Annals of Economics and Statistics, Number 142, June 2021

DIFFUSION OF COVID-19 IN SOCIAL AND PRODUCTION NETWORKS: SIMULATION EVIDENCE FROM A DYNAMIC MODEL

VICTOR AGUIRREGABIRIA[a], JIAYING GU[b], YAO LUO[c] AND PEDRO MIRA[d]

This paper presents a dynamic model to evaluate economic and public health effects of the diffusion of COVID-19. Our framework combines a SIR epidemiological model of virus diffusion with a dynamic game of network production and social interactions. The economy comprises three types of geographic locations: *homes*, *workplaces*, and *consumption places*. Each individual has her own set of locations where she develops her life. The combination of these sets for all the individuals determines the economy's production and social network. Every day, individuals choose to work and consume either outside (with physical interaction with other people) or remotely (from home, without physical interactions). Working (and consuming) outside is more productive and generates stronger complementarities (positive externality). However, in the presence of a virus, working outside facilitates infection and the diffusion of the virus (negative externality). Individuals are forward-looking. We calibrate the model and implement numerical experiments to evaluate the health and economic impact of several counterfactual public policies: subsidies for working at home; testing policies; herd immunity; and changes in the network structure. These policies generate substantial differences in the propagation of the virus and its economic impact.

JEL Codes: C57, C73, L14, L23, I18.
Keywords: Covid-19, Virus Diffusion, Dynamics, Production and Social Networks, Production Externalities, Public Health.

1. INTRODUCTION

The COVID-19 pandemic has generated important challenges and uncertainties in our societies. The academic economics profession has responded with notorious engagement to these challenges. We, as many other academics, believe that COVID-19 is a unique opportunity to make progress on some research issues which are important for the economics of pandemics. There is much uncertainty about the evolution of the pandemic, its economic effects, and the effectiveness of different public policies. There are also uncertainties in the academic profession – more than usual – about which models can be more useful for policy analysis in this new scenario. This paper is one of the many attempts by economists to provide a structural model of dynamic decision-making that could potentially be useful for policy analysis.

In this paper, we present a framework that combines an epidemiological model of COVID-19 diffusion with a dynamic game of network production and social interactions. The model emphasizes several aspects typically absent from epidemiological models that

We would like to thank comments from the Editor, two anonymous referees, Connor Campbell, Marc-Antoine Chatelain, Luca Flabbi, Donna Gilleskie, Andrés Hincapié, Matt Mitchell, Adonis Yatchew, and from participants in seminars at University of North Carolina - Chapel Hill and University of Toronto.

[a]University of Toronto and CEPR. victor.aguirregabiria@utoronto.ca
[b]University of Toronto. jiaying.gu@utoronto.ca
[c]University of Toronto. yao.luo@utoronto.ca
[d]CEMFI. mira@cemfi.es

are work horses in economic models: (i) individuals make choices to maximize their own welfare and respond to incentives affecting this welfare; (ii) individuals interact with each other and their choices and contributions to economic and social outcomes depend on the behavior of coworkers, suppliers, clients, family members, and friends. The production (social) system is a network. Importantly, production and social networks also determine physical links between individuals that can facilitate infection. Network links can generate positive spillover effects in production and consumption, but also negative externalities because of infections. This tradeoff between the positive and negative externalities from network links plays a fundamental role in the diffusion of a virus and its economic impact. In this paper, we develop a dynamic model that emphasizes the relationship between the production/social network in an economy and the diffusion of COVID-19 and its economic impact. The model can be used to evaluate the impact that different public policies have on the propagation of a virus and its economic effects.[1]

Our model incorporates the following features.

Production and social network. The economy comprises a set of geographic locations and a set of individuals. We distinguish three types of locations: *homes*, *workplaces*, and *consumption places*. Each individual has her own set of locations where she develops her life: her *home(s)*, *workplaces*, and *consumption places*. The combination of these sets for all the individuals determines the economy's production and social network. The structure of the network can be determined using data on individuals' mobility in the absence of COVID-19.

Endogenous individuals' choices. Every day, individuals choose to work and consume either outside (with physical interaction with other people) or remotely (from home, without physical interactions). Working (consuming) *outside* is more productive and generates complementarities. Therefore, in the absence of a virus, working outside generates a positive externality.

Epidemiological model. In the presence of a virus transmitted through physical contact, working or consuming *outside* facilitates the diffusion of the virus (negative externality). The epidemiological part of our model incorporates substantial extensions with respect to standard SIR models. First, the production/social network is an important component of our epidemiological model. Second, the probability of infection is endogenous as it depends on working and consumption decisions of the own individual and of other people in her production and social network.[2] This implies that the probability of infection varies across regions in the network. The model provides a landscape of this probability over a city, and this landscape evolves endogenously.

Information structure and testing. A special feature of COVID-19 is that asymptomatic individuals, some of which may never develop symptoms, are infectious. In the absence of testing, an individual without symptoms does not know whether she is healthy (noninfectious), or infected asymptomatic, or even already recovered without having developed

[1]The negative externality from production and consumption network links is less important for other viruses which are not transmitted as easily as Covid-19, such as HIV. For this other type of viruses, a production/social network model may be less helpful for predictions and policy analysis.

[2]Other types of individuals' choices – which are important for the spread of the virus and have economic implications – are hygiene and keeping a distance from others in personal interactions. So far, we have focused on working and consumption decisions because they have strong economic implications, and mobility data that measures them are readily available.

symptoms. This incomplete information facilitates the diffusion of the virus and is an important element in our model. In this context, testing asymptomatic individuals can reduce this uncertainty, both for the tested individual and economy-wide. The degree of testing is a policy variable chosen by the government.

We characterize an equilibrium of the dynamic game and present an algorithm for its computation. We calibrate the model and use it to evaluate the health and economic impact of factual and counterfactual public policies: subsidies for working at home; more testing; herd immunity; and changes in the structure of the production/social network.

This paper tries to contribute to a rapidly growing economic literature on the diffusion of COVID-19 and its economic impact. Our paper is closely related to the economic literature on *rational epidemics* that extends the SIR epidemiological model (Kermack and McKendrick, 1927) to take into account how individuals react to changes in prevalence: see Kremer (1996), Geoffard and Philipson (1996), Auld (2003), Chan, Hamilton, and Papageorge (2016), or Greenwood, Kircher, Santos, and Tertilt (2019), among others. Most of this literature has focused on the HIV epidemics. Recent papers apply this approach to study COVID-19. Alvarez, Argente, and Lippi (2020), Eichenbaum, Rebelo, and Trabandt (2020), and Hall, Jones, and Klenow (2020) combine an epidemiological model with a macro equilibrium model where individuals make consumption and labor supply decisions and are (intertemporal) utility maximizers. Jones et al. (2020) study the tradeoffs faced by a social planner who tries to mitigate the spread of COVID-19. They show that the social planner's solution implies a much more drastic reduction in consumption and output than in a decentralized equilibrium. In a similar vein, Acemoglu, Chernozhukov, Werning, and Whinston (2020) study Pareto optimal lockdown policies for COVID-19 –where the key tradeoff is between deaths and economic loss. They show that the Pareto frontier can be substantially improved if lockdown policies apply differently across age groups. In Farboodi, Jarosch, and Shimer (2020), forward-looking individuals choose a degree of social interaction. Similar to our paper but distinctive to many others, newly-infected individuals are unaware of their infection status and may spread the disease unknowingly. After calibrating their model, they quantify the expected cost of COVID-19 and the benefits of policy interventions.[3]

The standard SIR model and the macro equilibrium models mentioned in the previous paragraph assume that everyone has an equal chance of meeting the infectious population (i.e., random mixing). In practice, people interact more with those in their networks of family, work, or consumption, and this generates heterogeneity in the probability of infection. Our model captures this type of heterogeneity. Some economic papers on Covid-19 have also introduced a social network to understand the diffusion of the disease. Karaivanov (2020) shows that testing and contact tracing are more effective in the network model. However, his paper does not endogenize individual decisions. Azzimonti, Fogli, Perri, and Ponder (2020) build a city network model where social networks (family, work, school, and commute) determine production outputs and diffusion of the disease. Their model demonstrates how the effects of a pandemic highly depend on the severed links in social networks. Moreover, specifying the production sector is essential for assessing trade-offs of different policies and their complementarities. However, in their model, individuals follow their routines unless they get sick. In contrast, our model endogenizes

[3]Berger, Herkenhoff, and Mongey (2020) and Piguillem and Shi (2020) study the value of information from testing.

individuals' work and consumption choices.

Our paper is also related to the literature of learning in social network games. Computational tractability is a fundamental issue in this literature. Models with fully rational players with perfect Bayesian updating beliefs are intractable except in very stylized cases.[4] Authors have proposed different forms of adaptive or naive learning from neighbors (Bala and Goyal, 1998; Golub and Jackson, 2010). We follow this approach. More specifically, our assumptions on agents' information structure and beliefs updating are in the spirit of Acemoglu, Dahleh, Lobel, and Ozdaglar (2011); Acemoglu, Bimpikis, and Ozdaglar (2014) and Mossel, Mueller-Frank, Sly, and Tamuz (2020). In our model, agents combine local and economy-wide information and use an adaptive rule to update their beliefs about health and probability of infection. In contrast, to most models in this literature, agents in our model are forward-looking. However, for tractability, we need to impose restrictions on their beliefs about other agents in the game.

In contrast to the standard SIR model, the infection rate in our model is endogenous and heterogeneous across workplaces. Pichler (2015) proposes a model of endogenous sickness absences to study their procyclical behavior. In his model, the probability that a sick individual goes to work is higher in a boom than in a bust, and this implies a broader spread of a virus during periods of economic expansion. In our model, the risk of infection depends on the number of infected coworkers who decide to work in the workplace and not at home. This is an endogenous decision. There is complementarity in the production function between coworkers' choices of working in-site or at home. This implies that an increase in the risk of infection of an individual has a social multiplier effect on coworkers' decision of working at home.

In our model, the local structure of the production/social network plays a key role in the diffusion of the virus in a local community and across communities. Measures of social connectivity and mobility are important. Kuchler, Russel, and Stroebel (2020) use data from Facebook to measure the degree of social connectivity in Italy and in US. They present evidence on the relationship between an index of social connectivity and the density of COVID-19 cases. In an influential paper, Adda (2016) uses detailed weekly data on disease incidence in France covering a period of 25 years to measure how exogenous changes in social distancing – public transportation strikes, opening of new railway lines, school closure due to holidays – affect the probability of infection.

A motivation for our paper is to provide a framework to evaluate the economic impact of factual and counterfactual public policies to mitigate the spread of COVID-19. Recent papers present evidence for Japan, Italy, and France, respectively. Inoue and Todo (2020) use a large dataset with information from more than 1.6 million firms and almost 6 million supply-chain links in Tokyo to quantify the economic impact of a hypothetical lockdown policy in this city. Their estimates and experiments show a huge production loss of 309 billion yen per day. This effect would quickly spread to the whole Japanese economy such that in one month total output would be reduced by 86%. Boeri, Caiumi, and Paccagnella (2020) study the impact of COVID-19 on employment and on the type of jobs in Italy. Barrot, Grassi, and Sauvagnat (2020) propose measures on the degree of remote working for different industries in France and use these measures to estimate how social-distancing

[4]See the discussion on this issue in the recent paper by Mossel, Mueller-Frank, Sly, and Tamuz (2020) (2020; pages 1235-1236), and their citation to Gale and Karaivanov (2020): *"The computational difficulty of solving the model is massive even in the case of three persons"*.

policies have affected production. They conclude that a six weeks confinement reduces GDP by approximately 5.6%. The upstream sectors are the most negatively affected. The analysis emphasizes the importance of industrial composition for the aggregate economic impact.

The rest of this paper is organized as follows. Section 2 presents the model. Section 3 presents a calibration of the model and our policy experiments. We summarize and conclude in section 4.

2. MODEL

2.1. *The network*

The economy consists of a set of L geographic locations, $\mathcal{L} = \{1, 2, ..., L\}$, and a set of N individuals, $\mathcal{I} = \{1, 2, ..., N\}$. We index locations by ℓ and individuals by i. Time is discrete and indexed by $t \in \{0, 1, ...\}$. One period is one day. There are three types of locations: *homes*, *workplaces*, and *consumption places*. Each individual has her own set of locations where she develops her life: her home(s), $\mathcal{L}_i^H$, workplaces, $\mathcal{L}_i^W$, and consumption places, $\mathcal{L}_i^C$. Each of these individual-specific sets may contain one or multiple locations.

An individual's *household* consists of all the other individuals who share the same home: that is, the set $\mathcal{H}_i \equiv \{j : \mathcal{L}_i^H \cap \mathcal{L}_j^H \neq \varnothing\}$. Similarly, an individual's *production (consumption) team* consists of all the other people who share a workplace (consumption place) with her: that is, the set $\mathcal{W}_i \equiv \{j : \mathcal{L}_i^W \cap \mathcal{L}_j^W \neq \varnothing\}$ for production, and the set $\mathcal{C}_i \equiv \{j : \mathcal{L}_i^C \cap \mathcal{L}_j^C \neq \varnothing\}$ for consumption.

The combination of all these sets, $\{\mathcal{L}_i^H, \mathcal{L}_i^W, \mathcal{L}_i^C : i \in \mathcal{I}\}$ or equivalently $\{\mathcal{H}_i, \mathcal{W}_i, \mathcal{C}_i : i \in \mathcal{I}\}$, describes the *network* in this economy. This network is an exogenous primitive in the model.[5] The network can vary across economies because industrial composition, geography, transportation infrastructures, culture, etc. In this model, the network not only describes social and economic interlinks but also physical contacts. We assume that an individual with the virus can infect other individual only if they share a common location, either home, or workplace, or consumption place.

A network can be represented by a graph consisting of nodes and edges. Nodes correspond to the set of individuals while edges represent who they are connected to, either through homes, workplaces or consumption places. The *degree of a node*, d_i, is defined as the number of neighbors it has. In our model, d_i is the cardinality (number of elements) of the set $\{\mathcal{H}_i, \mathcal{W}_i, \mathcal{C}_i\}$. Three properties are often used to describe a graph: (i) the distribution of degrees d_i, which measures the heterogeneity of individual's connectivity; (ii) the *clustering coefficient*, which measures how often a triple of nodes forms a triangle; and (iii) the *average path length*, where a path length is defined as the smallest number of edges one needs to travel to connect two nodes. This statistic measures how connected the network is. A *regular graph* is a graph where each node has the same degree: the distribution of degrees is degenerate. Social networks in reality have a non-uniform distribution

[5] Although the network is an exogenous primitive, in our model an individual makes daily choices about which of her production and consumption places to visit on a given day. The model's solution determines a probability distribution over the locations she visits and the persons she meets. In this sense, the network is endogenous and it changes as the epidemic spreads. However, we assume that an individual's sets of production and consumption places remain fixed. We believe this is a plausible assumption because Covid-19 is not good time to expand your group of physical friends, consumption places, or workplaces.

of degrees, with some individuals having more connections than others.

2.2. *Health and diagnosis states and transitions*

Variable $x_{it} \in \mathcal{X}$ describes the health and diagnosis state of individual i at period t. It can take ten possible values: $\mathcal{X} \equiv \{H, AU, AD, SU, SD, RAU, RSU, RAD, RSD, Death\}$.

State H (for ***Healthy***) means that the individual has not been infected with the virus. States AU, AD, SU, and SD represent infected individuals at different states depending on the development of symptoms and on the existence of diagnosis.[6] State AU (for ***infected Asymptomatic Undiagnosed***) represents an individual who is infected but has not developed symptoms yet and has not been diagnosed. State AD means that the individual is infected and asymptomatic (A) but she has been diagnosed (D). State SU represents an infected individual who has developed symptoms (S) but has not been diagnosed (U). State SD represents an infected individual who has developed symptoms (S) and has been diagnosed (D).

States RAU, RSU, RAD, and RSD represent recovered individuals who are at different states depending on whether they developed symptoms and whether they were diagnosed.[7] States RAU and RSU represent recovered individuals who had never been diagnosed – asymptomatic and symptomatic, respectively – such that the individual does not know that she has been infected. In contrast, states RAD and RSD represent recovered individuals who had been previously diagnosed. Finally, $Death$ means death because of the virus.

[6]Throughout the paper, we use the term infected as synonymous of infectious. In reality, this is not exactly the case. A virus needs to replicate itself sufficiently in a person's body before this person becomes infectious. Our model can be trivially extended to include an additional state between states H and AU such that the model would distinguish between infected and infectious. This additional state – say E form "Exposed" – would represent an individual who had the virus but has not become infectious yet.

[7]We assume that "Recovered" implies not infectious and immune.

Figure 1: Flow Diagram of States and Transitions in Simplified Model

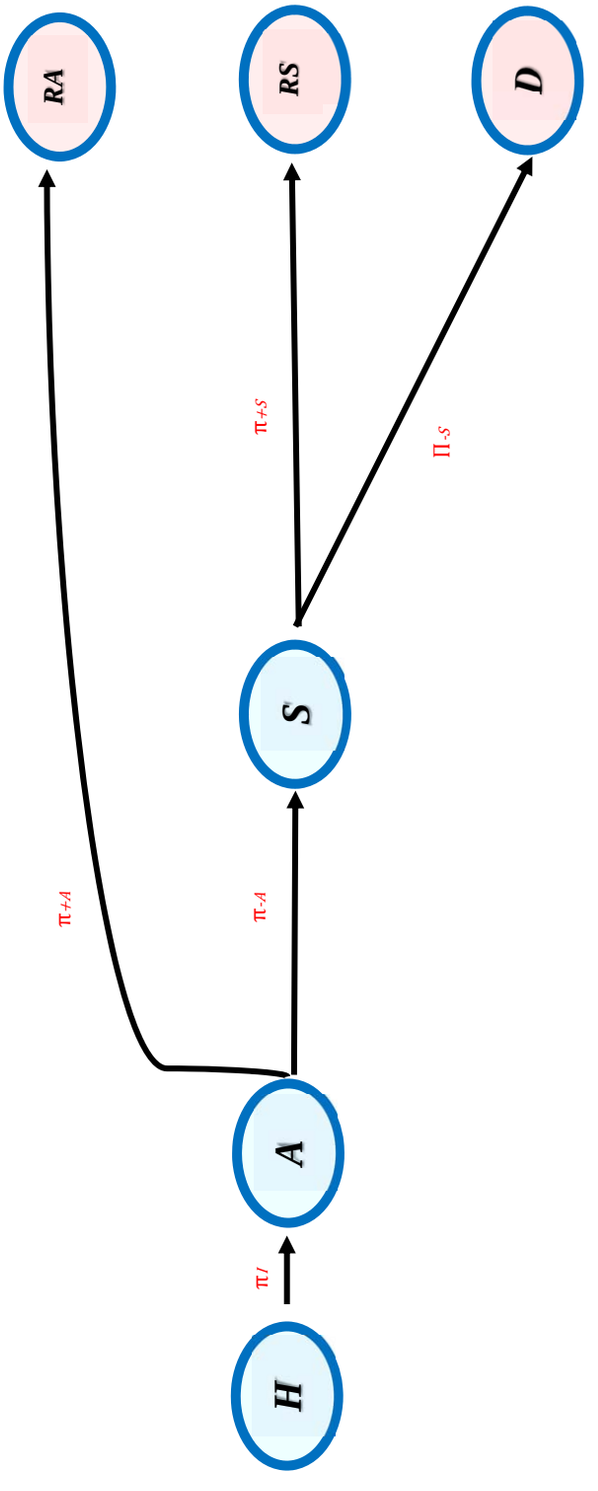

States: **H** *(Healthy)*; **A** *(Infected Asymptomatic)*; **S** *(Infected Symptomatic)*; **RA** *(Recovered from Asymptomatic)*; **RS** *(Recovered from Symptomatic)*; **D** *(Death)*.

Probabilties: π_I *=Infection;* π_{+A} *=Positive health shock (Recovery) for asymptomatic.;* π_{-A} *=Negative health shock (symptoms) for asymptomatic.;* π_{+S} *=Positive health shock (Recovery) for symptomatic.;* π_{-S} *=Negative health shock (Death) for symptomatic.*

Figure 2: Flow Diagram of States and Transitions

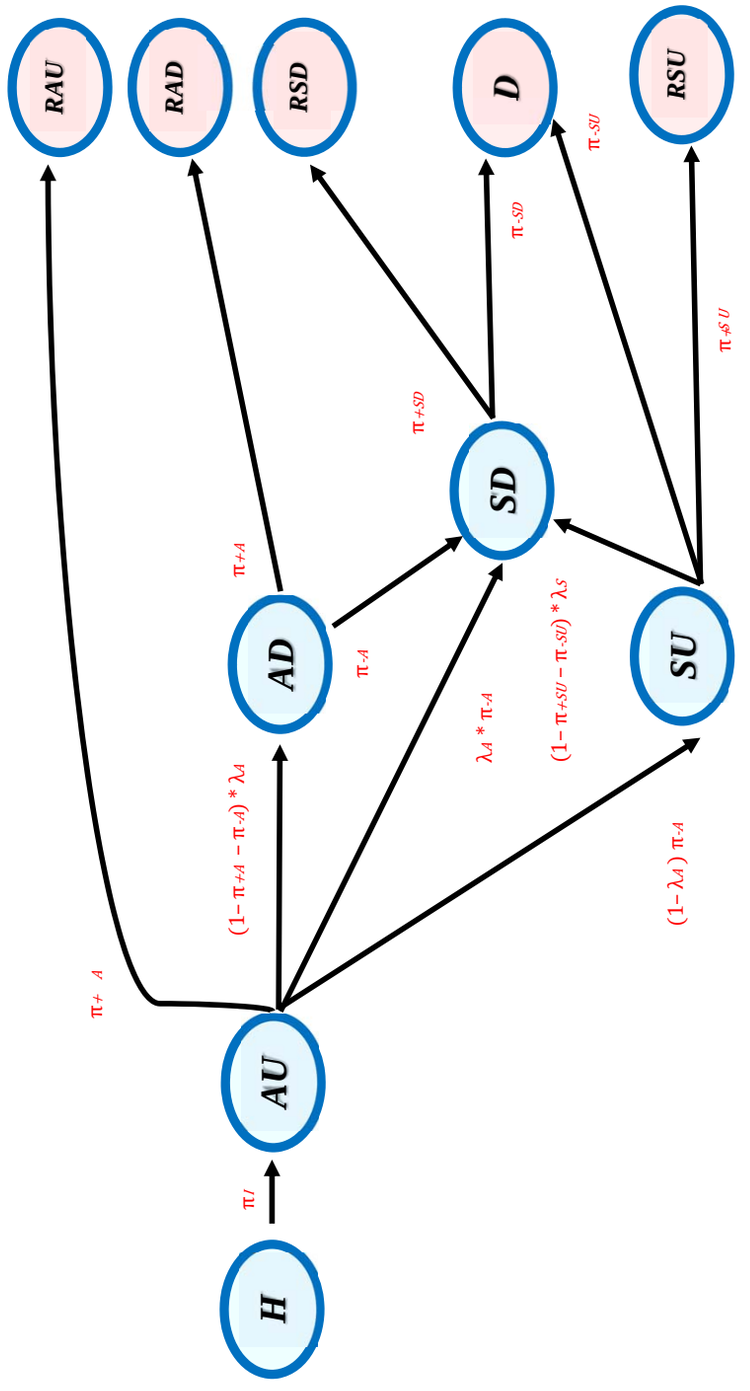

States: ***H*** *(Healthy);* ***AU*** *(Infected Asymptomatic Undiagnosed);* ***AD*** *(Infected Asymptomatic Diagnosed);* ***SU*** *(Infected Symptomatic Undiagnosed);* ***SD*** *(Infected Symptomatic Diagnosed);* ***RAU*** *(Recovered from AU);* ***RSU*** *(Recovered from SU);* ***RAD*** *(Recovered from AD);* ***RSD*** *(Recovered from SD);* ***D*** *(Death).*

Probabilties: π_I *=Infection;* λ_A *=Test to asymptomatic;* λ_S *=Test to symptomatic;* π_{+A} *=Positive health shock (Recovery) for asymp.;* π_{-A} *=Negative shock (symptoms) for asymp.;* π_{+SU} *=Positive shock (Recovery) for symp-undiag.;* π_{-SU} *=Negative shock (Death) for symp-undiag.;* π_{+SD} *=Positive shock (Recovery) for symp-diag.* π_{-SD} *=Negative shock (Death) for symp-diag.*

Transitions between states are based on two types of shocks: health shocks, and testing shocks. We consider that health shocks can take three values: positive, negative, and neutral. Moreover, we assume that tests are accurate such that they cannot provide neither false positives nor false negatives. Figure 1 presents a flow diagram of a simplified version of our model with only four health states. Figure 2 presents the flow diagram of our model.

(i) Infection [Transition $H \rightarrow AU$]. Every day t, a healthy individual (H) can become infected with a probability π_{it}^{I}. This probability is endogenous. It depends on the individual's behavior (confinement or not) and on the behavior of other individuals in her social group. We describe the form of this endogenous probability in section 2.5 below. Individuals in state H can be randomly selected to be tested. If the test result is negative, the individual remains in state H.

(ii) Transitions from AU. Every period, an individual in state AU (infected asymptomatic undiagnosed) receives a health shock and a testing shock. The health shock can take three possible values: positive (with probability π_{+A}), negative (with probability π_{-A}), or neutral. If the shock is positive, the individual recovers and becomes immune. If the shock is negative, the individual develops symptoms. We assume that individuals cannot transition within one day from asymptomatic to death: they need to develop symptoms before dying.

The testing shock is independent of the health shock and it determines whether the individual is tested for the virus (with probability λ_A) or not (with probability $1 - \lambda_A$). We also assume that an individual in AU cannot be tested positive on the same day that she receives a positive health shock and recovers.[8]

Under these conditions, there are five possible transitions from state AU. (1) Neutral health shock and no testing – with probability $(1 - \pi_{+A} - \pi_{-A})\ (1 - \lambda_A)$ – implies that the individual remains in the same state AU. (2) Neutral health shock and testing – with probability $(1 - \pi_{+A} - \pi_{-A})\ \lambda_A$ – implies that she remains asymptomatic but now is diagnosed: she moves to state AD. (3) Regardless testing, a positive health shock – with probability π_{+A} – means that she recovers and has not been diagnosed such that she does not know that she was infected. This corresponds to state RAU. (4) Negative health shock and no testing – with probability $\pi_{-A}\ (1 - \lambda_A)$ – implies that the individual develops symptoms but is still undiagnosed: she moves to state SU. (5) Finally, with a negative health shock and testing – with probability $\pi_{-A}\ \lambda_A$ – the individual moves to state SD where she is both symptomatic and diagnosed.

(iii) Transitions from SU. The transitions from state SU are similar to those from AU, but a main difference is that a negative health shock implies death. Furthermore, the probabilities of a positive and a negative health shock for a symptomatic individual – π_{+SU} and π_{-SU}, respectively – are different that for an asymptomatic. The government policy on testing can be different for symptomatic and asymptomatic, such that probability λ_S is

[8] We are assuming that all the testing is PCR testing. A PCR test detects the infection while it is active. For a PCR test, individuals in states A and S should test positive, and individuals in states H and R negative. For this test, the rates of false positives or false negatives are very low. An interesting extension of the model would be to include other tests that are emerging, including testing for antibodies and immunity.

different to λ_A. Similarly as for state AU, we assume that the individual cannot recover and test positive on the same day.

Health and testing shocks determine four possible transitions from state SU. (1) Under neutral health shock and no testing, the individual remains in state SU – with probability $(1 - \pi_{+SU} - \pi_{-SU})\ (1 - \lambda_S)$. (2) With neutral health shock and testing, she becomes diagnosed and moves to state SD – with probability $(1 - \pi_{+SU} - \pi_{-SU})\ \lambda_S$. (3) A positive health shock – regardless testing – means that she recovers and remains undiagnosed: she moves to state RSU with probability π_{+SU}. (4) Finally, regardless testing, with a negative health shock – with probability π_{-SU} – the individual dies.

(iv) Transitions from AD and from SD. For diagnosed individuals, testing does not matter and only health shocks determine the transitions from these states.

For an individual in state AD, the probability distribution of the health shocks is the same as under state AU. That is, we assume that diagnosis does not affect the health transition when the individual is asymptomatic. There are three possible transitions. Under a positive health shock, the individual recovers and arrives to state RAD – with probability π_{+A}. Under a negative health shock, she develops symptoms and moves to state SD – with probability π_{-A}. And with a neutral shock, she stays in state AD – with probability $1 - \pi_{+A} - \pi_{-A}$.

Being diagnosed can affect the distribution of health shocks if an individual is symptomatic: that is, SD individuals are more likely to receive some treatment than SU individuals. Therefore, the probabilities π_{-SD} and π_{+SD} can be different than the probabilities π_{-SU} and π_{+SU}.

There are three possible transitions under state SD. Given a positive health shock, the individual recovers and arrives to state RSD – with probability π_{+SD}. Under a negative health shock – with probability π_{-SD} – she dies. Finally, with a neutral shock she stays in state SD with probability $1 - \pi_{+SD} - \pi_{-SD}$.

Finally, we assume that all the recovered states – RAU, RAD, RSU, RSD – are absorbing states. Individuals in states RAU and RSU can be subject to random testing, but the test will be negative and individuals remain in the same state.

There are two relevant extensions of the model regarding the probability λ_A. First, it is interesting to allow for false negatives in the results of the test. In this extension, the probability λ_A could be interpreted as the product of two probabilities: the probability of being selected for testing times the probability of a positive result of test conditional on infection, i.e., a *true positive*. In that model, parameter λ_A measures the government testing effort in two different dimensions: the number of tests and the quality of the testing procedure. However, this extension of the model requires also some non-trivial changes in individuals' beliefs about their actual health status. A second relevant extension is to allow the probability that an asymptomatic individual is selected for testing to depend on the number of members in her social group who are diagnosed as infected. Finally, the model applies to an epidemic before the development of vaccines. Introducing the probability of being vaccinated is another interesting extension of this model.

2.3. *Individual decisions*

Every period t, individuals make two decisions: working at home or outside, and consuming at home or outside. For the rest of the paper, we focus on a simplified version

where the only decision is working either at home or outside. We represent this decision using the binary variable a_{it}, where $a_{it} = 0$ means working outside, and $a_{it} = 1$ means *confinement* at home. We also abstract from the home and consumption teams – they only include the own individual – and focus on the workplace team $\mathcal{W}_i$ that has size $|\mathcal{W}_i|$.

The set of feasible choices for an individual depends on her current state. In particular, diagnosed individuals have mandatory confinement. We use $\mathcal{A}(x_{it})$ to represent the choice set under state x_{it} such that $\mathcal{A}(AD) = \mathcal{A}(SD) = \{1\}$, $\mathcal{A}(Death) = \varnothing$, and at any other state $\mathcal{A}(x_{it}) = \{0, 1\}$. For simplicity, we assume that confinement means that the individual does not have physical relationship with any other member of the society.

The assumption that individuals who are undiagnosed or recovered have the freedom to decide to work at home or outside deserves some explanation. One may be concerned that this decision is taken by the firm's manager or, in the case of mandatory confinement policies, by the government. These are important concerns that we take into account. In fact, we consider that government confinement policies can be applied with very different degrees of flexibility and not uniformly in all the sectors and regions of the economy. Though we can evaluate a hypothetical policy where the government has the ability to lockdown every individual at home, we are interested in more realistic policies that consist of penalties for working or consuming outside, or subsidies for confinement at home. These penalties and subsidies may vary across industries and/or geographic locations.

2.4. *Information structure*

The assumptions about individuals' information are important for predicting individual behavior and diffusion of the virus.

(i) Information about the network. An individual knows the identity of the members of her social group but she does not have information about the structure of the network outside of her own units, e.g., coworkers of coworkers, etc. According to this condition, we assume below that individuals have only information about members of her social group and economy-wide aggregate information provided by government and media.

(ii) Information about own health status. We consider that, without a test, an individual cannot distinguish between being healthy (H), infected asymptomatic undiagnosed (AU), and recovered after being asymptomatic undiagnosed (RAU). We use $\widetilde{H}$ to represent the union of these three states: $\widetilde{H} \equiv H \cup AU \cup RAU$. We assume that an individual's information about her own health status is captured by the variable $\widetilde{x}_{it}$ such that:

$$(1) \qquad \widetilde{x}_{it} = \begin{cases} \widetilde{H} & \text{if} \quad x_{it} \in \{H, AU, RAU\} \\ \\ x_{it} & \text{if} \quad x_{it} \notin \{H, AU, RAU\} \end{cases}$$

For an individual in $\widetilde{H}$, it is important to know the likelihood of being in state H, or AU, or RAU. In particular, her confinement decision can have implications on her future health only if she is in state H, but it is completely irrelevant if she is already in states AU or RAU. Therefore, an individual in state $\widetilde{H}$ forms beliefs about the probability of being in each of the three specific states. We represent these beliefs as the probabilities $B_{it}^{H|\widetilde{H}}$, $B_{it}^{AU|\widetilde{H}}$, and $B_{it}^{RAU|\widetilde{H}}$ such that $B_{it}^{H|\widetilde{H}} + B_{it}^{AU|\widetilde{H}} + B_{it}^{RAU|\widetilde{H}} = 1$. These beliefs are part of the individual's information set at period t. In section 2.8 below, we describe our assumptions about the initial value and the updating rule of these beliefs.

An individual in state RSU knows that she has experienced symptoms in the past and

now does not have those symptoms, but – similarly as someone in state RAU – she does not know that is immune because she has not been diagnosed. From the point of view of an individual's information, state RSU is different to $\widetilde{H}$ only if the symptoms from COVID-19 are different to those from other diseases, like the common flu. For instance, if COVID symptoms were clearly distinguishable, then state RSU would be equivalent to state RSD. At the other extreme, if the symptoms were the same as those from a common flu, then state RSU would be part of $\widetilde{H}$. More generally, we can have a probabilistic belief that captures the informative content of COVID symptoms. In our numerical experiments in section 3, we have assumed that state RSU is equivalent to RSD.

(iii) Information about health statuses of members of the own team. An individual knows the value $\widetilde{x}_{jt}$ for any other individual in her social group, $\{\widetilde{x}_{jt} : j \in \mathcal{H}_i \cup \mathcal{W}_i \cup \mathcal{C}_i\}$.

(iv) Information about health statuses of individuals outside the own team. An individual does not know the health status of individuals outside her team. However, she has information at the aggregate level for the whole economy. In particular, for every state $x \in \mathcal{X}$, she knows the proportion of individuals in state x at period t. We represent this aggregate shares as $S_t(x)$, and $\mathbf{S}_t$ is the vector $\{S_t(x) : x \in \mathcal{X}\}$. The implicit assumption is that the Health Ministry collects this information and communicates it to the citizenship.

(v) Aggregate probability of confinement. An individual has rational beliefs on the equilibrium probability of confinement at period t for each state $\widetilde{x}$. We use $Q_t(\widetilde{x})$ to represent these average probabilities, and $\mathbf{Q}_t$ to represent the vector $\{Q_t(\widetilde{x}) : \widetilde{x} \in \widetilde{\mathcal{X}}\}$.

(vi) Previous day's own decision of confinement. Individual i knows her own choice at previous period, $a_{i,t-1}$. As we explain below in the description of the utility function, lagged choices are payoff relevant because there are costs of changing the form of working – outside or remotely. For computational simplicity, we assume that individuals do not use information on the lagged actions of team members.

(vii) Private information productivity shocks. Finally, individuals are subject to productivity shocks which are their own private information and are independently distributed across individuals and over time. We represent those shocks as $\varepsilon_{it}(0)$ – if working outside – and $\varepsilon_{it}(1)$ – if working at home.

Summarizing, the information set of individual i at period t is:

$$\Omega_{it} = (\widetilde{\mathbf{x}}_{it}, \mathbf{S}_t, \mathbf{Q}_t, \varepsilon_{it}(0), \varepsilon_{it}(1)) \tag{2}$$

where we use the bold letter $\widetilde{\mathbf{x}}_{it}$ to represent in a compact form the vector of state variables $(\widetilde{x}_{it}, B_{it}^{H|\widetilde{H}}, B_{it}^{AU|\widetilde{H}}, a_{i,t-1}, \{\widetilde{x}_{jt} : j \in \mathcal{W}_i\})$.

2.5. *Probability of infection*

Let $n_{it}^{(x,0)}$ be the number of members in i's team who are in state x and choose to work outside. The probability of infection is an increasing function of the number of infected people that individual i interacts with at period t: that is, a function of $n_{it}^{(AU,0)} + n_{it}^{(SU,0)}$. Let π_{it}^I be the probability of infection. Then,

$$\pi_{it}^I = \begin{cases} 0 & \text{if } a_{it} = 1 \\ 1 - (1 - \rho_I)^{n_{it}^{(AU,0)} + n_{it}^{(SU,0)}} & \text{if } a_{it} = 0 \end{cases} \tag{3}$$

where the parameter $\rho_I \in (0,1)$ measures the probability of getting infected from one infectious teammate. The expression assumes independence (and homogeneity) between the events of getting infected from each sick member. Note that π_{it}^I is zero if there are not infected team members.

The probability of infection depends on variables which are not part of the information set of individual i. In particular, individual i does not know the number $n_{it}^{(AU,0)} + n_{it}^{(SU,0)}$ because: (i) she cannot distinguish team members who are healthy from those who are infected but undiagnosed; and (ii) she does not know their current confinement decisions a_{jt}. Given her information set Ω_{it}, individual i forms expectations about her infection probability π_{it}^I. We describe these beliefs in section 2.9.

2.6. *Production function*

The amount of output generated by an individual (Y_{it}) depends on her own health status and confinement choice, and on the health statuses and confinement choices of her coworkers. If an individual is diagnosed with infection, she is isolated and does not participate in production such that her output is zero. Therefore, we have that $Y_{it} = 0$ if $x_{it} \in \{AD, SD, Death\}$. For the other states, the production function is:

$$Y_{it} = \alpha(a_{it}) + \beta(a_{it}, 0)\, n_{it}^{(a=0)} + \beta(a_{it}, 1)\, n_{it}^{(a=1)} \tag{4}$$

where $\alpha(0)$, $\alpha(1)$, $\beta(0,0)$, $\beta(0,1)$, $\beta(1,0)$, $\beta(1,1)$, $\gamma(0)$, and $\gamma(1)$ are structural parameters, and $n_{it}^{(a=0)}$ and $n_{it}^{(a=1)}$ are the numbers of other individuals in the production team who decide to work at the workplace and remotely from home, respectively.

Parameter $\alpha(a)$ represents the output of an individual when nobody else in the production unit works outside and her confinement choice is a. We expect $\alpha(0) > \alpha(1)$ since confinement reduces an individual's feasible actions. Parameter $\beta(a, a')$ measures the contribution of a coworker to the output of an individual when the coworker's confinement choice is a' and the individual's choice is a. We expect $\beta(a,0) > \beta(a,1)$ and $\beta(0,a') > \beta(1,a')$. Furthermore, we expect to have complementarity (supermodularity) between the working outside decisions of coworkers such that:

$$\beta(0,0) - \beta(0,1) - \beta(1,0) + \beta(1,1) > 0 \tag{5}$$

2.7. *Preferences*

An individual's utility depends on the utility from consumption, $u(C_{it})$, plus the utility from her health status, $\phi(x_{it})$, and minus adjustment costs $\omega(a_{it}, a_{i,t-1})$.[9] We do not consider intertemporal consumption smoothing, such that consumption is equal to output minus net taxes (taxes minus subsidies): $C_{it}(a_{it}) = Y_{it}(a_{it}) - \tau_i(a_{it}, x_{it})$. Net taxes, $\tau_i(a_{it}, x_{it})$, may depend on the individual's confinement decision and on her health/diagnosis state.[10] The utility function is:

$$U_{it}(a_{it}) \;=\; u(Y_{it}(a_{it}) - \tau_i(a_{it}, x_{it})) + \phi(x_{it}) - \omega(a_{it}, a_{i,t-1}) + \varepsilon_{it}(a_{it}) \tag{6}$$

[9] The cost of no change is zero, such that $\omega(0,0) = \omega(1,1) = 0$.

[10] For instance, we may think in different tax/subsidy policies for immune individuals.

where $\{\phi(x) : x \in \mathcal{X}\}$, $\omega(1,0)$, and $\omega(0,1)$ are parameters, and $\varepsilon_{it}(0)$ and $\varepsilon_{it}(1)$ are private information shocks in individual i's utility of working outside and confined, respectively, and they are independently and identically distributed across individuals and over time with an extreme value type I distribution. We assume that $u(.)$ is a linear function: $u(C) = C$.[11]

For the utility from health status, we assume that $\phi(Death) = 0$ and:

$$\phi(x) = \begin{cases} \phi_{alive} + \phi_{health} & \text{for } x \in \{H, AU, AD, RAU, RSU\} \\ \phi_{alive} + \phi_{health} + \phi_{immu} & \text{for } x \in \{RAD, RSD\} \\ \phi_{alive} & \text{for } x \in \{SU, SD\} \end{cases} \tag{7}$$

Parameter ϕ_{alive} represents the flow utility from being alive. Parameter ϕ_{health} represents the extra utility from being (or feeling) healthy. Since an individual cannot distinguish between states H, AU, or RAU, we assume that these states report the same utility. Parameter ϕ_{immu} captures the additional utility from the knowledge of being recovered and immune.

Changing the location for working involves adjustment costs. Parameter $\omega(1,0)$ is the cost of moving from working outside to working at home; similarly, $\omega(0,1)$ is the cost of moving from working at home to working outside. They capture actual sunk investment costs as well as habits. These costs can play an important role to explain persistence in individual behavior and slow transitions at the aggregate level.

2.8. *Best response under two simplifying assumptions*

The numerical experiments in this paper are based on a version of the model that incorporates simplifying assumptions $A1$ and $A2$.

(A1) Individuals in the recovered states RSU, RAD, and RSD always choose to work outside. Individuals in state SU always choose to work at home. This behavior is common knowledge. According to this assumption, the only individuals free to choose are those in state $\widetilde{H} = \{H \cup AU \cup RAU\}$.

(A2) Individuals are quasi myopic. They are forward looking only in terms of how today's decision of where to work affects their own risk of being infected next period. The intertemporal utility function is:

$$U_{it}(a_{it}) + \delta\, W(x_{i,t+1}[a_{it}]), \tag{8}$$

where δ is the discount factor and $W(x)$ represents a terminal present value of having health x.[12] We use the notation $x_{i,t+1}[a_{it}]$ to emphasize that health status at $t+1$ depends on the confinement decision at t.

[11] Alternatively, our specification can be interpreted as one where the utility function is logarithmic, $u(C) = \ln(C)$, the production function is Cobb-Douglas, $Y_{it} = \exp\{\alpha(a_{it}) + \beta(a_{it},0)\, n_{it}^{(a=0)} + \beta(a_{it},1)\, n_{it}^{(a=1)} + \gamma(a_{it})\, \overline{Q}_t\}$, and taxes are proportional, i.e., $C_{it} = Y_{it}(1-\tau_{it})$.

[12] Function $W(x)$ is the expected and discounted value of future utilities for states that will be visited in the future. A key restriction is that the individual ignores that the stream of future utilities depends on her future behavior. She believes that the transition between health states in the future is based on exogenous probabilities.

Under these conditions, for an individual in state $\widetilde{H}$ the best response is working at home if:

$$\mathbb{E}_{it}[U_{it}(1) - U_{it}(0)] + \delta\, \mathbb{E}_{it}[W(x_{i,t+1}[1]) - W(x_{i,t+1}[0])] \geq 0 \tag{9}$$

where the expectation $\mathbb{E}_{it}$ is taken over the individual's beliefs, that we specify in section 2.9 below. The term $\mathbb{E}_{it}[U_{it}(1)-U_{it}(0)]$ is equal to $\widetilde{\alpha}+\widetilde{\tau}+\widetilde{\beta}_0\, \mathbb{E}_{it}(n_{it}^{(a=0)})+\tilde{\beta}_1\, \mathbb{E}_{it}(n_{it}^{(a=1)})+\widetilde{\varepsilon}_{it}$, with $\widetilde{\alpha} \equiv \alpha(1)-\alpha(0)$, $\widetilde{\tau} \equiv \tau(1)-\tau(0)$, $\widetilde{\beta}_0 \equiv \beta(1,0)-\beta(0,0)$, $\widetilde{\beta}_1 \equiv \beta(1,1)-\beta(0,1)$, and $\widetilde{\varepsilon}_{it} \equiv \varepsilon_{it}(1) - \varepsilon_{it}(0)$. The probability distribution of next period health depends on today's confinement decision only if the individual is currently healthy ($x_{it} = H$). In that case, next period health can take two possible values: H or AU. If $a_{it} = 1$, we have that $x_{i,t+1} = H$ with probability one. If $a_{it} = 0$, we have that $x_{i,t+1}$ is equal to AU with probability $\mathbb{E}_{it}(1-(1-\rho_I)^{n_{it}^{(AU,0)}})$ and is equal to H with probability $\mathbb{E}_{it}((1-\rho_I)^{n_{it}^{(AU,0)}})$, where the expectation $\mathbb{E}_{it}(.)$ is taken over the the individual's beliefs about the distribution of $n_{it}^{(AU,0)}$. Therefore, we have that:

$$\mathbb{E}_{it}\left[W(x_{i,t+1}[1]) - W(x_{i,t+1}[0])\right] = B_{it}^{H|\widetilde{H}}\left[W_H - W_{AU}\right]\left[1 - \mathbb{E}_{it}((1-\rho_I)^{n_{it}^{(AU,0)}})\right] \tag{10}$$

When the parameter ρ_I is close to zero, we have that $1 - \mathbb{E}_{it}((1-\rho_I)^{n_{it}^{(AU,0)}})$ can be approximated well using $\rho_I\, \mathbb{E}_{it}(n_{it}^{(AU,0)})$.

Putting all these pieces together, we have that the best response is working at home if the following condition holds:

$$\widetilde{\alpha}+\widetilde{\tau}+\widetilde{\beta}_0\, \mathbb{E}_{it}(n_{it}^{(a=0)})+\tilde{\beta}_1\, \mathbb{E}_{it}(n_{it}^{(a=1)})+\widetilde{\varepsilon}_{it}+B_{it}^{H|\widetilde{H}}\left[W_H - W_{AU}\right]\, \rho_I\, \mathbb{E}_{it}(n_{it}^{(AU,0)}) \geq 0 \tag{11}$$

To complete the characterization of this best response condition, we need to specify individuals' beliefs and their evolution over time.

2.9. *Beliefs*

Equation (11) shows that, to make her best response, an individual needs to form beliefs about the probability of her actual health status ($B_{it}^{H|\widetilde{H}}$), and about the health statuses and confinement decisions of her team members ($\mathbb{E}_{it}(n_{it}^{(a=0)})$, $\mathbb{E}_{it}(n_{it}^{(a=1)})$, and $\mathbb{E}_{it}(n_{it}^{(AU,0)})$). This section describes our conditions on these beliefs.

(i) Current health status and confinement choices of other individuals in the team. For any state x different to $\widetilde{H}$, individual i knows the number of coworkers in that state – $n_{it}^{(x)}$ – but she does not know their current confinement choices: i.e., she does not know $n_{it}^{(x,0)}$ and $n_{it}^{(x,1)}$. We assume that individuals use the aggregate frequencies in $Q_t(\widetilde{x})$ to form probabilistic beliefs about the choices of team members in state $\widetilde{x}$. Individuals believe that $n_{it}^{(x,0)}$ has a Binomial distribution with arguments $n_{it}^{(x)}$ and $1 - Q_t(x)$. Accordingly, and using assumption (A1), the expected values of the number of team members working outside and at home are:

$$(12) \quad \begin{cases} \mathbb{E}_{it}\left(n_{it}^{(a=0)}\right) &= n_{it}^{(\widetilde{R})} + n_{it}^{(\widetilde{H})}[1 - Q_t(\widetilde{H})] \\ \mathbb{E}_{it}\left(n_{it}^{(a=1)}\right) &= n_{it}^{(AD)} + n_{it}^{(SU)} + n_{it}^{(\widetilde{H})}\, Q_t(\widetilde{H}) \end{cases}$$

where $\widetilde{R}$ represents the union of states $\{RSU \cup RAD \cup RSD\}$.

An individual also knows the number of coworkers in state $\widetilde{H}$, that we denote as $n_{it}^{(\widetilde{H})}$. But she does not the values of $n_{it}^{(x,a)}$ for $x = H,\, AU,\, RAU$ and $a = 0, 1$. We assume that individuals use the aggregate frequencies in $\mathbf{S}_t$ to form their beliefs about the actual health status of a team member at state $\widetilde{H}$. Therefore, for $x \in \widetilde{H}$, individuals believe that variable $n_{it}^{(x,a)}$ has a Binomial distribution with arguments $n_{it}^{(\widetilde{H})}$ and $p_{it}^{(x,a)|\widetilde{H}}$, where:

$$(13) \qquad p_{it}^{(x,a)|\widetilde{H}} \equiv \frac{S_t(x)}{S_t(\widetilde{H})}\,(1 - Q_t(\widetilde{H}))^{1-a}\, Q_t(\widetilde{H})^a.$$

In particular, individual i beliefs that $n_{it}^{(AU,0)}$ has a Binomial distribution with arguments $n = n_{it}^{(\widetilde{H})}$ and $p = \frac{S_t(AU)}{S_t(\widetilde{H})}\,(1 - Q_t(\widetilde{H}))$ such that:

$$(14) \qquad \mathbb{E}_{it}(n_{it}^{(AU,0)}) = n_{it}^{(\widetilde{H})}\, \frac{S_t(AU)}{S_t(\widetilde{H})}\,(1 - Q_t(\widetilde{H}))$$

These beliefs differ from the ones of a perfectly rational (Bayesian) individual in two aspects. First, she does not use the previous history of health statuses of team members to form beliefs about their current health statuses – H, AU, or RAU – and review these beliefs using Bayesian updating. Instead, she uses only the aggregate probabilities $S_t(x)$. Second, by using the aggregate probability of confinement, $Q_t(x)$, individual i is not taking into account that her own health status can affect her teammates' confinement decisions. For instance, the probability of confinement of a team member can be larger if i's own health is SD than if it is H.

(ii) Expected probability of infection. Conditional on working choice a_{it}, the expected probability of infection for individual i – that we denote as $\overline{\pi}_{it}^{I,own}(a_{it})$ – has the following expression:

$$(15) \quad \overline{\pi}_{it}^{I,own}(a_{it}) \equiv (1 - a_{it}) \sum_n BIN\left(n \mid n_{it}^{(\widetilde{H})}, \frac{S_t(AU)}{S_t(\widetilde{H})}(1 - Q_t(\widetilde{H}))\right)[1 - (1 - \rho_I)^n]$$

where $BIN(n|N, p)$ represents the density function of a Binomial distribution with parameters N and p.

(iii) Evolution of beliefs about own health status $\widetilde{H}$. An individual in state $\widetilde{H}$ uses probabilistic beliefs about her actual status, that we denote as $B_{it}^{H|\widetilde{H}}$, $B_{it}^{AU|\widetilde{H}}$, and $B_{it}^{RAU|\widetilde{H}}$. Every period, she updates these beliefs using new information. At period t, if the individual remains in state $\widetilde{H}$, the beliefs about her actual health status are updated using the

following natural formula. For $x = H$, AU, RAU:

$$B_{it}^{x|\widetilde{H}} = \frac{F_{it}^{i}(x \mid \widetilde{H}, a_{i,t-1})}{F_{it}^{i}(H \mid \widetilde{H}, a_{i,t-1}) + F_{it}^{i}(AU \mid \widetilde{H}, a_{i,t-1}) + F_{it}^{i}(RAU \mid \widetilde{H})} \tag{16}$$

where $F_{it}^{i}(x|\widetilde{H}, a)$ is individual i's (subjective) probability of $x_{it} = x$ given that $\widetilde{x}_{i,t-1} = \widetilde{H}$ and $a_{i,t-1} = a$. They have the following form:

$$\begin{cases} F_{it}^{i}(H \mid \widetilde{H}, a_{i,t-1}) &= B_{i,t-1}^{H|\widetilde{H}} \left(1 - \overline{\pi}_{i,t-1}^{I,own}(a_{i,t-1})\right) \\ F_{it}^{i}(AU \mid \widetilde{H}, a_{i,t-1}) &= B_{i,t-1}^{AU|\widetilde{H}} \left(1 - \pi_{+A} - \pi_{-A}\right)\left(1 - \lambda_A\right) + B_{i,t-1}^{H|\widetilde{H}} \overline{\pi}_{i,t-1}^{I,own}(a_{i,t-1}) \end{cases} \tag{17}$$

The first equation says that an individual is healthy at t if she was healthy at period $t-1$ – that has subjective belief $B_{i,t-1}^{H|\widetilde{H}}$ – and was not infected during that period – that has subjective belief $1 - \overline{\pi}_{i,t-1}^{I,own}(a_{i,t-1})$. The second equation establishes that she arrives to state AU at t either if she was at state AU at period $t-1$ and she gets a neutral health shock and no testing, or shes was at state H at period $t-1$ and gets infected. The rest of the transition probabilities from state $\widetilde{H}$ do not depend on $a_{i,t-1}$.
Note that this updating rule depends on the individual's previous confinement choice and on previous infection probabilities. In particular, $B_{it}^{H|\widetilde{H}}$ is greater with $a_{i,t-1} = 1$ than with $a_{i,t-1} = 0$.

2.10. *Equilibrium*

The best response of an individual in state $\widetilde{H}$ can be represented as a choice probability that results from integrating the condition in equation (11) over the distribution of the shocks $\widetilde{\varepsilon}_{it}$. Putting together equations (11), (12), and (14), the best response probability of confinement is:

$$P_{it}(\widetilde{H}) = \Lambda\left(\alpha_{it} + \beta_{it}\, Q_t(\widetilde{H})\right) \tag{18}$$

with

$$\begin{cases} \alpha_{it} &= \widetilde{\alpha} + \widetilde{\tau} + \widetilde{\beta}_0(n_{it}^{(\widetilde{R})} + n_{it}^{(\widetilde{H})}) + \widetilde{\beta}_1\, n_{it}^{(AD)} + \delta B_{it}^{H|\widetilde{H}} \left[W_H - W_{AU}\right] \rho_I\, n_{it}^{(\widetilde{H})} \frac{S_t(AU)}{S_t(\widetilde{H})} \\ \beta_{it} &= \left[\widetilde{\beta}_1 - \widetilde{\beta}_0 - \delta B_{it}^{H|\widetilde{H}} \left[W_H - W_{AU}\right] \rho_I\, \frac{S_t(AU)}{S_t(\widetilde{H})}\right] n_{it}^{(\widetilde{H})} \end{cases} \tag{19}$$

Under condition (A1), the probability of confinement $Q_t(\widetilde{H})$ is the only endogenous element in the vector $\mathbf{Q}_t$ such that the equilibrium mapping is a scalar function. This equilibrium depends on individuals' best response probabilities of confinement $P_{it}(\widetilde{H})$.

Let $\widetilde{\mathbf{x}}_t$ be the vector $(x_{it},\, a_{i,t-1},\, B_{it}^{H|\widetilde{H}},\, B_{it}^{AU|\widetilde{H}} : i = 1, 2, ..., N)$. That is, $\widetilde{\mathbf{x}}_t$ contains health status, current beliefs about status conditional on $\widetilde{H}$, and last period choice of every individual in the economy.[13]

[13] Since $\mathbf{S}_t$ is the vector of frequencies of each health status, we have that $\mathbf{S}_t$ is a deterministic function of $\widetilde{\mathbf{x}}_t$.

DEFINITION. Given $\widetilde{\mathbf{x}}_t$ (and the corresponding $\mathbf{S}_t$), an equilibrium is a vector of probabilities of confinement $P_{it}(\widetilde{H})$ – one for each individual – and an aggregate probability of confinement $Q_t(\widetilde{H})$ that satisfy the following conditions: (i) $P_{it}(\widetilde{H})$ is a best response probability that satisfies equation (18); and (ii) $Q_t(\widetilde{H})$ is the probability that results from the aggregation of individual choice probabilities:

$$Q_t(\widetilde{H}) = \frac{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}\ P_{it}(\widetilde{H})}{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}} \tag{20}$$

Equilibrium conditions (i) and (ii) can be combined into a single condition: the aggregate probability of confinement $Q_t(\widetilde{H})$ is a solution to the following fixed point mapping Ψ_t:

$$Q_t(\widetilde{H}) = \Psi_t\left(Q_t(\widetilde{H})\right) \equiv \frac{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}\ \Lambda\left(\alpha_{it} + \beta_{it}\ Q_t(\widetilde{H})\right)}{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}} \quad \square \tag{21}$$

Remark 1. Equilibrium existence. The equilibrium mapping Ψ_t in equation (21) is continuous on the compact set $[0, 1]$. By Brower's theorem, an equilibrium exists.

Remark 2. Strategic complementarity vs. substitutability in individuals' confinement decisions. If an individual's propensity to confinement $P_{it}(\widetilde{H})$ increases with the aggregate probability of confinement $Q_t(\widetilde{H})$, then we say that confinement decisions are *strategic complements* and we have a coordination game. Otherwise, confinement decisions are *strategic substitutes* and we have an entry game. Depending on the values of the parameters, this model can generate either complementarity or substitutability. The complementarity between individuals' confinement decisions in the production function can generate strategic complementarity in this game. However, individuals' concern for their future health can generate strategic substitutability. The larger the proportion of confined individuals, the lower the probability of getting infected and the smaller the expected health benefits of current confinement.

The analytical expression for the derivative $\Psi'_t\left(Q_t(\widetilde{H})\right)$ provides a simple approach to determine whether there is strategic complementarity or substitutability between individuals' confinement decisions. The sign of β_{it} determines the sign of Ψ'_t such that we have complementarity when $\beta_{it} > 0$ and substitutability when $\beta_{it} < 0$.

Strategic complementarity or substitutability in confinement decisions can have important policy implications. Under complementarity, small incentives to confinement may generate large changes in the aggregate probability $Q_t(\widetilde{H})$. In contrast, under substitutability, it may be difficult to achieve a high aggregate probability of confinement.

Remark 3. Equilibrium uniqueness. We can also use the analytical expression of the derivative Ψ'_t to obtain a sufficient condition for equilibrium uniqueness that only depends on values of the structural parameters such that it can be easily checked. First, it is clear that if $\Psi'_t < 0$, then the equilibrium is unique. The sign of Ψ'_t is equal to the sign of β_{it}. A sufficient condition for $\beta_{it} < 0$ is that $\widetilde{\beta}_1 - \widetilde{\beta}_0 < 0$, and this condition holds if the production function is submodular in team members' decisions of working outside. However, as established in equation (5) above, we expect the production function to be supermodular. Suppose this is the case, such that $\widetilde{\beta}_1 - \widetilde{\beta}_0 > 0$. Then, there

are always values of $B_{it}^{H|\widetilde{H}}$ and $S_t^{(AU)}/S_t^{(H)}$ close enough to zero such that $\beta_{it} > 0$. This means that, in this case, we cannot find values of the structural parameters that imply $\Psi'_t < 0$. However, we also have equilibrium uniqueness if $0 < \Psi'_t < 1$. Taking into account that the derivative of the logistic function is $\Lambda'(u) = \Lambda(u)\,[1 - \Lambda(u)] \leq 1/4$, and that $\beta_{it} < [\widetilde{\beta}_1 - \widetilde{\beta}_0] n_{it}^{\widetilde{H}} < [\widetilde{\beta}_1 - \widetilde{\beta}_0] \max_i |\mathcal{W}_i|$, we have that:

$$
\begin{aligned}
\Psi'_t &= \frac{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}\ \beta_{it}\ \Lambda'_{it}}{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}} \\
&\leq \left[\widetilde{\beta}_1 - \widetilde{\beta}_0\right] \frac{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}\ n_{it}^{(\widetilde{H})}\ \Lambda'_{it}}{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}} \\
&\leq \frac{1}{4}\left[\widetilde{\beta}_1 - \widetilde{\beta}_0\right] \frac{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}\ n_{it}^{(\widetilde{H})}}{\sum_{i\in\mathcal{I}} 1\{\widetilde{x}_{it} = \widetilde{H}\}} \\
&\leq \frac{1}{4}\left[\widetilde{\beta}_1 - \widetilde{\beta}_0\right]\ \max_i |\mathcal{W}_i|
\end{aligned}
\tag{22}
$$

The last expression in equation (22) is an upper bound to Ψ'_t. This upper bound depends only on primitives of the model such that it can be easily calculated before the computation of an equilibrium. If this expression is strictly smaller than 1, then the equilibrium value $Q_t(\widetilde{H})$ is unique. In the numerical experiments in section 3, we confirm that this condition holds for our parameterization / calibration of the model and we apply fixed point iterations to calculate an equilibrium.

Remark 4. Endogenous stochastic process of $\{\widetilde{\mathbf{x}}_t, \boldsymbol{a}_t : t \geq 1\}$. As defined above, the equilibrium concept that we use takes $\widetilde{\mathbf{x}}_t$ as given and it applies to one period. However, this equilibrium concept implies a stochastic process for the vectors of state and decisions variables. Given $\widetilde{\mathbf{x}}_t$, the equilibrium at period t implies an aggregate probability of confinement $Q_t(\widetilde{H})$ and the corresponding probabilities of confinement for every individual i: $P_{it}(\widetilde{H})$. These probabilities define the distribution of the vector of choices $\boldsymbol{a}_t$ conditional on $\widetilde{\mathbf{x}}_t$. Then, the transition probabilities of the health state variable and the updating rule of beliefs define the probability distribution of $\widetilde{\mathbf{x}}_{t+1}$ conditional on $\widetilde{\mathbf{x}}_t$ and $\boldsymbol{a}_t$.

3. NUMERICAL EXPERIMENTS

In this section, we present several numerical experiments to illustrate the properties and predictions of the model.[14] To make the results more transparent, we focus on the implications of different network structures while ignoring other heterogeneity forms. For the Covid-19 pandemic, there is substantial evidence on individual heterogeneity in behavior and in transition probabilities between health statuses. Perhaps, the most notorious is the effect of age on the likelihood of developing symptoms and on subsequent adverse health shocks. Production function parameters may also vary in the population, depending on the

[14]The code in R language for the replication of these numerical experiments can be downloaded at http://jiayinggu.weebly.com/uploads/3/8/9/3/38937991/replication.zip

TABLE I

PARAMETERS IN BENCHMARK SCENARIO (EXPERIMENT 1)

• Population: $N = 10,000$	**Production function** ($F = 8.1435$)
• # individuals infected at period 1: 10	$\alpha(0) + \beta(0,0)\ \|\mathcal{W}\| = F$
• Number of team members: $\|\mathcal{W}\| = 40$	$\alpha(0) + \beta(0,1)\ \|\mathcal{W}\| = 0.40\ F$
	$\alpha(1) + \beta(1,0)\ \|\mathcal{W}\| = 0.35\ F$
Epidemiological parameters	$\alpha(1) + \beta(1,1)\ \|\mathcal{W}\| = 0.20\ F$
• $\pi_{-A} = 1/6$; $\pi_{+A} = 1/7$	$\alpha(0) = 0.20\ F$
• $\pi_{+SU} = 1/14$; $\pi_{-SU} = (10/90)(1/14)$, implying mortality rate at $SU = 10\%$.	$\alpha(1) = 0.05\ F$
	$\Lambda\left(\alpha(1) - \alpha(0) + [\beta(1,0) - \beta(0,0)]\|\mathcal{W}\|\right) = 0.005$
• $\pi_{+SD} = 1/10$; $\pi_{-SD} = (5/95)(1/10)$ implying mortality rate at $SD = 5\%$	**Preferences**
• Infection rate: $\rho_I = 0.27$	$\Lambda(\alpha(1) - \alpha(0) + [\beta(1,1) - \beta(0,1)]\|\mathcal{W}\| + \delta\,[W_H - W_{AU}])$=0.99 and this implies $\delta\,[W_H - W_{AU}] = 6.223$

education of the individual or her occupation. In our model, we can allow all the structural parameters to be functions of time-invariant individual demographics that are observable to the researcher, e.g., age, gender, education. We believe that the structural estimation of our model using microeconomic data would involve this type of specification. However, in our numerical experiments, we omit this type of heterogeneity. In this paper, we are particularly interested in studying the implications of different network structures, and the results would be noisier if we included other forms of heterogeneity.

Based on our calibration, we solve the model and simulate the path of the endogenous variables under different experiments. Experiment 1 is our benchmark scenario and is characterized by a ring network structure, an initial herd immunity of 0%, and no public interventions – no testing and no subsidies. Each of the other experiments incorporates a specific modification with respect to this benchmark. In experiment 2, the initial level of herd immunity is 67%. In experiment 3, we incorporate a government subsidy to work from home. We present results for three levels of this subsidy: 20%, 30%, and 40% of an individual's earnings if her workplace works at full capacity, i.e., when all the workers are active and working in site. In experiment 4, we introduce testing. We present results for three different levels of testing rate: 2%, 10%, and 20% for asymptomatic individuals, and 80% for symptomatic. Finally, in experiment 5, we modify the structure of the network of social connections.

3.1. *Parameterization / Calibration*

Table I presents then parameters used for the benchmark scenario (Experiment 1).

Population. We consider a population with $10,000$ individuals.[15]

Number of individuals infected at the initial period. At day 1, there are 10 individuals (i.e., 0.1% of the population), randomly selected, who are infected and undiagnosed (state AU). The rest of the individuals are in state H.

[15]Computation time scales up with the size of the population, but the results with a larger population of $100,000$ individuals share very similar dynamics to the results we present below.

Network structure. In these experiments we consider four different types of networks: a ring lattice; a small world network; a caveman graph; and a randomly rewired caveman graph. Figure 3 presents examples of these network structures with 25 individuals.

A *ring lattice* is a regular graph – each node has the same degree – where nodes are arranged in a circle with each node connected to $|\mathcal{W}|$ nearest neighbors. A ring network with $N >> |\mathcal{W}|$ has a high clustering coefficient and high average path length. A ring lattice does not capture the observed degree heterogeneity that we find in actual social networks. The network at the bottom left of figure 3 is a ring lattice with 25 individuals each having 5 edges.

The *small world network* is a variation of a ring lattice where some nodes are randomly rewired. This random rewiring has the effect of reducing the average path length. An example of a small world network is shown at the lower right of figure 3.

A *caveman graph* consists of several local clusters where nodes within a cluster are highly connected but there is very little connection between clusters. In our model, this graph may represent an economy where there is very small overlapping between production teams. Compared to a ring lattice, a caveman network shares the feature of a large average path length, but their local structures are very different.

Finally, we consider a variation of the caveman graph where some nodes are randomly rewired. Figure 3 (upper right) shows an example, with 25 individuals, 5 local clusters, and rewiring probability of 0.5. Rewiring decreases average path length making the network more connected.

In our benchmark economy (experiment 1), the network consists of ring lattice with $N = 10,000$ and $|\mathcal{W}| = 40$. In experiment 5, we modify this network structure by changing the value of the parameter $|\mathcal{W}|$ and by considering the other three types of networks.

Epidemiological parameters for COVID-19. We let $\pi_{-A} = 1/6$ (i.e., average incubation period of 6 days), and $\pi_{+A} = 1/7$ (i.e., 7 days of average waiting time to recovery if no symptoms).

For the recovery of symptomatic-undiagnosed individuals (state SU), we set $\pi_{+SU} = 1/14$, i.e., 14 days of average waiting time to recovery after developing symptoms. Parameter π_{-SU} is set to match a mortality rate of 10% for these undiagnosed (untreated) individuals. That is, $\pi_{-SU}/\pi_{+SU} = 10/90$, and this implies $\pi_{-SU} = 0.0079$.

For the recovery of symptomatic-diagnosed individuals (state SD), we set $\pi_{+SD} = 1/10$, i.e., 10 days of average waiting time to recovery with diagnosis and treatment. Parameter π_{-SD} is set to match a mortality rate of 5% for these diagnosed (treated) individuals. That is, $\pi_{-SD}/\pi_{+SD} = 5/95$, and this implies $\pi_{-SD} = 0.0053$.

Production function. We consider that $\gamma(0) = \gamma(1) = 0$ such that the production function is $Y_i = \alpha(a_i) + \beta(a_i, 0)\ n_i^{(a=0)} + \beta(a_i, 1)\ n_i^{(a=1)}$, where $n_i^{(a=0)}$ is the number of other members physically present in the workplace, and $n_i^{(a=1)}$ is the number of those working from home. The selection of the parameters $\alpha(0)$, $\alpha(1)$, $\beta(0,0)$, $\beta(0,1)$, $\beta(1,0)$, and $\beta(1,1)$ is based on the following conditions that relate these parameters with the amount of output at full capacity, F.

a. Full capacity output: everybody works outside: $\alpha(0) + \beta(0,0)\ |\mathcal{W}| \equiv F$.
b. Own individual works outside, team members work at home: $\alpha(0) + \beta(0,1)|\mathcal{W}| = 0.40\ F$.
c. Own individual works at home, team members work outside: $\alpha(1) + \beta(1,0)|\mathcal{W}| = 0.35\ F$.

d. Everyone works from home: $\alpha(1) + \beta(1,1)|\mathcal{W}| = 0.20\ F$.
e. Own individual works outside, all members inactive: $\alpha(0) = 0.20\ F$.
f. Own individual works at home, all members inactive: $\alpha(1) = 0.05\ F$.

It is straightforward to verify that conditions (a) to (f) imply that function $\beta(a, a')$ is: monotonic in its two arguments, i.e., $\beta(0,0) > \beta(1,0)$ and $\beta(0,0) > \beta(0,1)$; and supermodular, i.e., $\beta(0,0) - \beta(0,1) - \beta(1,0) + \beta(1,1) > 0$.

The value for full capacity output F is chosen such that the probability of working at home when there is zero risk of infection and all the team members are working outside is 0.5%. That is, $\Lambda(\alpha(1) + \beta(1,0)|\mathcal{W}| - \alpha(1) - \beta(0,0)|\mathcal{W}|) = 0.005$, and given conditions (a) to (f), this implies that $\Lambda(-0.65F) = 0.005$, or equivalently, $F = 8.1435$.

Difference in present values of healthy and infected. The parameter $\delta\,[W(H) - W(AU)]$ – that captures the difference between the present values of being healthy and being infected – is chosen such that an individual with no active team members and with probability one of getting infected chooses working at home with probability 99%. That is, $\Lambda(\alpha(1) - \alpha(0) + [\beta(1,1) - \beta(0,1)]|\mathcal{W}| + \delta\,[W(H) - W(AU)]) = 0.99$. The daily discount factor δ can be interpreted as equal to one.

Checking for equilibrium uniqueness. We have checked the sufficient condition for equilibrium uniqueness presented in equation (22). Given the model parameters, the upper bound for the derivative of the equilibrium mapping, Ψ'_t is strictly smaller than one. This implies that for every period t and for all our experiments, the equilibrium value of $Q_t(\widetilde{H})$ is unique.

3.2. *Results*

3.2.1. *Experiment 1: No interventions*

Figure 4 presents the simulated paths of endogenous variables in our benchmark scenario with a ring network, without testing, and without subsidies.

Number of new infections per day (bottom row, column 1) and cumulative share of infected individuals (row 1, column 4). During the first 10 days, the rate of infection increases very rapidly. It then slows down mainly due to the prominent response of self-confinement. By day 80 practically the whole population has got infected. At its peak, the number of new infections per day reaches 400 (4% of the total population).

Probability of confinement (top row, column 1). This probability responds endogeneously – with some lag – to the rapid expansion of the virus. The response is quite promptly and has a peak at day 10 when about 15% of the workers decide (voluntarily) to be confined, and then declines slowly as daily new cases taper off.

Share of deceased (top row, column 2) and share of immune (top row, column 3). The rapid expansion of the virus generates a fast convergence to the new steady-state. This steady-state is practically achieved after only 130 days. The share of deaths is slightly above 5%.

Aggregate output (bottom row, column 2). The long-run (permanent) effect of the virus on aggregate output is a reduction by 7.5%. This long-run effect is due to 5% reduction

Figure 3: Examples of Network Structures (25 individuals)

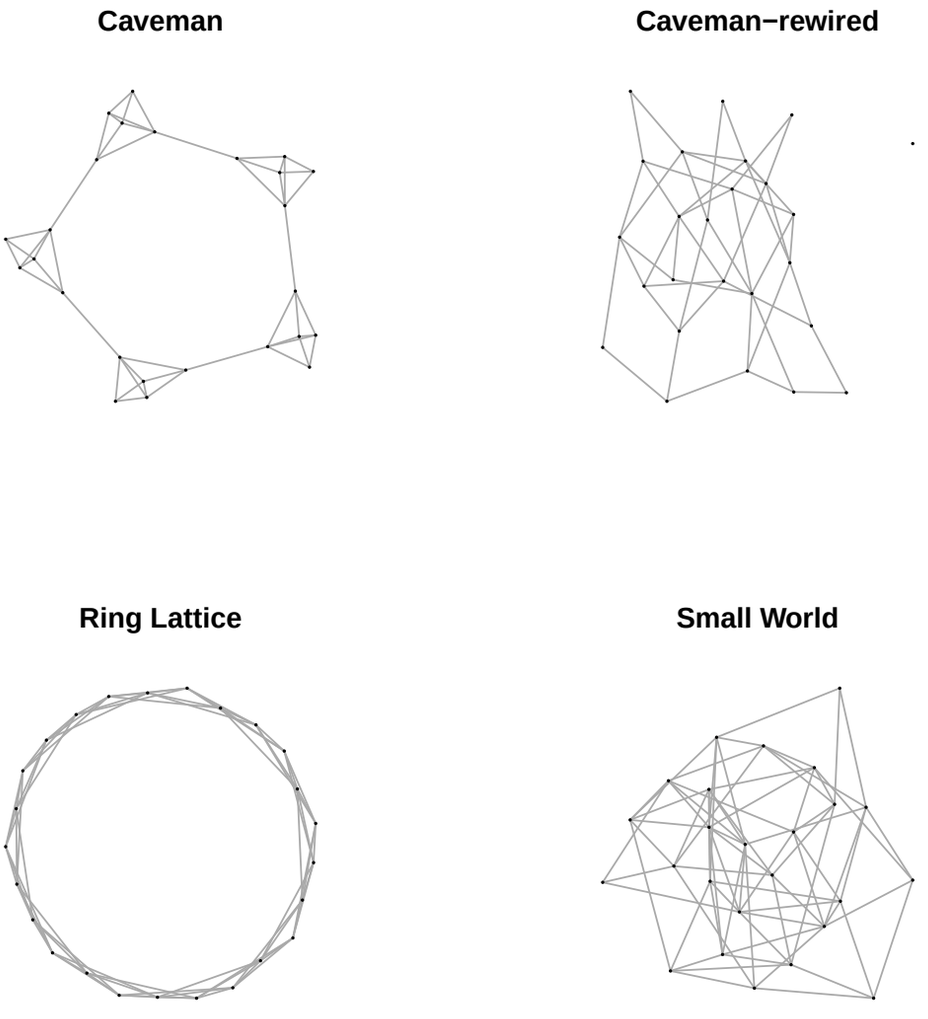

in the labor force (deaths) together with the complementarity in the production function. The steady-state amount of output is reached after 130 days periods, and it follows after a very deep recession that lasts approximately 10 days and in which output becomes 75% of its full capacity.

3.2.2. *Experiment 2: Herd immunity*

Figures 5 presents the results for experiment 2 where everything is the same as in our benchmark except that at period $t = 1$ the share of immune individuals is 67%.

The goal of this experiment to show the value of immunity. We can interpret this immunity as the result of vaccination. We can also interpret it as a second wave of the virus that arrives once a substantial proportion of the population is already immune because of previous recovery after infection. The evolution of all the endogenous variables is dramatically different than under the benchmark scenario. The diffusion of the virus is very slow. The number of infected individuals per day is always lower than 30 and the effect on output is negligible. Herd immunity has a nonlinear effect on the diffusion of the virus.

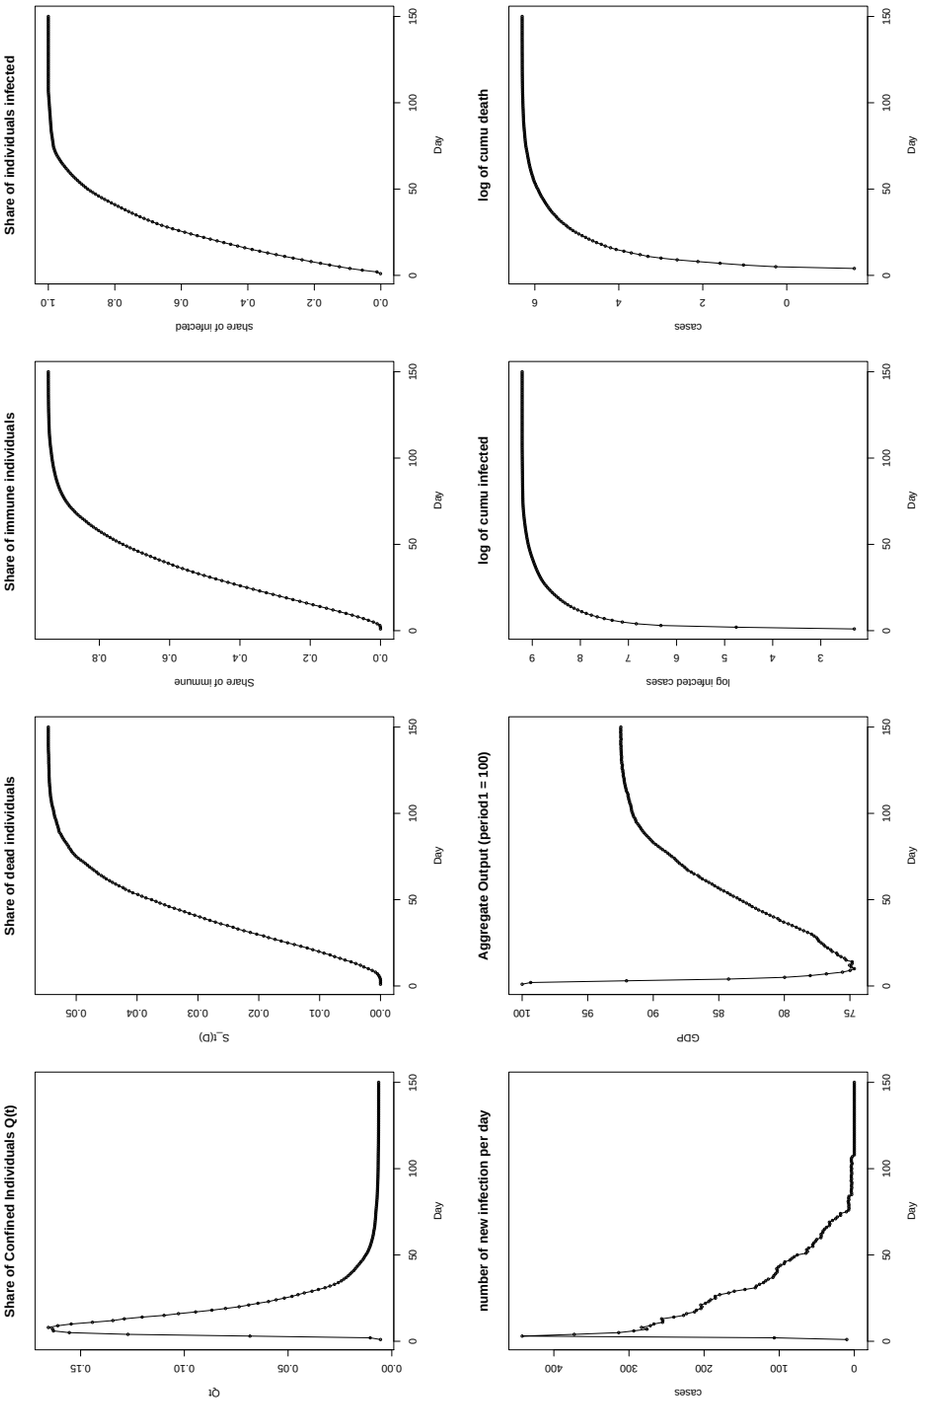

Figure 4: Experiment 1: No Public Interventions

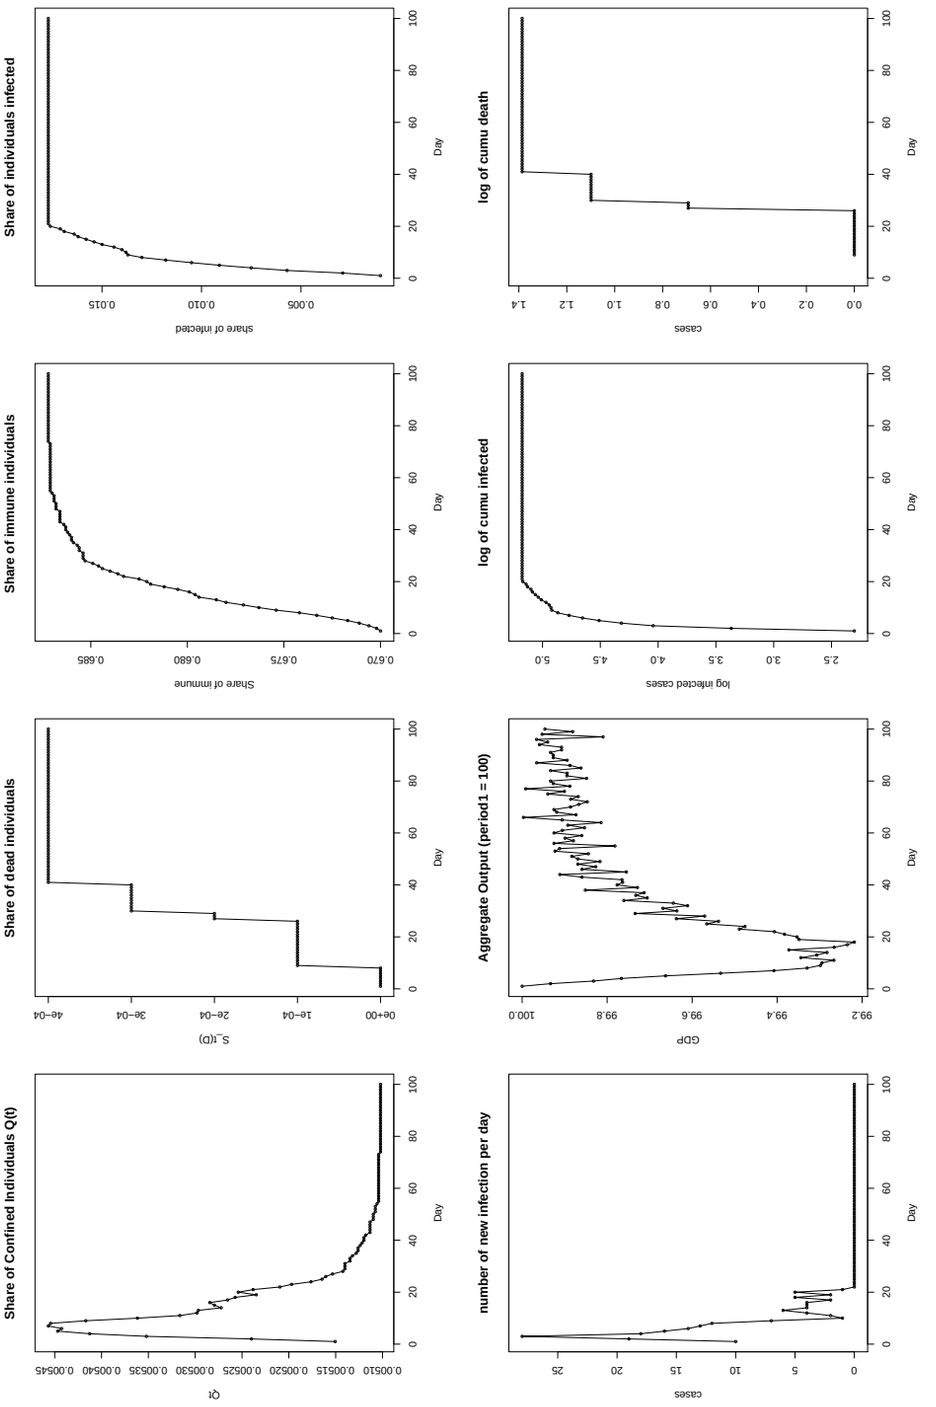

Figure 5: Experiment 2: Initial Herd Immunity of 67 Percent

3.2.3. *Experiment 3: Subsidy to work from home*

In experiment 3, we modify the benchmark scenario by including a subsidy to work from home. We implement experiments for three different values of the subsidy: 10%, 20%, and 30% of the full capacity output. That is, $\tau(0) = 0$ (no subsidy or tax for working outside), and $\tau(1) = 0.1F$, or $0.2F$, or $0.3F$, respectively.

Figure 6 shows the effects of each subsidy rate with respect to the benchmark. The subsidy – even at 10% which is not large as a percentage of the full salary – has a very strong positive effect on confinement decisions, especially during the peak of the virus expansion. At this peak, the probability of working at home increases from 15% without the subsidy to about 30%. It has also an important effect the total output, driven by workers' voluntary confinement decision. With 10% subsidy, the deepest point of the recession marks 62% of the output under full capacity. As expected, all these effects become larger when the subsidy rate increases.

3.2.4. *Experiment 4: Testing*

In experiment 4, we modify the benchmark scenario by introducing testing, both to symptomatic individuals – with a probability $\lambda_S = 80\%$ – and to asymptomatic individuals – with three experiments where the probability λ_A takes the values 2%, 10%, and 20%, respectively.

Figure 7 presents the effects with respect to the benchmark. Very interestingly, we find that the effects of introducing testing are basically the complements of the subsidy to confinement. Testing has practically zero effect on the number of infections per day. However, it has an important effect on the timing of confinement and on output, and especially in the number of deaths. Testing identifies infected individuals and removes them from the labor force. Due to complementarity in the production process, this has a positive incentive on confinement. Now, the peak of confinement occurs earlier than in the benchmark. This has also an effect on output: the recession is slightly not as deep and it is substantially shorter. The long run effects are also smaller, due to the savings of lives.

3.2.5. *Experiment 5: Modifying the network structure*

Here we include a set of experiments to illustrate how the structure of the social/production network influences virus diffusion and its economic impact. Very interestingly, we show that the network structure not only has an impact on the infection status of the population, but it also affects individuals' confinement decisions.

Figure 8 presents the simulated paths of the endogenous variables compared tothe benchmark model. The small world network (green curve) is generated from the same ring lattice as the benchmark but with a re-wiring probability of 1% using the Watts-Strogatz algorithm. This modification of the network (which reduces the average path length of the network) brings a very different virus diffusion. It starts off slower than the benchmark case, however the momentum of the virus takes much longer before it tapers off, leading to a faster trajectory to get the whole population infected. The response of the confinement decision also paints a very different picture compared to the benchmark case. Only a tiny fraction of workers decide to work from home.

The ring lattice with a smaller degree – $|\mathcal{W}| = 20$ instead of $|\mathcal{W}| = 40$ – also paints a different picture compared to the benchmark. In this network, with larger social distance, the virus diffuses more slowly. As a result, the share of individuals who choose to work

Figure 6: Experiment 3: Subsidy to Work from Home. Differences with Benchmark

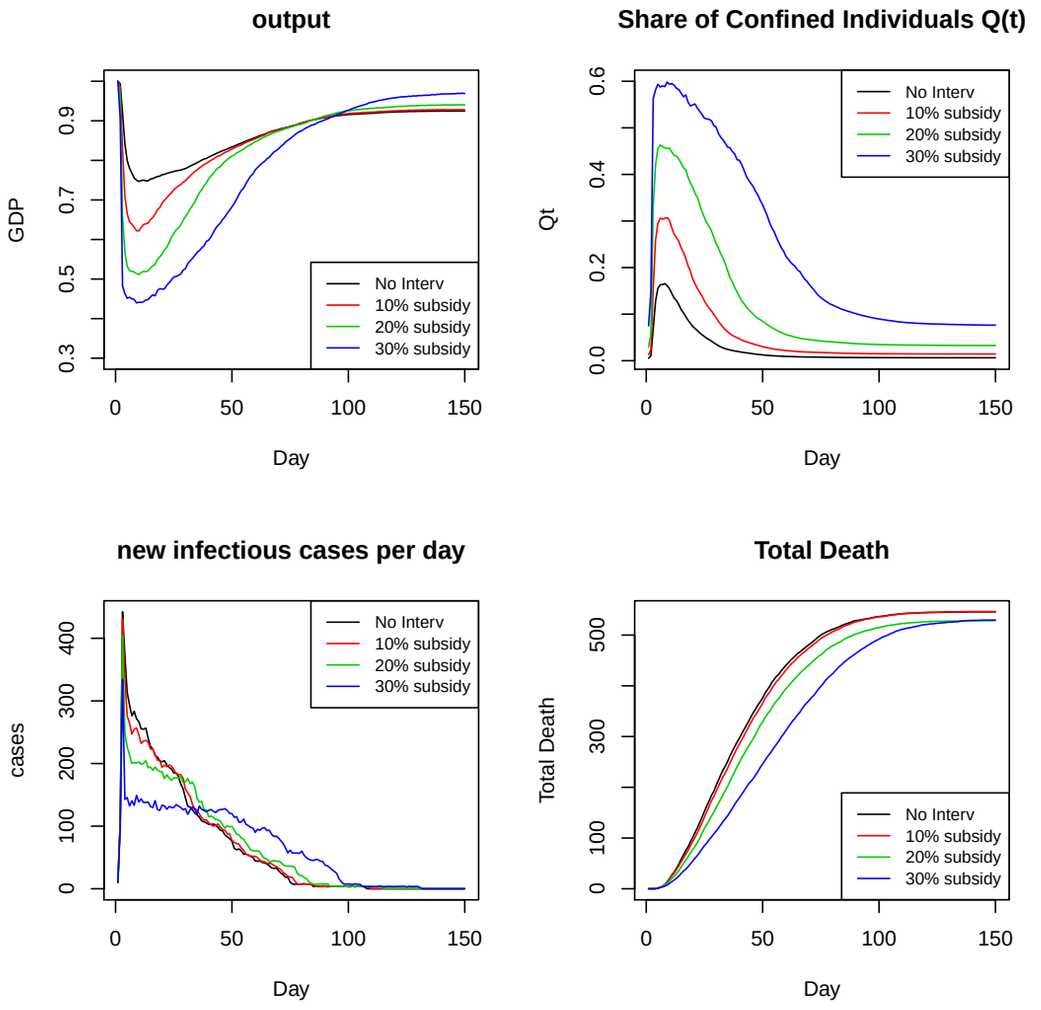

Figure 7: Experiment 4: Testing. Differences with Benchmark

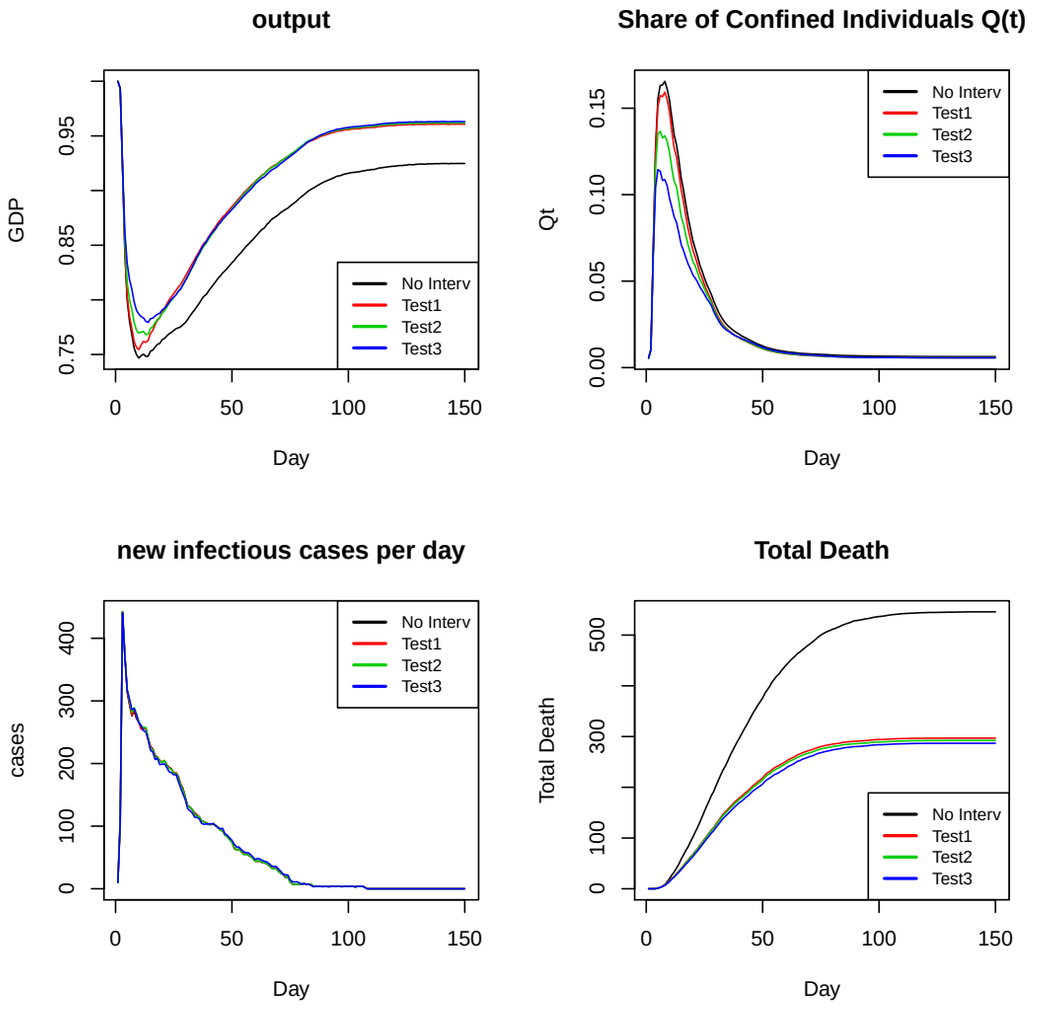

Figure 8: Experiment 5: Network Structures. Differences with Benchmark

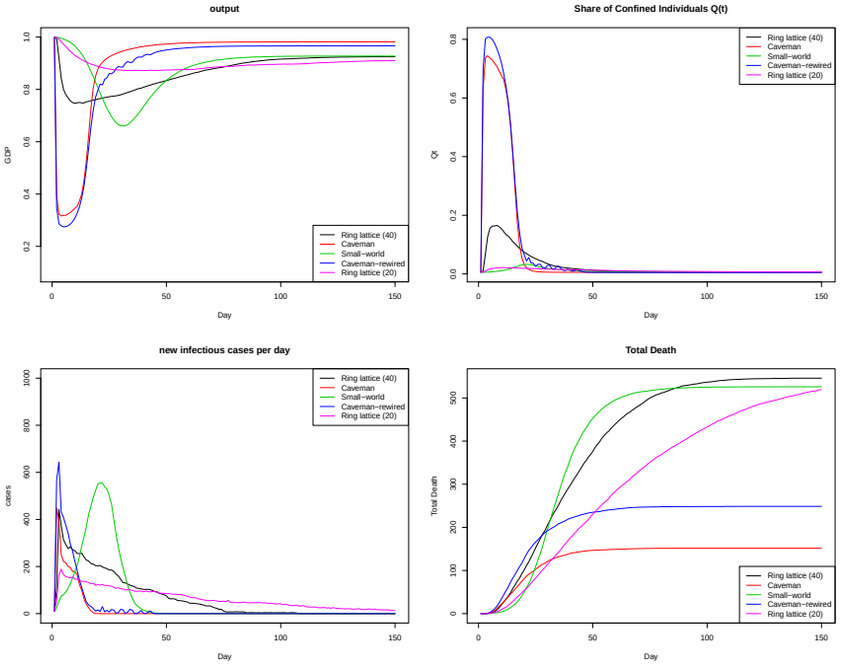

from home is smaller. Although eventually all population gets infected, this occurs much more slowly. Output also gets a much smaller hit.

The caveman network (red curve) shows also different diffusion dynamics than the benchmark. Since individuals are strongly connected within the local clusters, the virus spreads very swiftly in the beginning. But at the same time, there is also a much stronger response to work from home, which then leads to a slower spread of the virus compared to the benchmark case. In fact, the steady state reveals that less than 30% of the population gets infected before the virus dies off. The economy is hit very hard in the beginning given the large share of workers work from home, but it picks back up much faster and since the death toll is much smaller, the steady state of the output settles at a better position after the pandemic.

The variant of the caveman network, with a rewiring probability of 0.5 (blue curve) depicts a similar dynamics. Being more connected compared to the vanilla caveman network, it takes longer before the virus settles at a slow growth state and the total population gets infected is about 45%. The response in confinement decision as well as the output trajectory behaves quite similarly.

4. CONCLUSIONS

In this paper, we present a framework that combines an epidemiological model of COVID-19 diffusion with a dynamic game of network production and social interactions. The model can be used to study individuals' mobility, working and consumption decisions during the epidemic. Simulations based on a simple calibrated version of the model illustrate the potential of the framework as a tool to evaluate economic and health impacts of public policies such as subsidies to home work and testing. The simulations also showcase the important roles of behavioral responses and the structure of production networks for the dynamics of virus diffusion and output.

REFERENCES

ACEMOGLU, D., K. BIMPIKIS, AND A. OZDAGLAR (2014): "Dynamics of information exchange in endogenous social networks," *Theoretical Economics*, 9(1), 41–97.

ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): "A multi-risk SIR model with optimally targeted lockdown," Discussion paper, National Bureau of Economic Research.

ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian learning in social networks," *The Review of Economic Studies*, 78(4), 1201–1236.

ADDA, J. (2016): "Economic activity and the spread of viral diseases: Evidence from high frequency data," *The Quarterly Journal of Economics*, 131(2), 891–941.

ALVAREZ, F., D. ARGENTE, AND F. LIPPI (2020): "A Simple Planning Problem for COVID-19 Lockdown," Discussion paper.

AULD, M. C. (2003): "Choices, beliefs, and infectious disease dynamics," *Journal of Health Economics*, 22(3), 361–377.

AZZIMONTI, M., A. FOGLI, F. PERRI, AND M. PONDER (2020): "Pandemic control in econ-epi networks," Discussion paper, National Bureau of Economic Research.

BALA, V., AND S. GOYAL (1998): "Learning from neighbours," *The Review of Economic Studies*, 65(3), 595–621.

BARROT, J.-N., B. GRASSI, AND J. SAUVAGNAT (2020): "Sectoral effects of social distancing," *Available at SSRN*.

BERGER, D. W., K. F. HERKENHOFF, AND S. MONGEY (2020): "An seir infectious disease model with testing and conditional quarantine," Discussion paper, National Bureau of Economic Research.

BOERI, T., A. CAIUMI, AND M. PACCAGNELLA (2020): "Work versus safety," *COVID Economics*, (2).

CHAN, T. Y., B. H. HAMILTON, AND N. W. PAPAGEORGE (2016): "Health, risky behaviour and the value of medical innovation for infectious disease," *The Review of Economic Studies*, 83(4), 1465–1510.

EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2020): "The macroeconomics of epidemics," Discussion paper, National Bureau of Economic Research.

FARBOODI, M., G. JAROSCH, AND R. SHIMER (2020): "Internal and external effects of social distancing in a pandemic," Discussion paper, National Bureau of Economic Research.

GEOFFARD, P.-Y., AND T. PHILIPSON (1996): "Rational epidemics and their public control," *International Economic Review*, pp. 603–624.

GOLUB, B., AND M. O. JACKSON (2010): "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, 2(1), 112–49.

GREENWOOD, J., P. KIRCHER, C. SANTOS, AND M. TERTILT (2019): "An equilibrium model of the African HIV/AIDS epidemic," *Econometrica*, 87(4), 1081–1113.

HALL, R. E., C. I. JONES, AND P. J. KLENOW (2020): "Trading off consumption and covid-19 deaths," Discussion paper, National Bureau of Economic Research.

INOUE, H., AND Y. TODO (2020): "Lockdowns and Supply Chains," *Covid Economics*, (2).

KARAIVANOV, A. (2020): "A social network model of COVID-19," *Plos one*, 15(10), e0240878.

KERMACK, W. O., AND A. G. MCKENDRICK (1927): "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character*, 115(772), 700–721.

KREMER, M. (1996): "Integrating behavioral choice into epidemiological models of AIDS," *The Quarterly Journal of Economics*, 111(2), 549–573.

KUCHLER, T., D. RUSSEL, AND J. STROEBEL (2020): "The geographic spread of COVID-19 correlates with the structure of social networks as measured by Facebook," Discussion paper, National Bureau of Economic Research.

MOSSEL, E., M. MUELLER-FRANK, A. SLY, AND O. TAMUZ (2020): "Social learning equilibria," *Econometrica*, 88(3), 1235–1267.

PICHLER, S. (2015): "Sickness absence, moral hazard, and the business cycle," *Health Economics*, 24(6), 692–710.

PIGUILLEM, F., AND L. SHI (2020): "Optimal COVID-19 quarantine and testing policies," *Einaudi Institute for Economics and Finance (EIEF) Working Paper, No. 20/04*.