demographic characteristics of the agents. Substantiating a “Lucas critique” argument, we assess the cost of naive discretionary policies ignoring agents and firms’ behavioral responses.

[‡]Please check our websites for an updated version of this paper. Bisin: New York University, wp.nyu.edu/albertobisin/, alberto.bisin@nyu.edu. Moro: Vanderbilt University andreamoro.net, andrea@andreamoro.net.

1 Introduction

In this paper, we develop a diffusion model of infection to study the effectiveness of different public health policies on contagion dynamics. We extend the Spatial-SIR model introduced in [Bisin and Moro \(2021\)](#) to account for a demographic and a network structure of the social contacts driving the diffusion process. This allows us to interact population and location heterogeneity with the spatial dimension of contacts. We also account for agents' and firms' behavioral responses to the diffusion of the epidemics and to policy interventions.

In the model, agents belong to three demographic types: Young (students/workers), Not employed, Old. Each day, the types potentially visit different locations in a structured network: the City, School/Work locations, and Home, which differ by contact density. Consequently, the model displays heterogeneous rates of diffusion of the epidemic by demographic type and across the network of locations. Furthermore, depending on the spread of the epidemic, contagion-risk averse agents may reduce their social interactions, and firms may instruct their workers to operate remotely. We calibrate the model to reproduce several stylized facts about the diffusion of SARS-CoV-2 epidemics. We calibrate agent types and family compositions to match the U.S. demographics, particularly the family size distribution by age. Next, we simulate the calibrated model.

The model simulation identifies several important factors relevant to understanding the heterogeneous rates of diffusion of the epidemic, highlighting that local herd immunities may form not only geographically (the focus in [Bisin and Moro \(2021\)](#)) but also along socio-demographic dimensions depending on the network structure. For instance, the City originates more infections than School/Work and School/Work more than Home, even though Home is the densest location and the City is the least dense. This result occurs because the ranking of the locations in terms of density is opposite to the ranking in terms of size, therefore herd immunity is achieved earlier at Home and School/Work than in the City. Similarly, the epidemic is primarily concentrated on the Young because they are exposed at School/Work, where there are more contacts than at Home. These effects, which we say arise from *local herd immunity* are akin to a selection effect, whereby the epidemic spreads selectively by location and demographic structure.

Our analysis also illustrates the relevance of the indirect effects induced by agents' and firms' behavioral responses across locations and demographic types. For instance, the Youngs' behavioral responses, preventing exposure in School/Work locations, have a sizable substantial effect on the Old in the City, where most of the Old get infected.

We use this model to study the effects of several public health policy interventions. Simulations allow us to evaluate the outcomes of these interventions with respect to, e.g., the peaks of the first and second waves of infection and the spread of the infection by demographic structure. Specifically, we emphasize public health policy rules, that is, articulated interventions on an epidemic's whole course possibly developing over various waves. We study, e.g., lockdown rules, where the public health authorities set thresholds levels of active infection cases determining when social interactions are restricted and when they are

allowed. Lockdown rules interact in rather subtle ways with the dynamics of herd immunity. When the lockdown is placed too early with respect to the spread of the epidemic, the fraction of infected is too small, herd immunity is too far out, and a second wave of infections at reopening reaches a peak higher than the first one. If reopening is delayed to occur at a lower fraction of active cases than the one that triggered the lockdown, a strategy we labeled *cautious reopening*, herd behavior is more advanced upon reopening, substantially dampening the second wave.

We also simulate the ability of the government to exploit granular information on the epidemics by imposing lockdowns at the neighborhood level based on the local progression of the infection. Neighborhood-level lockdown policies can flatten the curve of infections and generate fewer infected at steady state, if compared to a city-wide lockdown. However, the effectiveness of these localized lockdowns crucially depends on how segregated neighborhoods are, in particular on the level of interaction across neighborhoods caused by agents visiting School/Work locations.

Moreover, we find that when neighborhoods are heterogeneous in their demographic structure, the location of the initial cluster is a crucial determinant of the infection dynamics and steady-state outcome. When the infection starts in a neighborhoods with smaller average household size, the infection progresses more slowly; therefore, it has time to spread to nearby neighborhoods, generating a higher city-wide peak and steady-state levels of infections. When the infection starts in neighborhoods with large families, it affects locally a large fraction of people, achieving quickly a local herd immunity that spares, to some extent, the rest of the city from a higher level of infections.

The model’s network and demographic structure also allow us to study various selective lockdown policies, where lockdowns are imposed by the location of social interactions and/or by the agents’ demographic characteristics. We concentrate on two such policies. In the first one, we limit the lockdown to the City, and we let firms/schools decide whether to operate remotely. In the second one, we limit the lockdown to the Old, a policy aiming at reducing total fatalities while bearing limited economic costs (the Old are not economically active but suffer a higher fatality rate than other demographic groups). Both these policies have the important property that they are much less costly economically because workers/students remain active. This benefit is traded-off with health-outcome costs at the steady-state or at the peak (relevant when there is a health care resource constraint, such as hospital/ICU beds). While these selective lockdowns have positive direct effects in terms of economic costs, indirect effects across locations and/or demographic types could, in principle, substantially limit their advantage over general lockdowns. This outcome does not appear to occur in our simulations. The City-Only lockdown does not induce a much larger fraction of infected at steady-state than the general lockdown, while the Old-Only lockdown is very successful in limiting Old agents’ deaths.

Last but not least, we illustrate the implications of what economists refer to as the Lucas Critique in the context of epidemiological models. Policy evaluations and interventions disregarding that agents’ and firms’ “decision rules vary systematically with changes in

the structure of series relevant to the decision-maker” (Lucas (1976)) might lead to policy decisions that are very costly in terms of their effects on the epidemic dynamics. The costs of disregarding behavioral responses are clearer if we consider that the lockdown and opening thresholds are calibrated to flatten the infection curve to avoid hitting a constraint on the available health care resources. In this case, our simulations imply under-utilization of health care resources after lockdown (that is, the lockdown will turn out to be stricter than necessary) and possibly an over-run of these resources after reopening.

2 Epidemic by Spatial, Demographic, Network, and Behavior Characteristics

We introduce a model of the dynamics of an epidemic which, while stylized, contains several important components needed to analyze the effects of different forms of lockdowns. Our goal is to maintain a parsimonious structure while replicating the features of general epidemics that are relevant for policy analysis, to produce relatively robust qualitative results.¹

We build on the Spatial-SIR model we have introduced in Bisin and Moro (2021). In the Spatial-SIR, agents are located and move in a two-dimensional space. With some probability, they become infected when they get close to an infected agent.² In this paper, we introduce heterogeneity in agents’ type (by age and occupation status) and in the characteristics of the social spaces where they interact so that they interact in distinct networks.³

We model agents responding to the dynamics of the epidemic, by choosing to limit their social contacts, and we assume that some schools/firms adopt behavioral responses similar to the agents’, that is, they choose to shut down operations or to allow workers to work remotely, to limit infections in the School/Work locations.⁴

¹Frontier research in epidemiology has proposed models of contagion, extending the classic SIR model in many directions, including detailed descriptions of the geographic or network characteristics of the agents’ social interactions - see, e.g., Eubank et al. (2004) and the research at GLEAM project, mobs-lab, and the Imperial’s college MRC Centre for Global Infectious Disease Analysis. While these are large-scale models fundamentally aimed at forecasting with accuracy and precision, their complexity makes it difficult to uncover the main driving forces of the epidemics and policy interventions’ effects.

²Several recent contributions to the study of the SARS-CoV-2 epidemic in economics have restricted their epidemiology component to the SIR model, without modeling the spatial dimension of the epidemics; see e.g., Fernandez-Villaverde and Jones (2020), Atkeson (2020), Eichenbaum et al. (2020), Giannitsarou et al. (2020), Engle et al. (2021), Rowthorn and Toxvaerd (2012) Weitz et al. (2020). As exceptions, Antràs et al. (2020), Brotherhood et al. (2020b), Cuñat and Zymek (2020), and Glaeser et al. (2020) introduce connections between different geographical units.

³Along these lines, Acemoglu et al. (2020), Alfaro et al. (2020), Babajanyan and Cheong (2020), Baqaee et al. (2020a), Baqaee et al. (2020b), Bognanni et al. (2020), Brotherhood et al. (2020a), Fischer (2020), Giagheddu and Papetti (2020), Gollier (2020) add a demographic structure like we do; Azzimonti et al. (2020) adds a network structure; Ellison (2020) allow for heterogeneity of the contact process between subpopulations.

⁴Models of rational agents limiting contacts to reduce the risk of infection include Fenichel (2013) Weitz et al. (2020) in epidemiology (see also Funk et al. (2010) and Verelst et al. (2016) for surveys); and Geoffard

2.1 State transitions

Let $\mathcal{S} = \{S, A, Y, R, D\}$ denote the individual state-space, indicating Susceptibles, Asymptomatic, sYmptomatic, Recovered, and Dead.⁵ We assume the following transitions: i) a Susceptible agent in a location within distance p from the location of an Asymptomatic becomes infected with probability π ;⁶ ii) an Asymptomatic agent infected at t , at any future period, can become sYmptomatic with probability ν , or can Recover with probability ρ ; iii) a sYmptomatic agent can Recover or Die at any period with probabilities ρ and δ , respectively; iv) Dead and Recovered agents never leave these states.

2.2 Demographic and network structure

There are three *types of* locations: the *City* at large, a set of *School/Work* locations shared by subsets of agents, and *Home* locations shared by family members. Agents differ by age and occupation in their access to these three types of locations.

There are three *demographic types* of agents:⁷ the *Young*, the *Not employed*, and the *Old*. The Young represent students/workers who spend part of their day in School/Work, that is, in a location where they interact daily with the same individuals. The Not employed represent adult individuals below 65 years of age that do not have access to a School/Work location. The Old are not working but differ from the Not employed because of their higher fatality rate and because they tend to live at Home with individuals of the same type. The demographic structure interacts with the location structure in that the Old do not visit School/Work locations.

2.3 Social interactions

Each day is divided into three periods dedicated to the location where agents interact: Period 1 is devoted to visiting the City, Period 2 to visiting the School/Work location and Period 3 to stay at Home. However, agents might choose or be structurally restricted from visiting the City and/or School/Work and hence stay Home, depending on their demographic

and Philipson (1996), Goenka and Liu (2012), Acemoglu et al. (2020), Aguirregabiria et al. (2020), Argente et al. (2020), Bethune and Korinek (2020), Farboodi et al. (2020), Fernandez-Villaverde and Jones (2020), Greenwood et al. (2019), Engle et al. (2021), Toxvaerd (2020) (as well as several of the papers cited in Footnote 2, modeling spatial extensions of SIR) in economics. Empirically, several papers use cellular phone data to document how fear of contagion has affected mobility or show that the economy deteriorated substantially even before or independently of lockdown orders. See, for example, Alfaro et al. (2020), Aum et al. (2020), Barrios et al. (2020), Bartik et al. (2020), Cicala et al. (2020), Coibion et al. (2020), Goolsbee and Syverson (2020), Gupta et al. (2020), Kahn et al. (2020), Rojas et al. (2020).

⁵There is no accepted denomination for the form of compartmental model we adopt in this paper. Alternative denominations would include A-SIRD and Stochastic Symptom-Response SIRD, following, respectively, Gaeta (2020); Leung et al. (2018).

⁶We assume sYmptomatic agents are isolated at home or in the hospital, therefore not contagious.

⁷The age structure of the population has specific relevance for the SARS-CoV-2 epidemic because the case-fatality rate is very skewed, higher for older agents.

Table 1: Location by time of day

Location type	Period 1	Period 2	Period 3
City	All except sYmptomatics	No-one	No-one
School/Work	No-one	Young except sYmptomatics	No-one
Home	sYmptomatics	sYmptomatics, Not employed, Old	All

and/or health state.⁸ SYmptomatic agents must stay at Home in all three periods. If not sYmptomatic, all agents visit the City in Period 1, and all Young agents visit the School/Work location in Period 2, when the other agents stay Home. Table 1 summarizes the locations visited by agents by the time of day.

Agents move randomly in the City: Every day $t = [0, T]$, agents travel distance μ toward a random direction of $d \sim U[0, 2\pi]$ radians. By doing so, they potentially meet new individuals every day; therefore, μ models the speed at which agents find new contacts, potentially in a different state than their previous neighbors. Each Young is randomly assigned to a School/Work location and does not move within it. School/Work locations are far from each other (there is no possibility of contagion across School/Work locations). Finally, agents are assigned to families located in Homes that are far from each other and far from School/Work locations.

The distinction we introduce across the three locations abstractly captures different relative densities in agents’ social interactions. The School/Work location represents a non-random location that is visited every day by the same set of people and is denser than the City (see the next section for details on how location density is calibrated). Home represents the location with the tightest interactions: agents from the same family have identical Home locations; that is, they are necessarily in contact with each other. On the other hand, at School/Work and in the City, the contagion may occur only if people are close in space.

2.4 Behavior

In this section, we model agents responding to the epidemic dynamics by choosing to limit their contacts. Following Engle et al. (2021), we introduce a reduced-form behavioral response, represented by a function $0 \leq \alpha(I_t) \leq 1$, representing the fraction of agents who

⁸Agents might be restricted to visit some location also as a policy choice. We discuss these restrictions in the next section.

Table 2: Calibrated parameter values

Parameter	Value
Number of people	26,600
Initially infected	30
Family size and age distr.	match 2018 ACS
School/Work population (avg.)	100
Group quarter population	50
School/Work side length	0.003304
Contagion radius	0.00805
Distance traveled in City	0.028175
Prob. of becoming symptomatic	0.09
Prob. of recovery	0.05
Fatality rate Young	0.000533
Fatality rate, Old	0.00533
Contagion probability, Home	0.064
Contagion probability, City School/Work	0.032
Behavioral responses, agents	$\phi : .88, \underline{I} : .001$
Behavioral response, firms	$\phi : .88, \underline{I} : .00062$

Note: ACS: American Community Survey. Group quarters are places where people live or stay in group living arrangements. The School/Work side length is relative to the City’s side length, which is normalized to 1. See Appendix B for details.

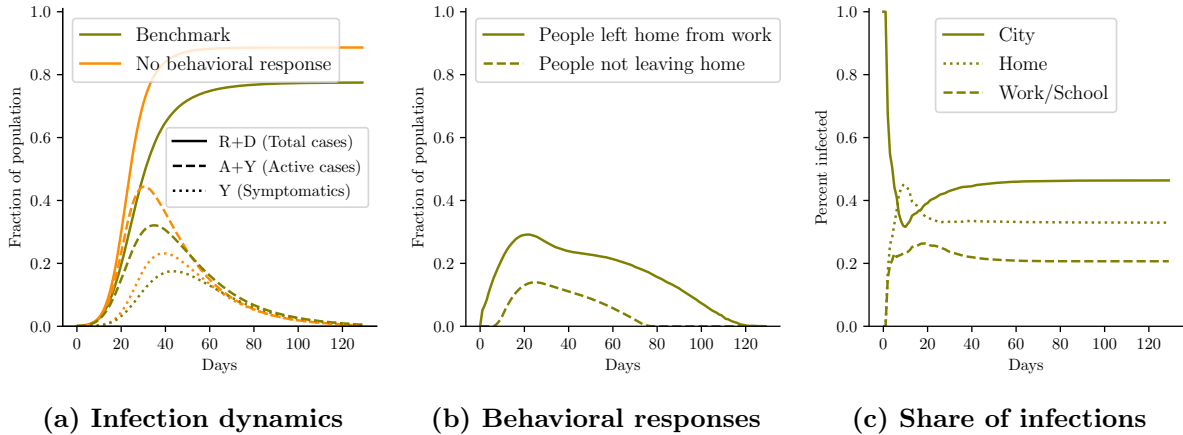
choose to interact in the City as a function of the number of infected I_t :⁹

$$\alpha(I_t) = \begin{cases} 1 & \text{if } I_t \leq \underline{I} \\ \left(\frac{\underline{I}}{I_t}\right)^{1-\phi} & \text{if } I_t > \underline{I} \end{cases} . \quad (1)$$

We rank the agents (and School/Work locations) by contagion-risk aversion. Agents choosing not to interact in the City (and School/Firms leaving agents at Home) are the most contagion-risk averse. We also assume some firms and schools adopt behavioral responses

⁹In our simulations, we assume the behavioral response depends on the fraction of Asymptomatics because the sYmptomatics are isolated at Home. Since the fraction of Asymptomatics is not observable, the behavioral response could only depend on the number of sYmptomatics, as a proxy; with Rational Expectations, however, the agents know (rationally infer) the equilibrium map from Y to A , say $A(Y)$, possibly with noise; see Bisin and Moro (2020) or Gans (2020).

Figure 1: Simulation of the benchmark model



Note: Infection dynamics, Benchmark model (green lines), and model without behavioral responses (orange lines)

similar to the agents’ according to (1), that is, they choose to shut down operations or to allow workers to work remotely, to limit infections at school or work.

3 Calibration and simulation

We use data from the 2018 American Community Survey (ACS) to classify agents by type and allocate them to households replicating the size and age distribution of American households. We calibrate the three locations’ relative density to match the data on the distribution of contacts reported by [Mossong et al. \(2008\)](#). Transition probabilities between states match various SARS-CoV-2 parameters from epidemiological studies, notably, e.g., [Ferguson et al. \(2020\)](#). We calibrate contagion rates and the agents’ movement speed in the City to fit the epidemic’s initial growth rate, proxied by the growth rates of deaths in Lombardy. Finally, behavioral response parameters are calibrated following [Engle et al. \(2021\)](#).¹⁰ Appendix B describes the calibration in detail. Table 2 reports the parameter used in our main specification.

Results from the simulation of the benchmark model are illustrated in Figure 1; Figure 1(a) reports the epidemic dynamics; Figure 1(b) describes the behavioral response, reporting the fraction of agents (out of the total population) choosing not to visit the city (dashed line) and the fraction of agents left home from work because their workplace shuts down; and

¹⁰We acknowledge the uncertainty in the literature concerning many epidemiological parameters pertaining to this epidemic. Our approach is less damaging when aiming at understanding mechanisms and orders-of-magnitude rather than at precise forecasts.

Figure 1(c) the locations of the infections. Table 3(a) reports the steady-state outcomes.¹¹, including the locations of infection by agent type and the percent of the population in three different states at the steady-state.

Generally, our simulation of the model demonstrates that *local herd immunity* - which in Bisin and Moro (2021) we argue explains the dependence of the dynamics of the epidemic from geographical characteristics - also explains the dependence of the epidemic from location and demographic characteristics. With respect to location, the City originates more infections, 46 percent (vs. 21 percent in School/Work location, and 33 percent at Home) even though Homes and School/Work locations are denser because more contacts with different people are generated in the City. Herd immunity is achieved *locally* earlier at Home and at School/Work than in the City. The infection dynamics by location shows a larger fraction of infections at Home than at School/Work and in the City after about ten days, when the City relative infection rate starts to pick up.¹²

More Young are infected because of their exposure to more contacts at School/Work than at Home. Old and the Not employed are hit differently because of family structure: the latter live in households with relatively more young individuals that are infected at School/Work locations. These effects are akin to a selection effect, whereby the epidemic spreads selectively by location and demographic structure.¹³

It is interesting to compare the simulation of the benchmark model to a model simulated without behavioral responses (see the orange lines in Figure 1(a) and Table 3(b)). The total fraction of Recovered and Dead at the Steady-state is 0.77 in the benchmark model but would have been .88 without the behavioral response. The fraction of infections occurring at Home is almost ten percentage points lower without behavioral response. The symmetric opposite effect is mostly concentrated at School/Work rather than in the City, whose fractions of infection increase, respectively, by 9 and 1 percentage points. The total rise in infections at School/Work affects mostly the Young, who get infected less at Home (from 31 to 20 percent) and more at School/Work (from 30 to 43 percent). At the steady-state, the proportion of Young Recovered or Dead increases from 82 percent to 94 percent.

The most important implication resulting from this simulation is the indirect effect of the behavioral response across locations and demographic types. Specifically, the behavioral responses of (mostly) the Young, (mostly) in the School/Work location have a sizable important effect on the Old in the City (the Old get infected mostly in the City, over 70 percent, with or without behavioral response). Notably, the compensating effect of higher infections at Home with the behavioral response is of reduced importance for the Old because of family composition (the Old tend to live with other Old). At the steady-state, the proportion of Old Recovered or Dead increases from 63 percent to 73 percent.¹⁴ The impact of behavioral

¹¹All results are obtained averaging over 20 replications of the simulation, randomizing over the initial location of individuals in the City and at School/Work

¹²We set the location of infection at $t = 0$ in the City.

¹³See Gomes et al. (2020) and Britton et al. (2020) for related theoretical analyses.

¹⁴A simulation - which we do not report in the text - where we only allow for a behavioral response by

Table 3: Benchmark model, infection location and steady state outcomes

	Steady-state outcomes			Share of infections		
	D	R	Peak A+Y	City	W/S	Home
All	0.008	0.766	0.323	0.463	0.205	0.332
Young	0.004	0.819	0.349	0.394	0.296	0.309
Not employed	0.004	0.728	0.290	0.537	-	0.463
Old	0.029	0.599	0.249	0.726	-	0.274
(a) Benchmark model						
All	0.010	0.876	0.446	0.472	0.294	0.234
Young	0.005	0.933	0.486	0.372	0.427	0.201
Not employed	0.005	0.815	0.383	0.626	-	0.374
Old	0.036	0.719	0.365	0.777	-	0.223
(b) Benchmark model without behavioral responses						

Note: The table reports steady-state outcomes as a fraction of total population and the number of infected by location as a fraction of the total number of infected $D + R$.

response on the Not employed is instead moderate because their family composition induces them to be infected more at Home.

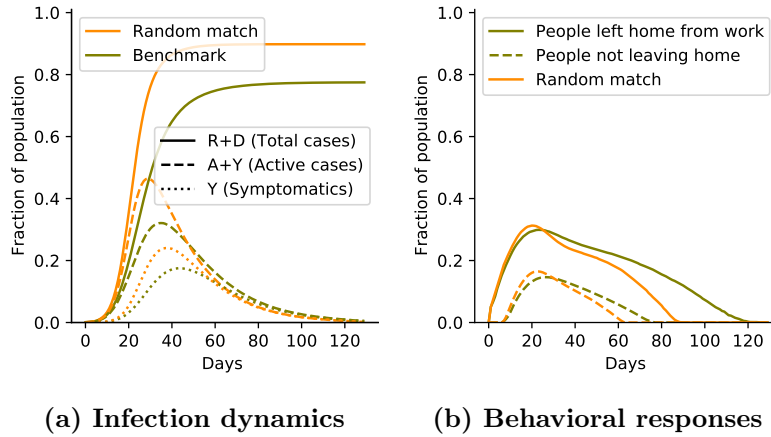
Finally, we assess the relevance of the spatial dimension, which induces local herd immunity, in conjunction with the network structure, which distinguishes City, School/Work and Home. To this end, we compare our benchmark spatial model with one characterized by random matching of agents, that is, with no spatial structure, but with the network structure. In the random matching model, each agent can be infected by anyone in the City and in the School/Work every day.¹⁵

Results are reported in Figure 2. The spatial component of our analysis matters first of all in that the steady state of infections is greater with random matching, as shown in Figure 2(a). Importantly, this effect is not crucially dependent on behavior, as it is

firms, has a much smaller effect on the Old, suggesting that the main impact on the Old goes through the reduction of infections in the City.

¹⁵Specifically, we expand the contagion radius so that the contagion circle includes the entire City. Similarly, we expand the contagion ratio in each School/Work to include the entire establishment. We recalibrate the contagion probability in the City and in each School/Work so that the average number of people one infected agent infects is kept constant across the benchmark spatial and random matching models.

Figure 2: Benchmark model comparison with random matching



	Fraction infected (R+D) in steady-state			
	City	Work/School	Home	Total
All	0.355 (-.004)	0.152 (-.007)	0.390 (+.133)	0.898 (+.124)
Young	0.283 (-.042)	0.234 (-.010)	0.440 (+.185)	0.957 (+.133)
Not employed	0.459 (+.066)	-	0.359 (+.020)	0.818 (+.086)
Old	0.523 (+.067)	-	0.227 (+.055)	0.750 (+.122)

(c) Fraction infected by location in random matching model and difference with benchmark model

Note: The figures illustrate the infection dynamics, Benchmark model (green lines), and random-matching model (orange lines). The table reports the number of infected in the random matching model as a fraction of the total population of each type, by location of infection. Difference between random matching model and benchmark model in parenthesis.

clear from Figure 2(b) that the behavioral reactions of both firms (leaving workers home) and agents (staying home from the City) is only slightly more pronounced under random matching. It is rather the local herd immunity phenomenon in our benchmark model which is fundamentally responsible for the reduction of the steady state fraction of infected with respect to the random matching model. The mechanism operates however through a rather subtle implication of the interaction of the spatial component of our benchmark model with the network structure. This is apparent from the distribution of infections by type of agent and by location reported in the table underneath Figure 2 (panel 2(c)). Indeed, even though random matching operates only in the City, and all types interact in the City,

the extra infected are largely at Home (+.133 percentage points) and are mostly the Young (+.133). Furthermore, the effect is larger for the Old (+.122) than for the Not employed (+.086), even though these two types are exposed in the same way to interactions in the City.¹⁶ In fact, the first order effect of the lack of local herd immunity in the City with random matching is to speed up the infection process in the early stages of the epidemic. This faster spreading of infections in the City in turn speeds up the infections at Home, at an even accelerating rate due to the fact that infections at Home are much faster than at any other location since agents live in close quarters (at zero distance from each other in the simulations). Of course, agents who are infected at Home are not infected in other locations, in particular not in the City. This explains why the effects of random matching are seen mostly at Home rather than in the City. This effect is more prominent for the Young because the mechanism we are illustrating is accelerated in their case by the fact that they can also be infected in Work/School. Furthermore, these effects are larger for the Old than for the Not employed because the demographic composition of the Home is such that the Old live mostly with other Old. In the benchmark model, the effect of local herd immunity is tamed for the Not employed because they are more likely to live with the Young, who infect them at Home even when there are no active cases locally. When local herd immunity is removed by random matching, the infections at Home increase more for the Old than for the Not employed.

4 Lockdown policies


This section explores the effects of various non-pharmaceutical interventions along the lines of those proposed and implemented by various administrations during the COVID-19 outbreak. Our goal is to exploit the spatial nature of the model, the structure of the network of interactions, and the (distinction by) demographic characteristics of the agents to analyze the different impact of these policies both in the aggregate and across types.

We study general lockdown rules, that is, interventions where the lockdown and reopening rules depend on the spread of the infections. We also study selective lockdowns, where only the City is locked down, or only Old individuals are restricted in movement across locations (either confined either at Home or in quarantine facilities of different sizes). In all the cases we study, policies overlap with the behavioral response; that is, they are applied to the population of agents and firms randomly, not accounting for their behavioral response. In other words, lockdown policies may apply to agents (resp. firms) that would have chosen to self-isolate (resp. shut down) even in the absence of the policy.¹⁷

¹⁶We thank one referee for pointing out this piece of evidence.

¹⁷This analysis is, to our knowledge, the first example of an application of Lucas Critique argument in this context; see Section 5 for a more articulated discussion and interesting examples.

Figure 3: Lockdown and reopen forever



	Steady-state outcomes			Share of infections		
	D	R	Peak A+Y	City	W/S	Home
(a) No lockdown	0.008	0.766	0.323	0.463	0.205	0.332
(b) 10% rule	0.007	0.703	0.158	0.363	0.225	0.412
(c) 5.0% rule	0.007	0.712	0.150	0.362	0.231	0.407
(d) 3.5% rule	0.007	0.715	0.159	0.363	0.233	0.404

Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population under different lockdown policies. The horizontal lines indicate the proportion of infected triggering a lockdown. The table reports steady-state outcomes as a fraction of total population and the number of infected by location as a fraction of the total number of infected $D + R$.

4.1 A general lockdown of the City and School/Work

We study the effects of a generalized lockdown of the City and 50 percent of School/Work locations. We consider policies where the lockdown is applied when the fraction of active cases, $A + Y$, reaches a given fraction of the population. We report results regarding different policies regarding when to reopen the City and the School/Work locations operating remotely. The results reveal the trade-off between flattening the the curve of active cases and delaying the reaching of global herd immunity.

4.1.1 Reopen forever after active cases return to the fraction which triggered the lockdown

We start by studying the case in which both City and School/Work reopen forever after active infections return to the fraction which triggered the lockdown (10, 5, and 3.5 percent). In this case, because the lockdown may have delayed the population reaching herd immunity, infections may increase again after reopening.

The results are displayed in Figure 3. We notice that the flattening of the infection curve is substantial. In all cases, the peak of active cases with lockdown is less than half than without lockdown. The reduction of Dead and Recovered at the steady-state is substantial. Importantly, the lockdown has rather large effects on the Old (not reported in the table). The fraction of Dead and Recovered decreases under all policies by about ten percentage points for the Old and only about half as much for the Young (and Not employed). While either lockdown reduces the fraction of Dead at the steady-state of only one-tenth of 1 percent, this reduction is four to five times as big for the Old.

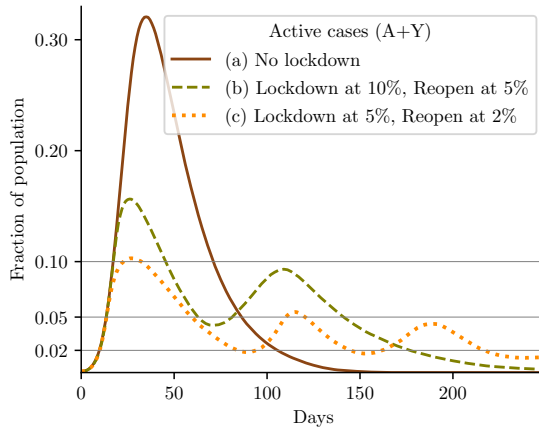
Notably, the lockdown strategy interacts in rather subtle ways with the dynamics of herd immunity. When the lockdown occurs too early (5 percent in our simulation), the fraction of infected is too small, herd immunity is too far out, and at reopening a second wave of infections reaches a peak higher than the first one.

4.1.2 Cautious reopening

We study rules in which reopening is set more cautiously: a 5 percent reopening threshold for the 10 percent lockdown and a 2 percent reopening threshold for the 5 percent lockdown. This is an interesting health policy strategy because, as we have seen in the previous sections, the lockdown strategy interacts with herd immunity dynamics, and an early reopening can induce a high second wave of infections.

Results are in Figure 4. Relative to the case in which reopening occurs at the same threshold as the lockdown, herd behavior is more advanced at reopening, and the second wave is substantially dampened. The peaks of the second wave occur at about 8 percent as opposed to 13 percent of the population, for the 10 percent lockdown. The result is even more striking for the 5 percent lockdown. In this case, with reopening at 5 percent, the second wave was higher than the first (15 percent). With cautious reopening at 2 percent, instead, the second wave is dampened to 5 percent. Interestingly, however this triggers another lockdown and hence a third wave whose peak is below 5 percent. Relatedly, the fraction of Dead and Recovered at steady-state is lower with cautious reopening, about 3 and 8 percentage points, respectively, for the 10 and the 5 percent lockdowns. Cautious reopening, we conclude, is particularly effective for the 5 percent lockdown, where reopening at 5 percent induced a large second wave and a fraction of Dead and Recovered at steady-state larger than with a less restrictive lockdown at 10 percent (with reopening at 10 percent).

Figure 4: Cautious reopening



	Steady-state outcomes			Share of infections		
	D	R	Peak A+Y	City	W/S	Home
(a) No lockdown	0.008	0.766	0.323	0.463	0.205	0.332
(b) 10-5% rule	0.007	0.674	0.158	0.339	0.235	0.426
(c) 5-2% rule	0.006	0.630	0.104	0.292	0.254	0.453

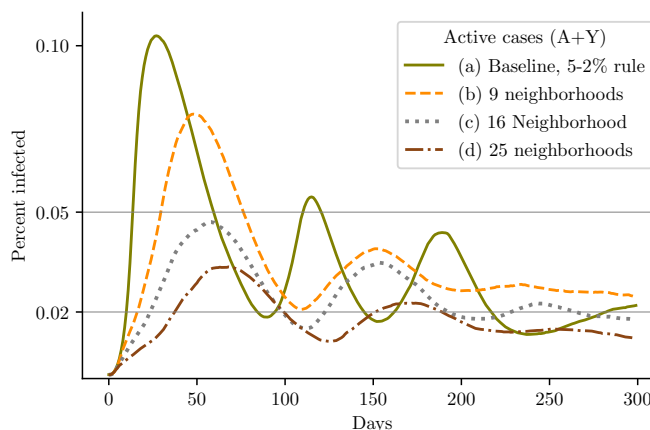
Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population under different lockdown policies. The horizontal lines indicate the proportion of infected triggering a lockdown or a reopening. The table reports steady-state outcomes as a fraction of total population and the number of infected by location as a fraction of the total number of infected $D + R$.

4.2 Local (Neighborhood) Lockdowns

In this section we allow public health authorities to define lockdown rules depending on local neighborhood-level information about the epidemic. The aim is to investigate whether such granular information is relevant to improve the effectiveness of NPI policies.

Specifically, we divide the city in $S = 9, 16, 25$ symmetric square-shaped neighborhoods. Home and School/Work are neighborhood specific and agents residing in a given neighborhood visit a School/Work in that same neighborhood. Agents visit the City initially in a random location of the neighborhood where they reside, but we let them move across neighborhood borders. The initial cluster of infection is in the south-west corner of the city. We simulate the effect of the cautious lockdown and reopening policy described in Section 4.1.2, but with thresholds and lockdowns that are neighborhood-specific. Results are shown

Figure 5: Neighborhood-specific lockdowns - Different number of neighborhoods




	Steady-state outcomes			Share of infections		
	D	R	Peak A+Y	City	W/S	Home
(a) Baseline 5-2% rule	0.006	0.630	0.104	0.292	0.254	0.453
(b) 9 neighborhoods	0.006	0.631	0.087	0.285	0.262	0.453
(c) 16 neighborhoods	0.004	0.462	0.065	0.257	0.286	0.457
(d) 25 neighborhoods	0.003	0.335	0.047	0.251	0.288	0.461

Note: The figure illustrates the dynamics of active cases ($A+Y$) as a fraction of the total population under a cautious reopening policy (5-2% rule) applied at the neighborhood level. The green solid line (Baseline) corresponds to a policy applied city-wide. The horizontal lines indicate the proportion of infected triggering a lockdown or a reopening. The table reports steady-state outcomes as a fraction of total population and the number of infected by location as a fraction of the total number of infected $D+R$.

in Figure 5: neighborhood-specific NPI policies are indeed very effective, as - other things equal - they induce substantial flattening of the epidemic infections and a smaller number of total infected and fatalities at steady state. Not surprisingly, but still importantly, this effect is monotonic in the number of neighborhoods, indicating the importance on the part of public health authorities of collecting and leveraging in policy design more and more granular information about infections.

In Figure 6 we allow a fraction of agents to attend School/Work outside of the neighborhood where their Home is located and they are initially placed. In both the 9 and 25 neighborhood simulations - respectively in Figures 6(a) and 6(b) - mixing of School-Work has a very large negative effect on the ability of NPI to flatten of the epidemic. Indeed, in both cases just a 5% of student/workers outside the neighborhood is enough to increase

Figure 6: Neighborhood-specific lockdowns - School/Work Mixing



Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population under a cautious reopening policy (5-2% rule) applied at the neighborhood level, under different levels of School/Work mixing across neighborhoods. The orange dashed line in panel (a) and the brown dash-dotted line in panel (b) correspond to the lines in Figure 5. The horizontal lines indicate the proportion of infected triggering a lockdown or a reopening.

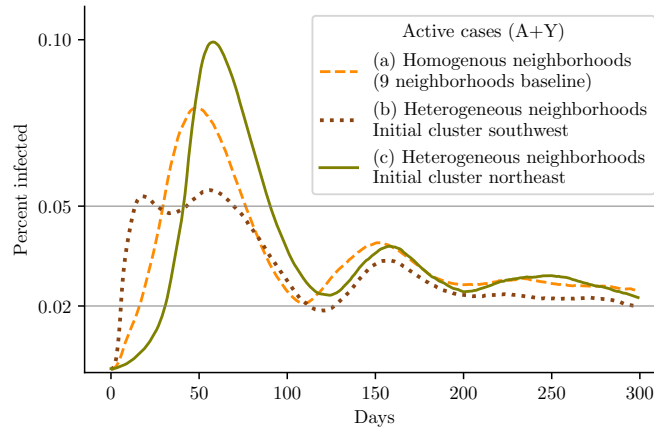
dramatically the peak of infections, keeping NPI policies unvaried.

4.2.1 Heterogeneous neighborhoods

To investigate the role of neighborhood heterogeneity in the dynamics of an epidemic, we simulate a city divided in 9 neighborhoods as in the orange line of Figure 5, but we place Homes with larger households - which we can interpret as poorer - closer to the southwest corner of the city (neighborhood 1, with average household size equal to 8.3), and progressively smaller household as the location of residence gets closer to the northeastern corner, where only singles live (neighborhood 9, with average household size about 1). We however maintain neighborhoods density constant across all neighborhoods, to isolate the effect of household size from the effect of density we studied in Bisin and Moro (2021), and we report simulations where School/Work locations for all agents are in their neighborhood of residence.

In Figure 7 we show that the dynamics of infections spreading from a cluster near the southwest corner (with largest average household size) display an earlier but lower peak than the corresponding dynamics starting from the northeast corner (with smallest average household size). This is an interesting composition effect of the heterogeneity in average household size. Figure 8(a) show that an infection spreading from the neighborhood with the largest average household size (neighborhood 1) has an early and large peak in this neighborhood, which spreads into into the other neighborhoods but dies off quickly, inducing local herd immunity at the level of the neighborhood before spreading to lower average

Figure 7: Neighborhoods with different household size




	Steady-state outcomes			Share of infections		
	D	R	Peak A+Y	City	W/S	Home
(a) Homogenous neighb.	0.006	0.631	0.087	0.285	0.262	0.453
(b) Heterogeneous, SW cluster	0.006	0.562	0.067	0.283	0.267	0.450
(c) Heterogeneous, NE cluster	0.006	0.639	0.104	0.307	0.248	0.446

Note: The figure illustrates the dynamics of active cases ($A+Y$) as a fraction of the total population under a cautious reopening policy (5-2% rule) in a city with 9 neighborhoods differing by household size. The dashed yellow line corresponds the 9-neighborhoods line in Figure 5. The table reports steady-state outcomes as a fraction of total population and the number of infected by location as a fraction of the total number of infected $D + R$.

household size neighborhoods (like 8 and 9). On the other hand, an infection spreading from the neighborhood with the lowest average household size (neighborhood 9) in Figure 8(b) starts slowly because this neighborhood includes only singles, and no-one can be infected at home. The infection builds up to induce relatively high peaks of infections at the same time in the neighborhood with mid-average household size, inducing an higher aggregate peak of infections. The different speed at which the infection spread in the two cases - when the cluster starts from the highest or the lowest average household size - is clearly shown comparing the state of the dynamics of infection after 10 days, as in the top two panels of Figure 8. The differences are dramatic even though, it is important to remark, the densities of the neighborhoods are maintained constant.

Figure 8: Dynamics by neighborhood




Note: The top two panels display the geographic location of individuals color-coded by their status at day 10. The bottom panels display the dynamics of active cases ($A + Y$) as a fraction of the total population in selected neighborhoods under a cautious reopening policy (5-2% rule) in a city with 9 neighborhoods differing by household size. Neighborhood 1 is located in the southwest corner of the city and is the neighborhood with highest average family size. Neighborhoods 2 . . . 8 are located progressively further away from neighborhood 1 and include households with progressively smaller family size. Neighborhood 9 is located in the northeast corner and is the neighborhood with the lowest average family size. The brown dotted line in panel (a) and the green solid line in panel (b) correspond to the aggregate behavior in Figure 7.

4.3 Selective Lockdowns

This section presents two selective lockdown policies. In the first one, we limit the lockdown to the City, and we let firms/schools decide whether to operate remotely. In the second one, we limit the lockdown to the Old, a policy aiming at reducing total fatalities while bearing limited economic costs (the Old are not economically active but suffer a higher fatality

Figure 9: City-only lockdown



	Steady-state outcomes			Share of infections		
	D	R	Peak A+Y	City	W/S	Home
(a) No lockdown	0.008	0.766	0.323	0.463	0.205	0.332
(b) Baseline 10-5% rule	0.007	0.674	0.158	0.339	0.235	0.426
(c) City-only lockdown	0.006	0.650	0.186	0.264	0.296	0.440

Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population under different lockdown policies. The horizontal lines indicate the proportion of infected triggering a lockdown or a reopening. The table reports steady-state outcomes as a fraction of total population and the number of infected by location as a fraction of the total number of infected $D + R$.

rate than other demographic groups). Both these lockdown policies have the important property that they are much less costly, economically, as worker/students remain active. This consideration needs to be traded-off with additional costs in terms of health outcomes at the steady-state or at the peak, e.g. if the society has a constraint with respect to health care resources like hospital/ICU beds. Furthermore, while these selective lockdowns have positive direct effects in terms of economic costs, indirect effects across locations and/or demographic types could, in principle, substantially limit their advantage over general lockdowns.

4.3.1 City-Only lockdown

Results for the City-Only lockdown are reported in Figure 9, with respect to a 10 percent lockdown with cautious reopening at 5 percent. Interestingly, in this case, with respect to the

general lockdown, the first peak in $A+Y$ is higher (18.6 percent rather than 15.8), but there is no significant second wave. Furthermore, the total fraction of Dead and Recovered at the steady-state is two percentage points lower. This is striking, as the City-Only lockdown is much less costly, economically, than the general lockdown. According to these simulations, the City-Only lockdown has extra health costs inasmuch as the first wave might be too high with respect to a possible constraint in terms of health care resources. In other words, indirect effects, e.g., from City to Home, do not appear to limit the advantages of the City-Only lockdown in our simulations.¹⁸

4.3.2 Old-Only lockdown

Results for the Old-Only lockdown are reported in Figure 10, simulating a 5 percent lockdown rule with a period of cautious reopening at 2 percent. Olds are isolated in their Homes, where they are in contact with their families. Comparing it with the case in which the lockdown is general (that is, includes the Young and the Not Employed as well), it is remarkable that the number of stationary Dead between the Old is more than halved, from almost 2 percent to less than 1 percent. The peak of active cases remains rather high, as the Young and the Not Employed keep interacting in the City and the Young also at School/Work. This high peak does not necessarily affect possible constraints on health care resources if the Old use them relatively more intensely. As for the City-Only lockdown, our simulations of the Old-Only lockdown suggest that indirect effects, e.g., from Young to Old or from School/Work to Home, do not substantially limit its advantage.

The simulations we have just reported refer to the case in which the Old are locked-down at Home, where they interact with Young and Not Employed agents. In Figure 10(b) we report $D + R$ at steady-state for different simulations where the Old are locked-down in nursing homes of different sizes (the dot is the case where the Old are locked-down at their Home). It is noteworthy that the number of Dead increases steeply with the size of the nursing homes, reaching 2 percent for homes with 50-agents each. It is also interesting that the lockdown of the Old at Home is equivalent in terms of infection and fatality rates to one where the Old are locked-down in nursing homes of about size 10. This is a consequence of the indirect interactions across locations hurting the more susceptible demographic group.

5 An Exercise in the Lucas Critique


This section illustrates the implications of what economists refer to as the *Lucas Critique*, in the context of epidemiological models.¹⁹ We have already shown, in the previous section,

¹⁸As we repeatedly noted, we do not aspire to quantitatively precise prediction because the calibrated model we simulate is very stylized.

¹⁹From Lucas (1976). In its original formulation, the Critique is summarized as follows:

“Given that the structure of an econometric model consists of optimal decision rules of economic agents and that optimal decision rules vary systematically with changes in the structure

Figure 10: Old-only lockdown



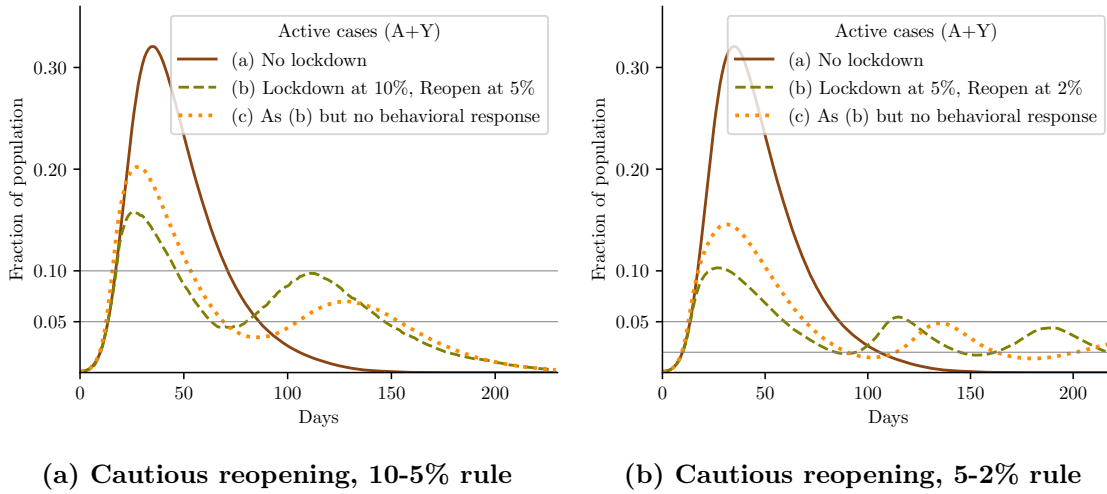
	Steady-state outcomes			Share of infections			
	D	R	Peak A+Y	City	W/S	Home	
(a)	All, no lockdown	0.008	0.766	0.323	0.463	0.205	0.332
	Old, no lockdown	0.029	0.599	0.249	0.726	-	0.274
(b)	Old, General lockdown	0.019	0.397	0.050	0.609	-	0.391
(c)	All, Old-only lockdown	0.004	0.635	0.241	0.436	0.149	0.415
	Old, Old-only lockdown	0.008	0.167	0.053	0.228	-	0.772

Note: The figure on the left and the table illustrate, respectively, the dynamics of active cases ($A + Y$) and the steady-state outcomes, under a cautious lockdown policy (5-2% rule) when Old are isolated at their home. The figure on the right illustrates the fractions of infected among all and Old (as a fraction of the population) when the Old are isolated in nursing homes of different size (the dots represent outcomes when Old are isolated at their homes).

that the behavioral response of agents and firms to the dynamics of an epidemic is bound to have important effects on the dynamics themselves. We now discuss how policy evaluations and interventions disregarding that agents' and firms' "decision rules vary systematically

of series relevant to the decision-maker, it follows that any change in policy will systematically alter the structure of econometric models."

Figure 11: Naive policymaker



Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population with naive policymaking. The horizontal lines indicate the proportion of infected triggering a lockdown or a reopening.

with changes in the structure of series relevant to the decision-maker” might lead to policy decisions that are very costly in terms of their effects on the dynamics of the epidemic.

5.1 Policymaker error


We simulate the prediction error of a policymaker evaluating different lockdown policies while not anticipating the behavioral response of agents and firms to the dynamics of the epidemic. We interpret the dynamics without behavioral response as the policymaker’s prediction of the effect of the lockdown and the case with response as the actual outcome of the lockdown.

We report the case of a lockdown at 10 percent with reopening at 5 and a lockdown at 5 percent with reopening at 2; respectively in Figure 11(a) and 11(b) (we do not report tables for these simulations). In both cases, the policymaker predicts a higher peak than it actually occurs. On the other hand, he/she will face an earlier and higher than predicted second wave.

5.2 Policymaker error when recalibrating the reopening policy

In this section we postulate a more articulate procedure to evaluate different lockdown policies on the part of a policymaker who does not anticipate the behavioral response of

Figure 12: Re-calibrated reopenings



Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population with naive policymaking. The dashed horizontal lines indicate the initial (naive) target peak. The dash-dotted horizontal lines indicate the recalibrated reopening. Panel (a) Lockdown threshold policy at 3.5% with a target peak at 12.9 (dashed line). Results from adjusting the reopening policy at 1.1% (dotted line) after the first peak. Panel (b): Lockdown threshold policy at 5% with a target peak at 14.7% (dashed line). Results from adjusting the reopening policy at 4.2% after the first peak (dotted line).

agents and firms. Consider a policymaker who evaluates lockdown interventions assuming no behavioral responses, but chooses the reopening threshold as a consequence of a policy re-valuation after the lockdown, still disregarding the behavioral response. More specifically, consider the case where, at the re-evaluation stage, the policymaker naively attributes the lower-than-predicted peak of infections to the effect of fewer people visiting the City and School/Work locations, generating lower contact rates and slower diffusion of the epidemics. The policymaker, however, continues to disregard the fact that behavioral responses are a function of the number of cases and re-evaluates the reopening policy assuming that people visiting the City and School/Work locations remain at the (lower) level observed at the re-evaluation.

The adverse effects and the costs of this procedure are clearly seen in the case in which a policymaker sets up a 3.5 percent lockdown rule with the goal of not experiencing a peak active cases over about 13 percent, as predicted by the model without behavioral response (see Figure 12(a)). When reaching the peak, which occurs at about 8 percent, the policymaker re-evaluates its reopening policy, re-calibrates its model and computes a new reopening rule. The recalibration and the new reopening rule are computed postulating still no behavioral response upon reopening. In our simulations, we run the model to find the

reopening rule that would generate 13 percent active cases at the peak in the policymaker’s recalibrated model.


The error induced by the policymaker recalibrating the reopening policy is evident in the Figure. Because there is not enough herd immunity at the first peak, the policymaker delays reopening until the fraction of active cases is about 1 percent. Nonetheless, the recalibration misses the behavioral response of agents and firms, who increase the contact rate in both the City and the School/Work due to the low infection rates. As a consequence, the second wave is much higher than the first and ends up reaching a peak of 14 percent, 9% higher than the government’s original target. If the lockdown and opening threshold are calibrated to avoid hitting a constraint on the available health care resources, this simulation implies that health care resources will be under-utilized after lockdown (that is, the lockdown will turn out to be stricter than necessary) and they will instead be over-run after reopening, with potentially high human costs.

Interestingly, the costs associated with this policy-making procedure disregarding behavioral responses are less apparent in the case of a lockdown at 5 percent so as to avoid a peak higher than about 15 percent of active cases (see Figure 12(b)). It is still the case that, after observing a smaller than expected first peak, the policymaker, expects a low infection rate for the future, induced by (what she expects to be a constantly) lower contact rate. She then decides to reopen too soon, when cases are low, and hence the behavioral response of agents and firms increases the contact rate. As a consequence, the epidemic picks up, giving rise to an unexpected second wave of infection as agents and firms. However, in this case, the second peak after recalibration is higher than the first, but near the 15 percent target, not higher as in the previous case. While the recalibration misses the behavioral response, leading agents and firms to increase the contact rate after the first peak, it also misses the herd immunity generated by this behavior. The effect of herd immunity is larger because initial lockdown is set at 5 percent rather than at 3.5 percent, because of the relatively higher first peak of infections, and because, after the government re-evaluation, when active cases decline more people than the government expects visit the City and get infected.

The policymaker error when naively recalibrating the reopening policy in an environment where lockdowns are neighborhood-specific illustrates another interesting dimension of application of the Lucas critique to NPI policies. In the model with 9 homogeneous neighborhoods and neighborhood-specific lockdown policy, the naive reopening, after naively recalibrating the model, induces a second peak overshooting the target of 6.3%, computed with the same criteria used in the previous cases, by over 40% of its level, reaching a second peak of infection at about 9% infected, as shown in Figure 13(a).

A qualitatively similar outcome is obtained if we repeat the exercise with heterogeneous neighborhoods differing by average household size, with the initial outbreak in the corner of the city with the largest household size (the same structure used in the simulations reported in Figure 8(a)). The peak after the government’s recalibration overshoots the target by 26%

Figure 13: Lucas Critique with Neighborhood-Specific Lockdowns



Note: The figure illustrates the dynamics of active cases ($A + Y$) as a fraction of the total population with naive policymaking. The dashed horizontal lines indicate the initial (naive) target peak. The dotted horizontal lines indicate the recalibrated reopening. Panel (a): model with 9 homogenous neighborhoods; Panel (b) model with 9 heterogeneous neighborhoods (initial cluster in the southwest corner).

(see Figure 13(b)).²⁰

6 Conclusions

In this paper, we developed a parsimonious stylized model which, in our objectives, adds several dimensions of interest to the standard epidemiological components of general economic models in which the policy trade-off between economic activity and epidemic diffusion is evaluated, as, e.g., in Atkeson (2020), Alvarez et al. (2020), Jarosch et al. (2020), Kaplan et al. (2020). Both the geographic and the network dimension appear very relevant in several empirical analyses of the effects of the SARS-CoV-2 epidemic, which tend to concentrate on economic activity granular level.²¹ The model is designed to qualitatively and quantitatively identify the role of the interactions between several fundamental “forces” driving the dynamics of an epidemic, like (global and local) herd immunity, network, and demographic het-

²⁰When the outbreak starts in the neighborhood with the smallest average household size, the recalibration does not result in overshooting.

²¹See e.g., along these lines, Alstadsæter et al. (2020), Chetty et al. (2020), Crossley et al. (2020), Davenport et al. (2020), and De Paula et al. (2020).

erogeneity, behavioral responses of agents and firms to infections, and Non-Pharmaceutical policy interventions. In this respect, we stress that the role of Lucas Critique arguments, which identify the adverse effects and the costs of policy interventions disregarding behavioral responses, is potentially first order, qualitatively and quantitatively, depending on the calibration.

A Appendix: Theoretical Structure of SIR and Spatial-SIR

We refer to Appendix A in [Bisin and Moro \(2021\)](#) for a formalization of the Spatial-SIR we extend in this paper. The extension we adopt in this paper adds demographic and network structure as described but does not alter the mechanics of the epidemic state transitions.

B Appendix: Calibration

We calibrate the parameters of Spatial-SIR to the dynamics of the SARS-CoV-2 epidemic. We acknowledge the substantial uncertainty in the literature with respect to even the main epidemiological parameters pertaining to this epidemic. As we noted in the introduction, this is less problematic when aiming at understanding mechanisms and orders-of-magnitude rather than at precise forecasts. The parameters we choose in the calibration of the aggregate model are summarily reported in [Table 2](#).

Geography

The City is a square of area equal to 1. We choose $N = 25600$ so that our simulations converge in a reasonable time (see [Bisin and Moro \(2021\)](#) for a description of how the model scales in size). People are located in the City by randomly drawing their x and y coordinates independently from a Uniform distribution $\sim U[0, 1]$. At all $t > 0$, individuals are relocated at distance μ from their location at $t - 1$, in a direction drawn randomly from a Uniform distribution $\sim U[0, 2\pi]$. School/Work locations are squares of size calibrated as described below. The location of individuals within School/Work locations is randomly drawn as in the City, but does not change over time. Homes have no dimension. All School/Work and Home locations are far from each other and far from the City to prevent contagion across different location types, across different schools/workplaces and across different homes.

We calibrate the relative density of the three locations to match the data on the distribution of contacts reported by [Mossong et al. \(2008\)](#).²² In this data, on average, people make 13.5 contacts per day. Furthermore, the total number of contacts is distributed as follows: 34 percent at work or at school, 34 percent in other locations besides home.²³ Assuming that interactions at Home necessarily induce contacts between family members, and with an average family size of 3.3 from the ACS data, this implies 2.3 contacts at home.²⁴ Of the remaining 11.2 contacts, according to the data in [Mossong et al. \(2008\)](#), about half occur at Work and School and half in what we have categorized as the City. Because a fraction of the population is structurally restricted not to visit the School/Work location (the Old and

²²They survey mixing patterns on a sample of the population in eight European countries. We do not exploit cross-country heterogeneity in this section and consider averages.

²³We ignore the 7 percent of contacts occurring in “multiple” locations.

²⁴This is lower than implied by [Mossong et al. \(2008\)](#), but it’s an upper bound for the size distribution of families from ACS we use.

the Not employed), we need to rescale the number of contacts per day in each location. In the ACS this fraction is 15 percent and therefore we calibrate that a young and occupied agent (a schoolchild or a worker) receives 17.6 percent (100/85 percent) more contacts than anyone in other locations. We allocate therefore the 11.2 contacts per day not occurring at Home proportionally, 6 at School/Work and 5.2 in the City. We therefore calibrate the contagion radius to match 5.2 contacts on average in the City, given its population. We keep the same contagion radius in the School/Work locations, and calibrate the size of such locations so as to generate 6 contacts, on average.²⁵

Agent types and family composition

We use data from the 2018 American Community Survey (ACS) to classify agents by type and allocate them to households replicating the size and age distribution of American households. Specifically, we calculate from the ACS the distribution of family compositions by Young/Old age of respondents not living in Group Quarters, and construct households of sizes and Young/Old compositions replicating this distribution. The age cutoff for being classified as Old is 65. We then create households of size 50 containing either only Young or only Old individuals to replicate the fraction of Young and Old ACS respondents living in group quarters (2.9 percent of Old and 2.4 percent of Young). In the 2018 ACS, approximately 16 percent of individuals are above 65 years of age. We divide the remaining population to reflect the proportion of the population that are young and either employed or in school, 65 percent, and define the rest, 19 percent, as Not employed. We randomly assign all Young to a number of School/Work locations such that their average size is 100 individuals.

Infection and contact rates


We calibrate the contagion rate π and the mean travel distance μ to match the daily growth rates of the dynamics of fatalities, g_t , observed in the first 35 days of epidemics in Lombardy, Italy (data from Italian government, Dipartimento di Protezione Civile available at <https://github.com/pcm-dpc/COVID-19>); see Figure 14, left panel.²⁶ We chose to match data from Lombardy, the location where one of the first large clusters of infections occurred. We assume that contagion rates are twice as large at Home locations than elsewhere and that Symptomatic individuals remain contagious at Home but are not contagious elsewhere.

We calibrate the daily fatality rates by type imposing that the Old have 10 times higher fatality rates than the Young and scale their average so that the aggregate infection fatality rate (IFR) is 1 percent, consistently with, e.g., Ferguson et al. (2020).

²⁵With a population of 26,600 individuals on a square of unit side, the contagion radius we obtain is 0.00805. With 100 agents at each School/Work location on a square, the side we obtain is 0.0547.

²⁶Since the number of infections is not observed, we match the growth rates of infection in the model with the growth rate of deaths in the data. This is justified when, as we assumed, the case fatality rate is constant and Death follows infection after a constant lag on average.

Figure 14: Growth rates of infections in simulation, fatalities in Lombardy, Italy



Note: Solid green line: average of 20 simulations of the calibrated benchmark model. Dotted yellow line: COVID-19 fatalities growth rate in the Lombardy region (Italy), 3-day moving average, starting from February 29, 2020. Data from the Italian government, Dipartimento di Protezione Civile, available at <https://github.com/pcm-dpc/COVID-19>.

Transitions away and between the infected states, A, Y, D, R .

The probability to transition away from the Asymptomatic state A , is $\rho + \nu$. We assumed agents are infective only in state A (we assumed that all sYmptomatic agents reduce to zero social contacts). The average time an agents stays in state A is then $T_{inf} = \frac{1}{\rho + \nu}$. We set $\rho + \nu$ to match a theoretical moment which holds exactly at the initial condition in the basic SIR model. Recall R_0 denotes the number of agents a single infected agent at $t = 0$ infects, on average. Let g_0 denote the growth rate of the number of infected agents at $t = 0$. Then, in SIR,

$$\frac{(\mathcal{R}_0 - 1)}{T_{inf}} = g_0. \quad (2)$$

For the current SARS-CoV-2 epidemics, R_0 is reasonably estimated between 2.5 and 3.5.²⁷ The daily rate of growth of the infection g is estimated = .35 by Kaplan et al. (2020).²⁸ This implies, from Equation (2), that T_{inf} is between 4 and 7 days (respectively for \mathcal{R}_0 between 2.5 and 3.5). Ferguson et al. (2020) use 6.5 days. We set $\rho + \nu = .14$, so that $T_{inf} = \frac{1}{\rho + \nu} = 7$.²⁹ Furthermore, the average time from infection to death or recovery is reasonably estimated to be 20 days; see e.g., Ferguson et al. (2020). Therefore we set $\rho = .05$ so that $1/\rho = 20$. This implies $\nu = .09$.

²⁷For details, see Footnote 32 in Bisin and Moro (2021)

²⁸Alvarez et al. (2020) have .2; Ferguson et al. (2020) have 0.15.

²⁹We thank Gianluca Violante for suggesting this calibration strategy.

The case-fatality rate, the probability of death if infected, is estimated between .005 and .01; see e.g., [Ferguson et al. \(2020\)](#). Since agents remain sYmptomatic in the model, before Recovering, on average $\frac{1}{\nu} = 11$ days, we set the probability of Death for a sYmptomatic, δ , to be 0.001.³⁰

Behavioral responses

We set $\phi = 0.88$ as estimated by [Engle et al. \(2021\)](#) using Swine flu data, and assume people start responding to the spread of the contagion very soon by setting $\underline{I} = 0.01$. The fraction of firms that choose to let workers work remotely is calibrated using the same functional form (eqn. 1). We use the same calibration for $\phi = 0.88$, but we calibrate \underline{I} so that the fraction of School/Work locations that choose to shut down is 50 percent when the fraction of Asymptomatics is 20 percent (close to the peak of the Asymptomatics curve in most of our simulations³¹). We believe 50 percent to be a reasonable guess for the peak given that 24 percent of the population is reported by the 2018 ACS to be in school and that [Alon et al. \(2020\)](#) compute from ATUS data that about 25 percent of workers work in highly telecommutable jobs; see also [Dingel and Neiman \(2020\)](#), who report that the share of jobs that can be done at home goes from 50 percent to 30 percent, across metropolitan areas in the U.S.

Initial conditions

At time $t = 0$ we set 30 individuals in Asymptomatic state; all others are Susceptibles. the Asymptomatics at $t = 0$ are those initially located in a position closest to City location $[x = 0.25, y = 0.25]$. We assume they have been infected in the City. All others are Susceptibles.

³⁰We assume fatalities cannot occur to an agent less than 3 days before she becomes sYmptomatic and that every infected individual recovers with certainty after 100 days.

³¹Solving $(x/0.20)^{0.12} = 0.5$ gives $x = 0.00062$, therefore we use this value

References

- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston**, “A Multi-Risk SIR Model with Optimally Targeted Lockdown,” National Bureau of Economic Research 2020. (Cited on pages 4 and 5)
- Aguirregabiria, Victor, Jiaying Gu, Yao Luo, and Pedro Mira**, “A Dynamic Structural Model of Virus Diffusion and Network Production: A First Report,” CEPR Discussion Paper No. DP14750 2020. (Cited on page 5)
- Alfaro, Laura, Ester Faia, Nora Lamersdorf, and Farzad Saidi**, “Social Interactions in Pandemics: Fear, Altruism, and Reciprocity,” Technical Report, National Bureau of Economic Research 2020. (Cited on pages 4 and 5)
- Alon, Titan M, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt**, “The impact of COVID-19 on gender equality,” *Covid Economics: Vetted and Real-Time Papers*, April 2020, 4, 62–85. (Cited on page 31)
- Alstadsæter, Annette, Bernt Bratsberg, Gaute Eielsen, Wojciech Kopczuk, Simen Markussen, Oddbjorn Raaum, and Knut Røed**, “The first weeks of the coronavirus crisis: Who got hit, when and why? Evidence from Norway,” National Bureau of Economic Research 2020. (Cited on page 26)
- Alvarez, Fernando E, David Argente, and Francesco Lippi**, “A simple planning problem for covid-19 lockdown,” National Bureau of Economic Research 2020. (Cited on pages 26 and 30)
- Antràs, Pol, Stephen J Redding, and Esteban Rossi-Hansberg**, “Globalization and Pandemics,” Technical Report, Harvard University Working Paper 2020. (Cited on page 4)
- Argente, David O, Chang-Tai Hsieh, and Munseob Lee**, “The Cost of Privacy: Welfare Effect of the Disclosure of COVID-19 Cases,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Atkeson, Andrew**, “What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios,” National Bureau of Economic Research 2020. (Cited on pages 4 and 26)
- Aum, Sangmin, Sang Yoon Tim Lee, and Yongseok Shin**, “COVID-19 Doesn’t Need Lockdowns to Destroy Jobs: The Effect of Local Outbreaks in Korea,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Azzimonti, Marina, Alessandra Fogli, Fabrizio Perri, and Mark Ponder**, “Pandemic Control in ECON-EPI Networks,” Federal Reserve Bank of Minneapolis Staff report 609 August 2020. (Cited on page 4)

- Babajanyan, S. G. and Kang Hao Cheong**, “Age-structured SIR model and resource growth dynamics: A preliminary COVID-19 study,” *medRxiv*, 2020. (Cited on page 4)
- Baqae, David, Emmanuel Farhi, Michael J Mina, and James H Stock**, “Reopening Scenarios,” Working Paper 27244, National Bureau of Economic Research May 2020. (Cited on page 4)
- , – , **Michael Mina, and James H Stock**, “Policies for a second wave,” *Brookings Papers on Economic Activity*, 2020. (Cited on page 4)
- Barrios, John M, Efraim Benmelech, Yael V Hochberg, Paola Sapienza, and Luigi Zingales**, “Civic Capital and Social Distancing during the Covid-19 Pandemic,” Working Paper 27320, National Bureau of Economic Research June 2020. (Cited on page 5)
- Bartik, Alexander W, Marianne Bertrand, Zoë B Cullen, Edward L Glaeser, Michael Luca, and Christopher T Stanton**, “How are small businesses adjusting to covid-19? early evidence from a survey,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Bethune, Zachary A and Anton Korinek**, “Covid-19 infection externalities: Trading off lives vs. livelihoods,” *National Bureau of Economic Research Working Paper*, 2020. (Cited on page 5)
- Bisin, Alberto and Andrea Moro**, “Notes on Rational Forward Looking SIR,” mimeo, NYU 2020. (Cited on page 7)
- and – , “JUE Insight: Learning Epidemiology by Doing: The Empirical Implications of a Spatial-SIR Model with Behavioral Responses,” *Journal of Urban Economics*, 2021. (Cited on pages 2, 4, 9, 17, 28, and 30)
- Bognanni, Mark, Doug Hanley, Daniel Kolliner, and Kurt Mitman**, “Economic Activity and COVID-19 Transmission: Evidence from an Estimated Economic-Epidemiological Model,” mimeo, IIES 2020. (Cited on page 4)
- Britton, Tom, Frank Ball, and Pieter Trapman**, “A mathematical model reveals the influence of population heterogeneity on herd immunity to SARS-CoV-2,” *Science*, 2020. (Cited on page 9)
- Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt**, “An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies,” *CEPR Discussion Paper No. DP14695*, 2020. (Cited on page 4)
- , **Tiago Cavalcanti, Daniel Da Mata, and Cezar Santos**, “Slums and pandemics,” CEPR Discussion Paper No. DP15131 2020. (Cited on page 4)

- Chetty, Raj, J Friedman, Nathaniel Hendren, and Michael Stepner**, “The economic impacts of COVID-19: Evidence from a new public database built from private sector data,” Opportunity Insights Working Paper. 2020. (Cited on page 26)
- Cicala, Steve, Stephen P Holland, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates**, “Expected Health Effects of Reduced Air Pollution from COVID-19 Social Distancing,” Working Paper 27135, National Bureau of Economic Research May 2020. (Cited on page 5)
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber**, “The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Crossley, Thomas F, Paul Fisher, and Hamish Low**, “The Heterogeneous and Regressive Consequences of COVID-19: Evidence from High Quality Panel Data,” *Available at SSRN 3703018*, 2020. (Cited on page 26)
- Cuñat, Alejandro and Robert Zymek**, “The (Structural) Gravity of Epidemics,” CESifo Working Paper 2020. (Cited on page 4)
- Davenport, Alex, Robert Joyce, Imran Rasul, and Tom Waters**, “Spending and saving during the COVID-19 crisis: evidence from bank account data,” The Institute for Fiscal Studies briefind note BN308 2020. (Cited on page 26)
- Dingel, Jonathan I. and Brent Neiman**, “How Many Jobs Can Be Done at Home?,” Mimeo, U. of Chicago 2020. (Cited on page 31)
- Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt**, “The macroeconomics of epidemics,” National Bureau of Economic Research 2020. (Cited on page 4)
- Ellison, Glenn**, “Implications of Heterogeneous SIR Models for Analyses of COVID-19,” National Bureau of Economic Research Working Paper 27373 June 2020. (Cited on page 4)
- Engle, Samuel, Jussi Keppo, Marianna Kudlyak, E Querciolo, Lones Smith, and Andrea Wilson**, “The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics,” Working paper 2021. (Cited on pages 4, 5, 6, 8, and 31)
- Eubank, Stephen, Hasan Guclu, VS Anil Kumar, Madhav V Marathe, Aravind Srinivasan, Zoltan Toroczkai, and Nan Wang**, “Modelling disease outbreaks in realistic urban social networks,” *Nature*, 2004, 429 (6988), 180–184. (Cited on page 4)
- Farboodi, Maryam, Gregor Jarosch, and Robert Shimer**, “Internal and External Effects of Social Distancing in a Pandemic,” Working Paper 27059, National Bureau of Economic Research April 2020. (Cited on page 5)

- Fenichel, Eli P**, “Economic considerations for social distancing and behavioral based policies during an epidemic,” *Journal of health economics*, 2013, 32 (2), 440–451. (Cited on page 4)
- Ferguson, Neil, Daniel Laydon, Gemma Nedjati Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, ZULMA Cucunuba Perez, Gina Cuomo-Dannenburg et al.**, “Imperial College COVID-19 Response Team: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand,” Imperial College London 2020. (Cited on pages 8, 29, 30, and 31)
- Fernandez-Villaverde, Jesus and Charles Jones**, “Estimating and Simulating a SIRD Model of COVID-19,” mimeo 2020. (Cited on pages 4 and 5)
- Fischer, Carolyn**, “Six degrees of separation for Covid-19: The external costs and benefits of extended distancing for different social groups,” *Covid Economics: Vetted and Real-Time Papers*, June 2020, 25, 1–22. (Cited on page 4)
- Funk, Sebastian, Marcel Salathé, and Vincent AA Jansen**, “Modelling the influence of human behaviour on the spread of infectious diseases: a review,” *Journal of the Royal Society Interface*, 2010, 7 (50), 1247–1256. (Cited on page 4)
- Gaeta, Giuseppe**, “A simple SIR model with a large set of asymptomatic infectives,” *arXiv preprint arXiv:2003.08720*, 2020. (Cited on page 5)
- Gans, Joshua S**, “The Economic Consequences of $\hat{R} = 1$: Towards a Workable Behavioural Epidemiological Model of Pandemics,” Working Paper 27632, National Bureau of Economic Research July 2020. (Cited on page 7)
- Geoffard, Pierre-Yves and Tomas Philipson**, “Rational epidemics and their public control,” *International economic review*, 1996, pp. 603–624. (Cited on page 4)
- Giagheddu, Marta and Andrea Papetti**, “The Macroeconomics of Age-Varying Epidemics,” *Available at SSRN 3651251*, 2020. (Cited on page 4)
- Giannitsarou, Chryssi, Stephen Kissler, and Flavio Toxvaerd**, “Waning Immunity and the Second Wave: Some Projections for SARS-CoV-2,” CEPR Discussion Paper No. DP14852 2020. (Cited on page 4)
- Glaeser, Edward L, Caitlin S Gorbach, and Stephen J Redding**, “How Much does Covid-19 Increase with Mobility? Evidence from New York and Four Other U.S. Cities,” mimeo 2020. (Cited on page 4)
- Goenka, Aditya and Lin Liu**, “Infectious diseases and endogenous fluctuations,” *Economic Theory*, 2012, 50 (1), 125–149. (Cited on page 5)

- Gollier, Christian**, “Cost-benefit analysis of age-specific deconfinement strategies,” *Covid Economics: Vetted and Real-Time Papers*, June 2020, *24*, 1–22. (Cited on page 4)
- Gomes, M Gabriela M, Ricardo Aguas, Rodrigo M Corder, Jessica G King, Kate E Langwig, Caetano Souto-Maior, Jorge Carneiro, Marcelo U Ferreira, and Carlos Penha-Goncalves**, “Individual variation in susceptibility or exposure to SARS-CoV-2 lowers the herd immunity threshold,” *medRxiv*, 2020. (Cited on page 9)
- Goolsbee, Austan and Chad Syverson**, “Fear, lockdown, and diversion: comparing drivers of pandemic economic decline 2020,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michèle Tertilt**, “An equilibrium model of the African HIV/AIDS epidemic,” *Econometrica*, 2019, *87*(4), 1081–1113. (Cited on page 5)
- Gupta, Sumedha, Laura Montenovo, Thuy D Nguyen, Felipe Lozano Rojas, Ian M Schmutte, Kosali I Simon, Bruce A Weinberg, and Coady Wing**, “Effects of social distancing policy on labor market outcomes,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Jarosch, Gregor, Maryam Farboodi, and Robert Shimer**, “Internal and External Effects of Social Distancing in a Pandemic,” 2020. (Cited on page 26)
- Kahn, Lisa B, Fabian Lange, and David G Wiczer**, “Labor Demand in the Time of COVID-19: Evidence from Vacancy Postings and UI Claims,” Working Paper 27061, National Bureau of Economic Research April 2020. (Cited on page 5)
- Kaplan, Greg, Ben Moll, and Gianluca Violante**, “Pandemics According to HANK,” *University of Chicago*, 2020. (Cited on pages 26 and 30)
- Leung, Ka Yin, Pieter Trapman, and Tom Britton**, “Who is the infector? Epidemic models with symptomatic and asymptomatic cases,” *Mathematical biosciences*, 2018, *301*, 190–198. (Cited on page 5)
- Lucas, Robert E**, “Econometric policy evaluation: A critique,” in “Carnegie-Rochester conference series on public policy,” Vol. 1 1976, pp. 19–46. (Cited on pages 4 and 21)
- Mossong, Joël, Niel Hens, Mark Jit, Philippe Beutels, Kari Auranen, Rafael Mikolajczyk, Marco Massari, Stefania Salmaso, Gianpaolo Scalia Tomba, Jacco Wallinga, and et al.**, “Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases,” *PLoS Medicine*, Mar 2008, *5*(3), e74. (Cited on pages 8 and 28)

- Paula, Áureo De, Martin O’Connell, Kate Smith et al.**, “Preparing for a pandemic: Spending dynamics and panic buying during the COVID-19 first wave,” CEPR Discussion Papers n. 15371 2020. (Cited on page 26)
- Rojas, Felipe Lozano, Xuan Jiang, Laura Montenovo, Kosali I Simon, Bruce A Weinberg, and Coady Wing**, “Is the cure worse than the problem itself? immediate labor market effects of covid-19 case rates and school closures in the us,” Technical Report, National Bureau of Economic Research 2020. (Cited on page 5)
- Rowthorn, Bob RE and Flavio Toxvaerd**, “The optimal control of infectious diseases via prevention and treatment,” CEPR Discussion Paper No. DP8925 2012. (Cited on page 4)
- Toxvaerd, Flavio**, “Equilibrium social distancing,” Faculty of Economics, University of Cambridge 2020. (Cited on page 5)
- Verelst, Frederik, Lander Willem, and Philippe Beutels**, “Behavioural change models for infectious disease transmission: a systematic review (2010–2015),” *Journal of The Royal Society Interface*, 2016, 13 (125), 20160820. (Cited on page 4)
- Weitz, Joshua S, Sang Woo Park, Ceyhun Eksin, and Jonathan Dushoff**, “Moving Beyond a Peak Mentality: Plateaus, Shoulders, Oscillations and Other’Anomalous’ Behavior-Driven Shapes in COVID-19 Outbreaks,” *medRxiv*, 2020. (Cited on page 4)