Epp, Markus; Jäger, Marius

Conference Paper

Network Exposure in the Propagation of the COVID-19 Pandemic

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2021: Climate Economics

Provided in Cooperation with:
Verein für Socialpolitik / German Economic Association

Suggested Citation: Epp, Markus; Jäger, Marius (2021) : Network Exposure in the Propagation of the COVID-19 Pandemic, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2021: Climate Economics, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg

This Version is available at:
https://hdl.handle.net/10419/242465

Network Exposure in the Propagation of the COVID-19 Pandemic

Markus Epp[†], Marius Jäger[‡]

September 10, 2021

Abstract. This paper develops a parsimonious model of individual exposure and of public health policy that can be used to study the evolution of an epidemic and the optimal use of lockdown policies and other non-pharmaceutical interventions. At the heart of individual exposure choices are the trade-offs between private utility derived from exposure and the risks of infection. Agents' utilities from private exposure are amplified by network benefits from social and economic activities, which amplifies the infection externality and constitutes a role for coordination policies due to multiple equilibria. We find that lockdowns do not prevail endogenously in equilibrium and that containment policy can significantly enhance welfare. Unpopularity of such contact restrictions imply that policy might face time inconsistency problems.

Keywords. Pandemics; Endogenous exposure; SIR-macro; Network effects; Externality; Time inconsistency

JEL classification. E52; E58; E61; E63

[†] Albert-Ludwigs Universität Freiburg, Rempartstr. 16, 79085 Freiburg, markus.epp@vwl.uni-freiburg.de, +49 761-203-2117.

[‡] ibid., marius.jäger@vwl.uni-freiburg.de, +49 761-203-2329

1 Introduction

When COVID-19 started spreading outside of China in the end of March 2020, countries around the world decided to react differently to the spreading of the virus. While some countries, such as South Korea, Taiwan, Germany, and many more behaved very carefully on impact other countries adopted a more or less laissez-faire policy. Sweden even became famous for its unique "Swedish way", choosing to let the pandemic surge, only protecting vulnerable groups. Meanwhile in the US, president Trump also decided to take no measures when the first wave sweeped across the country. Outcomes in terms of infection numbers per capita and, more regrettably, the overall death counts differed as much as the strictness of the imposed measures. There is an ongoing discussion about the optimal approach to containment policy. The common SIR model has generally no explicit role for policy and assumes that infection peaks are reached when satiation occurs. Thus an epidemic ends when the susceptible part of the population (which is considered exogenous) has undercut a certain threshold by getting infected, immunized or, alas, killed by the disease. This is a very unsatisfying implication of a model given the fact that many diseases, such as COVID-19, last for cycles without “burning through their own fuel” and return more infectious or become endemic. We thus develop a model that has a non-trivial role for human behavior and thus containment policy in order to study optimal containmnent policy.

Our model follows Toxvaerd (2020), Jones et al. (2020) and (Eichenbaum et al. 2020) in that it models endogenous exposure choices of individual. Our focus in the present paper lies on two aspects of individual behavior and the public response to an epidemic: firstly, we put emphasis on the complementarity of individual and public exposures in social and economic activity. Secondly, we stress that containment measures are subject to feasibility constraints: taxing individuals in a Pigouvian way to address the infection externality will generally not be possible.

We are interested in differences in the welfare implications of length, strictness and the threshold infection rate characterizing an optimal lockdown regime under a parsimonious set of assumptions. Our results imply that non-linearities in all of these characteristics exist.

In our model, the susceptible group of individuals chooses to expose themselves endogenously, facing trade-offs between the social or economic benefit from exposure on the one hand and the risk of infection on the other hand. In this sense our paper is closely related to Jones et al. (2020), Toxvaerd (2020) and Collard et al. (2020). We introduce network effects into our model by modeling agents’ utility from exposure as a function of aggregate exposure. Thus individuals exposure decisions affect other agents’ utility not only through their effect on future infection rates, but also by affecting their current period utility. Empirical evidence for the relevance of these network effects has been provided

by Chang and Velasco (2020).
Our use of dynamic programming tools is closest in spirit to Eichenbaum et al. (2020). In contrast to Collard et al. (2020), Toxvaerd (2020), agents' behavior in Eichenbaum et al. (2020) is based on explicit consumption and working decisions, whereas our paper abstracts from the particular activities that are associated with varying degrees of exposure. Moreover, we do not use a macroeconomic framework but highlight the game-theoretic interaction between susceptible agents when deciding on private exposures.
Our main findings are that network effects and endogenous exposure can lead to multiple equlibria. The indeterminacy of these equilibria constitutes a coordinating rule for policy makers since individuals will generally not coordinate themselves in the presence of externalities. The numerical approach in the present paper allows us to incorporate various state-contingencies in policies and to calculate consistent equilibrium sequences over a large number of periods. The private decisions of agents are highly tractable but closed form solutions will generally not be available so that we develop an algorithm that solves for consistent expectations and equilibrium sequences.
The paper proceeds as follows: In the next section we first introduce the SIR-model as it can be found in Jones et al. (2020). We present a behavioral model using an interactive equilibrium model in which susceptible individuals maximize their utility given the state of the pandemic as well as the behavior of other individuals, giving rise to multiple equilibria. Further, we dissect the optimal behavior of our models' subgroups. Section 3 sets the stage for our numerical analysis by presenting the calibration of the full model. Section 4 presents the simulation results. The behavioral model allows us to conduct ordinal welfare-comparisons of these equilibrium sequences, thus allowing us to compare the welfare implications of different policy scenarios. Section 5 concludes.

2 A SIR-model with Endogenous Exposure

In accordance with the available data and the goal of this paper, we consider a discrete-time variant of the SIR model with the following deviations from the well-known default by Kermack and McKendrick (1927): in decreasing importance, the model explicitly accounts for (i) an endogenously determined number of exposed individuals E_t; (ii) behavioral patterns for subgroup's exposure consistent with the fact that infected agents differ along the lines of their health state: infected agents consider themselves susceptible until they contract the disease, which is when they reveal their type (iii) an exogenous immunization rate ν_t due to vaccination and testing/quarantining; (iv) demised agents D_t are withdrawn from interaction, while recovered agents are not. This denomination of aggregate exposure E_t implies more realistic contact rates

2.1 Epidemiological structure

The standard SIR model is implied by a pooling equilibrium where susceptible agents cannot distinguish infected agents from other susceptible agents, since they behave observationally equivalently.[1] The dynamics are primarily determined by the pools of susceptibles S_t and infected I_t but removed agents R_t and demised agents D_t matter to denominate interaction of exposed and infected individuals by $N_t = S_t + I_t + R_t$ which falls short of total population $P_t = 1$ by D_t.[2]
We work with the following epidemiological system

$$N_t = S_t + I_t + R_t \tag{1}$$

$$S_{t+1} = (1 - \nu_t)S_t - \psi_t \varepsilon_t \frac{I_t S_t}{N_t} \tag{2}$$

$$I_{t+1} = (1 - \delta_t - \rho_t)I_t + \psi_t \varepsilon_t \frac{I_t S_t}{N_t} \tag{3}$$

$$R_{t+1} = R_t + \rho_t I_t + \nu_t S_t \tag{4}$$

$$D_{t+1} = D_t + \delta_t I_t \tag{5}$$

where we let infection rate at contact ψ_t, case-fatality-rate (CFR) δ_t, vaccination rate ν_t and the recovery rate ρ_t be potentially time-dependent or state-dependent processes. We endogenize the transmission of the disease by providing a behavioral aggregator-model of exposure ε_t which scales the population moments S_t and I_t.
The infection rate (i.e. the number of expected new infections normalized by the number of susceptible agents) is thus given by

$$\phi_t = \psi_t \varepsilon_t^S \varepsilon_t^I \iota_t \tag{6}$$

where $\iota_t \equiv \frac{I_t}{N_t}$ can be understood as the *adjusted prevalence*, which we will simply refer to as *incidence*. The reproduction number $\mathcal{R}_t \equiv \frac{I_{t+1}}{I_t}$ implied by the model is

$$\mathcal{R}_t = 1 - \delta_t - \rho_t + \psi_t \varepsilon_t \sigma_t \tag{7}$$

where $\sigma_t = S_t/N_t$.

2.2 Aggregation of exposure

We treat aggregate exposure ε_t as a blackbox of economic and social activity. To reduce complexity, we assume that there exists an aggregator F that maps the exposure of

[1] A separation equilibrium does not produce infections and thus extinguishes the virus in few periods.
[2] Total population is normalized to unity w.l.o.g. so that the variables S_t, I_t, R_t and D_t can be understood as population shares.

the sub-populations S_t, I_t, R_t collected in $E_t = \{\varepsilon_t^S, \varepsilon_t^I, \varepsilon_t^R\}$ into the interval $\mathcal{E}$.[3] This aggregator is parametrized by the population moments $\mathcal{P}_t = \{S_t, I_t, R_t, D_t\}$. In short, we require

$$\varepsilon_t = F(E_t; \mathcal{P}_t) \tag{8}$$

where F has the following properties:

Assumption 1 *Monotonicity: F is increasing in each sub-group's exposure*

$$\partial_{\mathcal{P}} F(E; \mathcal{P}_t) \geq 0. \tag{9}$$

where $\partial_{\mathcal{P}} F$ denotes the partial derivative of F with respect to the exposure of subgroup $\mathcal{P} = S, I, R$ for given population moments $\mathcal{P}_t$.
Our second assumption on F is that subgroup's exposures are complementary to infected agents' exposure in bringing about aggregate exposure:

Assumption 2 *Superspreading:*

$$\frac{\partial^2 F}{\partial \varepsilon^I \partial \varepsilon^{\mathcal{P}}} > 0, \mathcal{P} = S, I, R \tag{10}$$

Lastly, we assume

Assumption 3 *Boundedness: Let $\bar{E} = \{\bar{\varepsilon}, \bar{\varepsilon}, \bar{\varepsilon}\}$ and $\underline{E} = \{\underline{\varepsilon}, \underline{\varepsilon}, \underline{\varepsilon}\}$. Then*

$$F(\bar{E}; \mathcal{P}_t) \leq \bar{\varepsilon}^2, \quad F(\underline{E}; \mathcal{P}_t) \geq \underline{\varepsilon}^2, \tag{11}$$

Consistent with assumptions (A1-A3) and with full control of infection probabilities by the subgroups' exposures as well as the epidemiological model above, we will employ the most simple product-aggregator:[4]

$$F(E_t; \mathcal{P}_t) = \varepsilon_t^S \cdot \varepsilon_t^I \tag{13}$$

Note that this aggregator has the intuitive property that total new infections are determined by the sub-groups exposures multiplied by their respective expsosures:

$$\frac{\varepsilon_t^I I_t \cdot S_t \varepsilon_t^S}{N_t} \tag{14}$$

[3]The subgroups' exposures are assumed to be bounded by the interval $\mathcal{E}^{\mathcal{P}} = [\underline{\varepsilon}, \bar{\varepsilon}]$.
[4]We also ran experiments with the aggregator

$$F(E_t; \mathcal{P}_t) = \varepsilon_t^I \cdot [\iota_t \varepsilon_t^I + \sigma_t \varepsilon_t^S + (1 - \iota_t - \sigma_t)\varepsilon_t^R] \tag{12}$$

as well as the Cobb-Douglas or CES-type aggregator, with vastly equivalent qualitative and similar quantitative results.

Taken together, the epidemiological structure so far defines an operator H, i.e. a law of motion that takes the state of the epidemic at time t and exposures E_t and spits out a new state of the pandemic,

$$\mathcal{P}_{t+1} = H(E_t, \mathcal{P}_t) \tag{15}$$

The exposure model below takes this law of motion and produces consistent choices of exposures E_t.

2.3 Heterogeneity and Expectations

For the main part of our analysis, we assume that individuals consider the SIR-model's coefficients $\Pi_t = \{\nu_t, \psi_t, \delta_t, \rho_t\}$ as representative for their private parameter counterparts $\pi_{j,t} = \{\nu_{j,t}, \psi_{j,t}, \delta_{j,t}, \rho_{j,t}\}$. Moreover, we assume that that the direct period-utility loss in case of sickness, $\kappa_{j,t}$, is expected to be given by κ_t.[5]

Nevertheless, susceptible agents understand that they can scale their private probability of infection $\phi_{j,t}$ and that the parameters exogenous to their own behavior can be contingent on the aggregate state of the epidemic. Hence, we allow individual parameters in Π_t to be determined by time-invariant functional forms that only depend on the population aggregates $\mathcal{P}_t = \{S_t, I_t, R_t, D_t\}$.[6] Hence we suppress agent-specific characteristics in $\kappa_{j,t}$, and $\pi_{j,t}$ by assuming that personal expectations about these coefficients coincide across agents. We keep the agents-index j to distinguish personal and aggregate exposures as well as personal perception of coefficients $\phi_{j,t}$ from the respective counterparts on the aggregate.

Summing up, we maintain the assumption that individuals' estimates of population shares and the relevant coefficients of the epidemiological model are accurate but assume that agents mistake the population rates ν_t, ψ_t, ρ_t and δ_t as representative for their private rates $\nu_{j,t}$, $\psi_{j,t}$, $\rho_{j,t}$ and $\delta_{j,t}$.

2.4 State-contingencies

The probability of fatality δ_t is directly matched with the CFR in the data and thus we use functional forms

$$\delta_{j,t} := \delta(I_t). \tag{16}$$

These functional forms reflect that case-fatality rates are sensitive to bottlenecks in the health sector, which is captured by the number of infected agents.

[5] The misperceptions and behavioral biases in individuals' estimates are a promising avenue for future extensions but are not the focus of the present analysis.

[6] The model can be extended easily to take into account demographic and geographic features of the population as long as they preserve the upper-triangular structure of the system of the continuation values. Hence it is possible to use known distributions of age and other relevant characteristics to trace out individual parameter values.

In our baseline model, the transmission coefficient ψ_t will be held constant but the solution method allows us to introduce seasonality (or pharmaceutical protection mechanisms) into the model by varying ψ_t. In our extensions below, we study the following form of seasonality:

$$\psi_t = \psi \left[1 + \alpha \sin\left(\frac{t}{\omega} \cdot \frac{\pi}{2}\right)\right] \tag{17}$$

where ψ is the benchmark estimate employed in the epidemiological literature, α scales volatility over the course of a yearly cycle and χ is the number of periods in a year.
Our focus in the present analysis are non-pharmaceutical interventions and the behavioral response to contact restrictions in the absence of treatments and vaccines. Thus we do not model vaccinations extensively, but we let the immunization rate ν_t end the pandemic after a certain number of periods in all our simulations. It is straightforward to relax this assumption and to introduce more state-contingency (e.g. endogenous development of vaccines) to embody additional motives of precautionary behavior.

2.5 Non-pharmaceutical policies

The equilibrium will not be Pareto-optimal since agents do not take into account the (positive and negative) effects of their exposure choices on other agents' risk of infection and utility from activity. Since Pigouvian policies (e.g. a tax on exposure that reflects the social costs and benefits of exposure) that attain the first-best solution will generally not be available, we do not study such policies. Instead, we look at contact restrictions that curtain maximum exposure of agents by imposing upper bounds of individuals' exposures. In the absence of contact restrictions $\mathcal{E}_t^{\mathcal{P}} = \mathcal{E}_t = [\underline{\varepsilon}, \bar{\varepsilon}]$, $\mathcal{P} = S, I, R$. We study non-pharmaceutical interventions that are characterized by sequences of contact restrictions $\{\tau_t\}$ so that $\max \mathcal{E}_t = \tau_t$ is imposed $\forall t$. Moreover, policy cannot respond to E_t directly but controls maximum individual exposure directly and contingent only on the state of the epidemic $\mathcal{P}_t$. Hence, policies follow rules $\tau_t = \tau(\mathcal{P}_t)$. We explore two such rules in our simulations.
Let $\tau_L \geq \underline{\varepsilon}$ be the strictness of contact restrictions in lockdown and $\bar{\iota}$ a threshold incidence. We consider the following regimes:

1. *Lockdowns*: if the share of infected individuals breaches the threshold $\bar{\iota}$, individuals are restricted to exposure τ_L for n periods.[7] Formally,

$$\{\tau_k\}_{t \leq k \leq t+n} = \tau_L \;\Leftrightarrow\; \iota_t > \bar{\iota}, \tag{18}$$

and

$$\tau_t = \bar{\varepsilon} \;\Leftrightarrow\; \iota_t \leq \bar{\iota}. \tag{19}$$

[7] Note that the lockdown-rule renews contact restrictions automatically if incidence ι_t does not decline while the lockdown is in action.

2. *Gradual contact restrictions*: exposure choices are restricted according to the current prevalence of the disease. Formally,

$$\tau_t = \zeta_t \underline{\varepsilon} + (1 - \zeta_t)\bar{\varepsilon} \tag{20}$$

where

$$\zeta_t = \min\{\iota_t/\bar{\iota}, 1\} \tag{21}$$

In addition, each regime prescribes quarantining for infected agents, which (imperfectly) truncates the choice set of infected agents to $\mathcal{E}_t^I = \{\underline{\varepsilon}, \bar{\varepsilon}_t^I\}$, $\bar{\varepsilon}_t^I = q\underline{\varepsilon} + (1-q)\tau_t$ and where $q \in [0,1]$ capture the effectiveness q of quarantining rules and containment strategies.

2.6 Preferences

We model agents' exposure choices by considering the dynamic optimization schemes of susceptible, infected and removed agents which are living inside their respective continuums of size S_t, I_t and R_t.
The recursive system of continuation values is thus given by

$$V_{j,t}^S = v(\varepsilon_{j,t}^S, \varepsilon_t) + \beta\{\nu_t V_{j,t+1}^R + (1-\nu_t)[(1-\phi_{j,t})V_{j,t+1}^S + \phi_{j,t}V_{j,t+1}^I]\}, \tag{22}$$

$$V_{j,t}^I = v(\varepsilon_{j,t}^I, \varepsilon_t) + \beta(1-\delta_t)[(1-\rho)V_{j,t+1}^I + \rho V_{j,t+1}^R)] - \kappa_t, \tag{23}$$

$$V_{j,t}^R = v(\varepsilon_{j,t}^R, \varepsilon_t) + \beta V_{j,t+1}^R \tag{24}$$

where $v(\cdot)$ is a twice continuously differentiable function that governs utility derived from individual exposure $\varepsilon_{j,t}^{\mathcal{P}} \in \mathcal{E}_t^{\mathcal{P}}$, $\mathcal{P} = S, I, R$, and aggregate exposure $\varepsilon_t \in \mathcal{E}_t$. The structure of these continuation values produces the typical incentives: susceptible agents will limit their exposure in order to avoid the drop-off in utility when contracting the disease. Sick and removed face no explicit rewards for restricting their contact rate. Consistent with the aggregator specified above, susceptible individuals consider $\phi_{j,t} = \psi_t \varepsilon_{j,t} \varepsilon_t^I \iota_t$ as their individual infection probability.[8]
In addition to C^2, we make the following assumption on the functional form of $v(\cdot,\cdot)$:

Assumption 4 *Monotonicity: v is weakly increasing in both arguments and satisfies*

$$\lim_{\varepsilon_j \to \bar{\varepsilon}} \partial_1 v(\varepsilon_j, \varepsilon) = 0, \quad \lim_{\varepsilon \to \bar{\varepsilon}} \partial_1 v(\varepsilon_j, \varepsilon) \geq 0. \tag{25}$$

where $\partial_x v(\cdot,\cdot)$ denotes the partial derivative of v with respect to the x-th argument.
Our second assumption is that individual and aggregate exposure are gross complements:

[8] Note that this infection probability does not have to coincide with the average infection rate, but will do so in an equilibrium.

Assumption 5 *Supermodularity:*

$$\frac{\partial^2 v(\varepsilon_j, \varepsilon)}{\partial \varepsilon_j \partial \varepsilon} > 0. \tag{26}$$

The direct cost of sickness κ_t is assumed to be constant and will be calibrated to reflect the relative disutility from sickness.[9] We do not model long-term repercussions from infections explicitly. This allows us to not further distinguish removed agents who got vaccinated from removed agents who had to undergo the disease. This comes at some cost of realism that we are willing to accept since it points in the direction of more aggressive exposure.[10]

2.7 Partial equilibrium: optimal exposure choices

The partial equilibrium of the model is a triplet of of exposure choice sequences that we denote by $\{\varepsilon^S, \varepsilon^I, \varepsilon^R\} = \{\varepsilon_t^S, \varepsilon_t^I, \varepsilon_t^R\}_{t\geq 0}$. Since agents' are homogeneous across subgroups, we solve their representative problems by backward induction and retrieve agent-specific subscripts when needed.
Removed agents choose

$$\varepsilon_t^R = \underset{\varepsilon_{j,t}^R \in \mathcal{E}_t^R}{\arg\max} \left[v(\varepsilon_{j,t}^R, \varepsilon_t) + \beta V_{j,t+1}^R \right] \tag{27}$$

By monotonicity, optimum is achieved at maximum exposure, $\max \mathcal{E}^R \equiv \tau_t$. Thus we set

$$\varepsilon_t^R = \tau_t, \quad V_t^R = \sum_{s=t}^{\infty} \beta^{s-t} v(\tau_s, \varepsilon_s), \ \forall t \tag{28}$$

which allows us to fully trace out removed agents' utility and optimal exposure by the state-contingency of τ_t, i.e. once $\mathcal{P}_t$ is known.
In turn, exposure of infected agents is given by

$$\varepsilon_t^I = \underset{\varepsilon_{j,t}^I \in \mathcal{E}_t^I}{\arg\max} \left[v(\varepsilon_{j,t}^I, \varepsilon_t) + \beta(1-\delta_t)[(1-\rho)V_{j,t+1}^I + \rho V_{t+1}^R)] - \kappa_t \right] \tag{29}$$

Again, monotonicity implies that infected agents choose $\max \mathcal{E}_t^I = \bar{\varepsilon}_t^I$ so that both V_t^I, V_t^R and ε_t^I, ε_t^R can be traced out once $\mathcal{P}_t$ is known. Susceptible agents' exposure involves

[9] The constancy of these utility costs across agents helps us to abstract from expectations over idiosyncratic variables and does not change the fundamental decisions of susceptible agents, but only the ex-post distributions of utility losses.

[10] Personal estimates of the probability and the long-term repercussions of (health-)hazards are usually underestimated, as has been illustrated by an extensive literature (see e.g. Slovic (1987) and Tversky and Kahneman (2013) for discussions and evidence of behavior under uncertainty.

a trade-off:

$$\varepsilon_t^S = \underset{\varepsilon_{j,t}^S \in \mathcal{E}_t^S}{\arg\max} \left[v(\varepsilon_{j,t}^S, \varepsilon_t) + \beta\{\nu_t V_{j,t+1}^R + (1-\nu_t)[(1-\phi_{j,t})V_{j,t+1}^S + \phi_{j,t} V_{j,t+1}^I]\} \right], \tag{30}$$

We are interested in the network externalities so we first look at their exposure choice under the assumption that the continuation values $V_{j,t+1}^S$, $V_{j,t+1}^I$ and $V_{j,t}^R$ are known, which is w.l.o.g. since the exposure choice does not update continuation values in the partial equilibrium.

Given the continuation values, agents solve a canonical problem under uncertainty:

$$\varepsilon_t^S = \underset{\varepsilon_{j,t}^S \in \mathcal{E}_t^S}{\arg\max} \left[v(\varepsilon_{j,t}, \varepsilon_t) + \beta\{\nu_t V_{j,t+1}^R + (1-\nu_t)[(1-\phi_{j,t})V_{j,t+1}^S + \phi_{j,t} V_{j,t+1}^I]\} \right] \tag{31}$$

In contrast to removed and infected agents, susceptible agents face the trade-off between getting infected (facing the lower continuation values of infected agents) and the marginal utility of exposure.

Individuals do not take into account the impact of their choices $\varepsilon_{j,t}$ on the collective exposure ε_t nor on the overall prevalence of the disease, giving rise to the typical infection externality. In addition to the infection externality, susceptible agents are subject to higher marginal utilities given higher aggregate exposure by other agents. The infection externality implies that overall exposure is usually too high, while the second externality can both amplify and mute the infection externality in an equilibrium.

As we will illustrate in what follows, susceptible agents implicitly play (anti-)coordination games with other susceptible agents when finding their optimal exposure. To see this, note that their first-order conditions for interior solutions are given by

$$\frac{\partial v(\varepsilon_{j,t}^S, \varepsilon_t)}{\partial \varepsilon_{j,t}^S} = \beta(1-\nu_t)\psi_t \varepsilon_t^I \iota_t \gamma_{j,t} \tag{32}$$

i.e. susceptible agents balance the benefits of increased exposure with the expected cost of getting infected which is given by the differences in further continuation values defined by:

$$\gamma_{j,t} = V_{j,t+1}^S - V_{j,t+1}^I \tag{33}$$

A *symmetric* pooling equilibrium $\gamma_{j,t} = \gamma_t$ immediately implies $\varepsilon_{j,t}^S = \varepsilon_t^S$ and thus that the solution to the interior optimality condition in t defined by (32) plus the boundary restrictions imposed by $\varepsilon_t^S \in \mathcal{E}_t^S$ generate a correspondence of best-response function of susceptible exposure choices to other susceptible agents' exposures, given their perceived infection probability:

$$\varepsilon_{j,t}^S = BR(\varepsilon_t^S; \phi_{j,t}, \gamma_{j,t}) \tag{34}$$

Let $\hat{\varepsilon}^S_{j,t}$ be the solution to (32). Then (34) has the functional form

$$BR(\varepsilon^S_t; \phi_{j,t}, \gamma_{j,t}) = \min\{\bar{\varepsilon}, \max[\hat{\varepsilon}^S_{j,t}, \underline{\varepsilon}]\} \tag{35}$$

2.8 Equilibrium

An equilibrium of the epidemiological system under endogenous exposure is given by the sequences of optimal exposures $E = \{\varepsilon^S_t, \varepsilon^I_t, \varepsilon^R_t\}_{t\geq 0}$ and the consistent evolution of the population aggregates $\mathcal{P} = \{S_t, I_t, R_t, D_t\}_{t\geq 0}$. Each equilibrium can be then be indexed by sequences of the parameters and time-variant functions in $\Pi_t = \{\nu_t, \psi_t, \rho_t\}$ as well as the starting values $\mathcal{P}_0$. Each such equilibrium is associated with consistent sequences of value functions of susceptible, infected and susceptibles $\mathcal{V} = \{V^S_t, V^I_t, V^R_t\}_{t\geq 0}$, policy parameters $\tau_t = \tau(\mathcal{P}_t)$ and case-fatality rates $\delta_t = \delta(\mathcal{P}_t)$.
Generally, the collection of C^2 best-response functions (34) are continuous and time dependent maps from $\mathcal{E}^S_t$ into itself, $\mathcal{E}^S_t \mapsto BR_t(\mathcal{E}^S_t)$, so that standard definition of local stability and the associated fixed point theorems apply within periods.[11] The set of strategies of agents is restricted by the non-repeated nature of interaction between susceptible agents, hence they cannot commit to sequences of exposure choices $\{\varepsilon^S_t\}$ and seek equilibria in strategies that are characterized by period-per-period actions.
We use a Nash-equilibrium concept that involves two requirements: (i) equilibria are robust to perturbations of public exposure, i.e. they are stable under aggregator F (ii) perceived infection probabilities and aggregate infection rates coincide. (iii) agents' exposure choices are mutual best-responses.
Formally, each equilibrium is thus characterized by

1. The law of motion of the epidemiological model

$$\mathcal{P}_{t+1} = H(E_t, \mathcal{P}_t) \tag{36}$$

2. Equivalence of perceived and actual infection probabilities

$$\phi_{j,t} = \phi_t = \phi(E_t, \mathcal{P}_t) \tag{37}$$

3. Susceptible agents play their mutual best-responses, while removed and infected agents maximize their exposures:

$$\varepsilon^R_t = \tau_t, \quad \varepsilon^I_t = \bar{\varepsilon}^I_t, \quad \varepsilon^S_t = BR(\varepsilon^S_t, \phi_{j,t}, \gamma_{j,t}), \ \forall t. \tag{38}$$

[11] Since the best-response functions map from a compact subset of Euclidean space $\mathbb{R}^n$ into itself, the set of candidate equilibria is non-empty according to Brouwer's fixed point theorem.

Since we will focus on symmetric equilibria where agents have equivalent perceptions of $\phi_{j,t}$ and their expected continuation values coincide ($\gamma_{j,t} = \gamma_t$), we will suppress agent specific arguments for the analysis below and use shorthand

$$BR_t(\varepsilon_t^S) \equiv BR(\varepsilon_t^S, \phi_{j,t}, \gamma_{j,t}) \tag{39}$$

for notational convenience.

2.9 Multiplicity of laissez-faire equilibrium: the case for coordination

The decentralized equilibrium in the presence of network externalities might not be unique. This observation implies a case for coordinating susceptible agents on a more desirable equilibrium. To see this consider the following example utility function:

$$v(\varepsilon_{j,t}, \varepsilon_t) = \upsilon \cdot \left[(\varepsilon_{j,t} - \underline{\varepsilon})\varepsilon_t\right]^{1-\theta} \tag{40}$$

where υ is a scaling parameter and $\theta \in [0, 1)$ governs the strength of the network effect. For notational convenience, let

$$\varphi_t := \frac{\upsilon(1-\theta)}{\beta(1-\nu_t)\varepsilon_t^I \psi_t \iota_t}. \tag{41}$$

Optimality condition (32) determines susceptible agents' best-response exposure in interior solutions as

$$\hat{\varepsilon}_{j,t}^S = \underline{\varepsilon} + \left[\frac{\varphi_t \varepsilon_t^{1-\theta}}{\gamma_{j,t}}\right]^{1/\theta} \tag{42}$$

which for symmetric equilibria $\gamma_{j,t} = \gamma_t$ conveniently produces a version of the best-response function (34):

$$BR_t(\varepsilon_t^S) = \min\left\{\bar{\varepsilon}, \underline{\varepsilon} + \left(\frac{\varphi_t}{\gamma_t}\right)^{1/\theta} \varepsilon_t^{(1-\theta)/\theta}\right\} \tag{43}$$

where aggregator $\varepsilon_t = \varepsilon_t^S \varepsilon_t^I$ implies that mutual best-responses $BR_t(\varepsilon_t^S) = \varepsilon_t^S$ can result in global pooling equilibria of the type $\varepsilon_{j,t}^S = \bar{\varepsilon}, \forall j$ or, with sufficient curvature, $\varepsilon_{j,t}^S = \check{\varepsilon}_t^S,\ \forall j$ with $\check{\varepsilon}_t^S$ being determined by the roots of the polynomial equation

$$\underline{\varepsilon} + \left(\frac{\varphi_t}{\gamma_t}\right)^{1/\theta} (\check{\varepsilon}_t^S \varepsilon_t^I)^{(1-\theta)/\theta} - \check{\varepsilon}_t^S = 0 \tag{44}$$

Note that the slope of the best-response function (43) is monotonically increasing and strictly smaller than unity for exactly one value $\varepsilon_t := \check{\varepsilon}_t$ and strictly larger else. Thus the

equation (44) has at most two roots, which we call ε_t^L and ε_t^H, with $\varepsilon_t^L < \varepsilon_t^H$. These rest points, if in $\mathcal{E}_t^S$, are associated with derivatives smaller and larger than unity in ε_t^S-space, respectively, so that ε_t^L is generally stable and ε_t^H is unstable. Except for the case when $\varepsilon_t^H = \bar{\varepsilon}$, $\bar{\varepsilon}$ is only a candidate for an equilibrium when desired exposure of all agents is larger than the upper bound of $\mathcal{E}_t^S$ so that $\bar{\varepsilon}$ is a stable equilibrium.

Summing up, there always exists one stable equilibrium, which might not be unique. The two stable equilibria are ε_t^L and $\bar{\varepsilon}$, which coexist if and only if $\varepsilon_t^L < \bar{\varepsilon}$. We refer to ε_t^L as the (decentralized) *lockdown* equilibrium, whereas $\bar{\varepsilon}$ indexes a *business-as-usual* reference. Due to the infection externality, the high-exposure equilibrium will generally not be Pareto-optimal and therefore constitutes a coordination role for a public health policy in addition to its role in addressing the infection externality. Such an agency faces both the task to coordinate agents on the more desirable equilibrium, and, in case of a unique equilibrium, to enforce the optimal level of aggregate exposure in the business-as-usual benchmark by restricting $\mathcal{E}_t^S$ accordingly. We show in our simulations below that lockdown policies can help to prime susceptible agents on the (often more desirable) lockdown equilibrium.

3 Calibration

Our calibration strategy involves three steps: (i) calibrating epidemiological parameters and functional forms for ψ_t, δ_t using estimates from the literature; (ii) calibrating model moments to ensure vaccination rate ν_t fully eradicates the epidemic after at least $T = 300$ periods;, sensitivity of CFR δ_t to health-sector bottlenecks; (iii) calibrating the cost of sickness to match a loss in continuation values consistent with model statistic $\mathcal{P}_t$ close to observed moments in the data.
All parameters used for the numerical simulations below can be found in Table 1.

3.1 Epidemiological parameters

We use the following values of the epidemiological parameters in our simulations below:
In our calibration of ψ we use the estimated values provided in Lounis and Bagal (2020). In their analysis, using the early stages of the pandemic in Algeria Lounis and Bagal (2020) estimate ψ around 0.35, We are aware of the fact, that in the standard SIR-framework (Kermack and McKendrick (1927)), contact rates and contact transmission rates are hard to disentangle, due to endogeneities as our model also suggests.
Our calibration of the baseline death rate, follows the the predicted deat rate in the absence of any healthcare bottlenecks provided by Rossman et al. (2021).
The calibration of the weekly recovery rate ρ follows the assumption that patients recover after an average of two weeks, this is in line with most of the literature such as Baud et al. (2020) and Lounis and Bagal (2020). Thus we set our weekly recovery rate to $\rho = 0.28$. It also provides the nice advantage that infection numbers in any point in time are comparable to observed 14-day incidences. Note that we do not take heterogeinities in the population into account. However this could be done in future research.
The immunization rateν is calibrated such that herd immunity$(\sigma_h : I_{t+1} < I_t \forall \varepsilon)$ is reached after T periods. The vaccine arrives after 100 periods. This ensures the end of the pandemic at the end of our simulation. In this particular case this means $\nu = 0.005$.
Calibrating the boundaries for private exposure $\underline{\varepsilon}$ and $\bar{\varepsilon}$ is a challenging task for multiple reasons. Firstly, epidemiological data is available on the aggregate which makes it difficult to deduce individual upper and lower bounds for exposure, in particular in the presence of superspreading. Secondly, data on reproduction numbers vastly differs across countries, time and estimation strategies. Moreover, severe issues of endogeneity arise due to the heterogeneity in risk perceptions, lack of information and, in accordance with the predictions of our model, the contact restrictions and other NPIs imposed by the respective (local) governments.
We use the following tentative strategy to pick numbers of $\bar{\varepsilon}$ and $\underline{\varepsilon}$: given estimates of (i) base reproduction number $\mathcal{R}_0$, (ii) CFR $\hat{\delta}$, (iii) average duration of the disease $\hat{\rho}$ (iv) mean estimates of the infection rates at contact $\hat{\psi}$, we use $S_0 = N_0 = P_0 = 1$ to obtain

an estimate for the upper bound of aggregate infections $\bar{\varepsilon}$ from the model's equation for the starting reproduction number

$$\mathcal{R}_0 = 1 - \hat{\rho} - \hat{\delta} + \hat{\psi}\varepsilon_0 \cdot 1 \tag{45}$$

Consistent with our model above, we use the following identification assumptions to trace out the subgroups' maximum exposure in the absence of contact restrictions: (i) estimates of $\mathcal{R}_0$ coincide with $\varepsilon_0 = \varepsilon_0^S \varepsilon_0^I = \bar{\varepsilon}^2$ under the aggregator F; (ii) mean estimates of the reproduction number $\mathcal{R}_0$ are independent of seasonality in ψ_t. Using the conservative estimate of $\mathcal{R}_0 = 2$ (following Achaiah et al. (2020)), we can determine the upper bound for aggregate exposure at $\sqrt{\varepsilon_0} = 1.91$.

Table 1: **Parametrization***

Definition	Notation	Calibrated by	Value
Contact transmission rate	ψ	Lounis and Bagal (2020)	0.350
Case-fatality rate	δ	Rossman et al. (2021)	0.005
Recovery rate	ρ	Lounis and Bagal (2020)	0.280
Immunization rate	ν	Herd immunity after T	0.005
Maximum exposure	$\bar{\varepsilon}$	Maximum reproduction rate	1.916
Minimum exposure	$\underline{\varepsilon}$	Minimum reproduction rate	0.493
Seasonal volatility of	α	Liu et al. (2021)	0.5
Containment effectiveness	q	Detection rate of 10%	0.9
Triage incidence	ι^{min}	Rossman et al. (2021)	0.001
Full triage incidence	ι^{min}	Rossman et al. (2021)	0.05
Maximum CFR	δ^{max}	Ritchie et al. (2020)	0.025
Discount factor	β	Long-run market returns	$0.95^{\frac{1}{52}}$
Cost of sickness	κ	5% difference $V^S - V^I$	11.77
Concavity of utility	θ	5% difference $V^S - V^I$	0.6
Rel. weight on social utilitiy	υ	5% difference $V^S - V^I$	0.65

*Rates are on weekly basis.

3.2 Functional Forms

We calibrate four functional forms: (i) the exposure aggregator F; (ii) the utility function v; (iii) seasonality of transmission psi_t (iv) bottleneck-sensitivity of the CFR δ.

For the utility function we use

$$v(\varepsilon_{j,t}, \varepsilon_t) = \upsilon \varepsilon_{j,t}^{1-\theta} + (1-\upsilon)[(\varepsilon_{j,t} - \underline{\varepsilon})\varepsilon_t]^{1-\theta} \tag{46}$$

This form satisfies assumptions (A4) and (A5). Our calibration for parameters υ and θ. The functional form of our seasonal effect on ψ assumes that changes in temperature and solar radiation over the year largely outweigh other fluctuation. This finding is also supported by own calculations based on data from Munoz (2020). Thus our ψ_t follows a sine function centered around the average ψ, peaking in winter while generating the lowest ψ_t in the summer. We calibrate the amplitude of this function in line with the findings of Liu et al. (2021) to $\alpha = 0.5$. This means that $\psi_t = 1.5\psi$ in winter and $\psi_t = 0.5\psi$ in summer.

We calibrate the state-dependent function δ_t to match the results of Rossman et al. (2021). We assume that δ takes a quadratic form implying that, from a threshold i_{min}, δ increases until the higher estimates for case fatality rates in western Europe (**?**) are reached. Thus $\delta(\cdot)$ takes the functional form $\delta(I_t) = \min\{\delta^{max}, 408I_t^2 - 0.4I_t + \delta\}$. We are aware of the fact, that $\delta(\cdot)$ can influence the optimal lockdown measures significantly, at least on a quantitative level. However, qualitative results are not affected in our simulations.

3.3 Preferences and cost of sickness

A vast literature has used market returns to calibrate the discount factor β in agents optimization so we stick to the value that is common. Obtaining values for θ, υ and the cost of sickness κ is more intricate. In our calibration these parameters are commonly chosen to produce infection rates and death counts under the baseline scenario that are consistent with observations in countries that did not impose any contact restrictions. Hence, we calibrated κ to match a loss of 5% of the steady-state continuation values of infected agents. Moreover, we choose conservative values of $\upsilon = 0.2$ and $\theta = 0.6$ to contain the effects of curvature and network effects in the utility derived from private and aggregate exposure activity.

4 Results

In what follows, we discuss our three main scenarios: (i) a baseline scenario: no public restriction measures are in place and agents coordinate decentrally on their optimal exposure ("No measures"); (ii) "optimal lockdowns": length, duration and trigger-threshold of lockdowns maximize ex-ante utility of susceptible agents, across all such lockdown policies (iii) optimal gradual contact restrictions: exposure is constrained in dependence of aggregate incidence such that ex-ante utility of susceptible agents is maximized ("optimal gradual/smooth measures").

Figure 1 illustrates the common development of the number of susceptibles, infected, removed and demised agents for all three scenarios. Table 3 reports our findings for the optimal policy in comparison to the baseline. Interestingly, agents ex-ante prefer the

Table 2: **Welfare***

Scenario	Threshold trigger	Max. exp. restrictions	Welfare
Baseline: no measures	-	$\bar{\varepsilon} = 1.9161$	0.9703
Opt. lockdowns	0.0050	$\tau_L = 1.2045$	0.9828
Opt. gradual measures	0.0054	$\tau_t = 1.4356$	0.9795

*Welfare is given by ex-ante utility per st. st. utility of susceptible agents: V_0^S/V_{ss}^S.

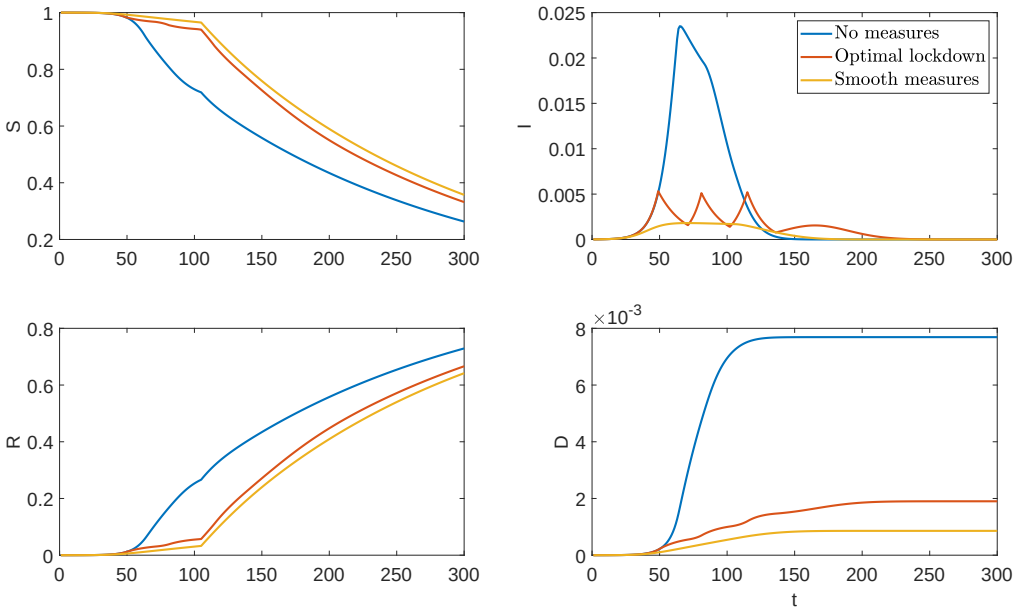

Figure 1: Evolution of epidemic for all three scenarios.

lockdown policy over both the baseline scenario without interaction and the policy regime where contact restrictions are lifted gradually. In the following we explain this surprising

result and try to derive plausible policy conclusions as well as limitations of the presented analysis.

4.1 Lockdown Scenarios

Figure 2 illustrates how optimal lockdowns constrain individual exposure. In contrast to the laissez-faire case, agents are send into lockdown significantly sooner than when agents decide to reduce their private exposure endogenously. Moreover, the exposure restrictions in this optimum require agents to limit their private exposure much stronger. Virtually all exposure in the optimal lockdown scenario is determined by public health measures, i.e. there is severe intervention necessary to address the infection externality that would prevail else. Somewhat counterintuitively, optimal lockdowns last comparably long (20

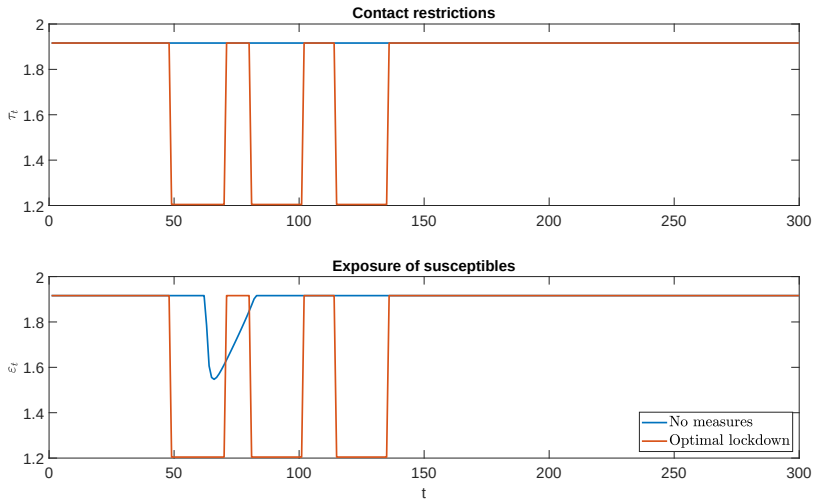

Figure 2: Contact restrictions and implied exposure by agents.

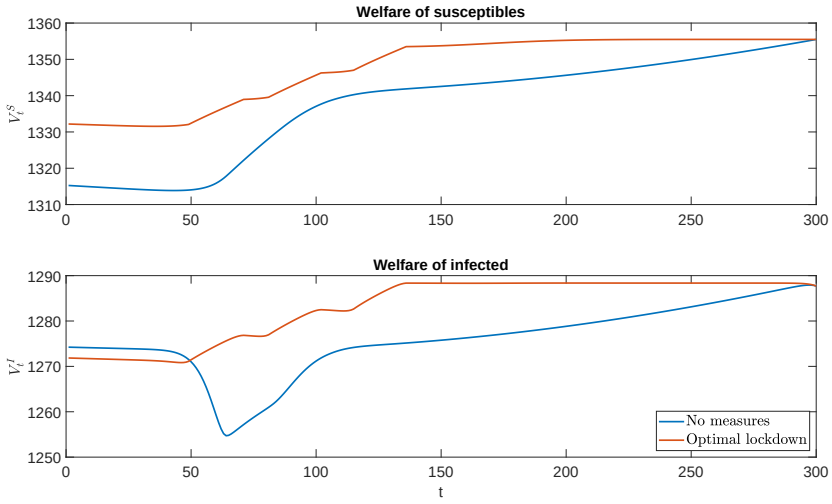

Figure 3: Present value of infected and susceptible agents' utilities.

weeks in our simulations) and are imposed starting at incidence values of 0.005 (i.e. at

500 infected per 100,000 agents). Testing for a broad range of values for the triggering threshold $\bar{\iota}$, the length of lockdowns n and the strictness of contact restrictions τ_L, we find that agents in the model ex-ante prefer more prolonged lockdowns that are associated with decisive thresholds. Appendix B lists our computed scenarios and the associated welfare implications.

Looking at the evolution of agents' welfare in figure 3 shows that much of the welfare gain can be explained by the sharp drop of welfare of infected agents in the scenario without interventions, which by means of infection probabilities also constitutes a significant part of susceptible agents' utility. Prolonged lockdowns that are scheduled early and that decisively restrict agents' exposure prevent this sharp decline in agents utility and thus "flatten the curve" until the vaccine arrives.

4.2 Gradual NPI

Figure 4 illustrates the contact restrictions and the associated exposure of susceptibles. When constrains on exposure are gradually adjusted, public health measures will start earlier to constrain individual exposures. As in the lockdown policy scenarios, agents' exposure will be constrained and thus fully determined by the public measures.

As before, agents are willing to trade the sharp decline in expected utilities during the

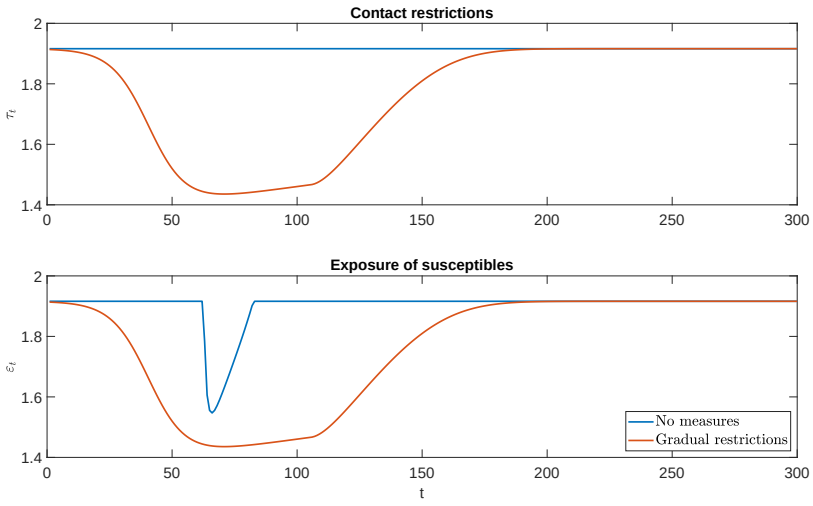

Figure 4: Contact restrictions and implied exposure by agents.

most severe weeks of the epidemic against the high exposure they are inducing in the beginning of the epidemic. The utility difference between gradual NPIs and the strict lockdown policies (which both dominate the baseline scenario without any restructions) is partially explained by the earlier sacrifices of exposure: susceptible agents dislike the gradual prescription of low exposures (when incidence is low) but are much in favor of very pro-active policies when incidence is high.

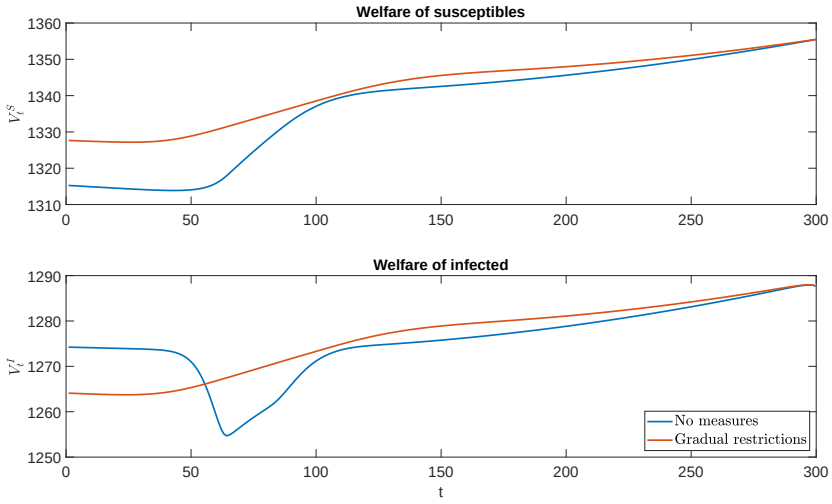

Figure 5: Present value of infected and susceptible agents' utilities.

These observations have an interesting implication: contact/exposure restrictions are unpopular *ex-ante*, i.e. when they are taking place in the near future — even when they are more effective in fighting the overall death and infection counts *ex-post*, as can be seen in figure 1. This is why political actors might be tempted to delay interventions into the future (even when they are less stringent) and will find broad support for reluctant policies while incidence is low.

5 Conclusion

In this paper we stressed the role of individual aversion to health hazards and studied the importance of network effects in bringing about endogenous exposure in the decentralized equilibrium.

Our main results are as follows: first, we confirmed the conventional wisdom that higher infection rates reduce individual's willingness to expose themselves and that exposure in equilibrium rests on coordination between individuals due to network externalities. This is also why public health measures that restrict certain social activities facilitate decentralized coordination.

Second, we evaluated two different styles of containment strategies, optimal lockdown policies and gradual contact restrictions in lockstep with rising values of incidence. Interestingly, harsh lockdowns are preferred to gradual increases in the stringency of contact restrictions. We attribute this finding to the fact that optimal lockdowns hit later and thus are discounted ex ante.

Across all scenarios, we find that agents generally prefer intervention over the laissez-faire scenario even when the available NPIs are not first-best, i.e. do not address the infection externality directly. The expected utility benefit mostly accrues to infected agents who

face a significantly lower probability of demise which amplifies the utility of susceptible agents via their continuation values.

Our results also suggest that there are important politico-economic challenges involved with conducting stringent policies when incidence rises, since such policies are the least popular when they are close to the present, making agents oppose smaller interventions upfront, i.e. when they are very effective. Focusing on the overall death toll or infection count as a measure of welfare, we find that gradual policies are more promising than strict and prolonged lockdowns.

References

Achaiah, N. C., S. B. Subbarajasetty, and R. M. Shetty (2020). R0 and re of covid-19: Can we predict when the pandemic outbreak will be contained? *Indian Journal of Critical Care Medicine: Peer-reviewed, Official Publication of Indian Society of Critical Care Medicine 24*(11), 1125.

Baud, D., X. Qi, K. Nielsen-Saines, D. Musso, L. Pomar, and G. Favre (2020). Real estimates of mortality following covid-19 infection. *The Lancet infectious diseases 20*(7), 773.

Bhattacharya, J., S. Chakraborty, and X. Yu (2021). A rational-choice model of covid-19 transmission with endogenous quarantining and two-sided prevention. *Journal of Mathematical Economics*, 102492.

Chang, R. and A. Velasco (2020). Economic policy incentives to preserve lives and livelihoods. Technical report.

Chudik, A., M. H. Pesaran, and A. Rebucci (2020). Voluntary and mandatory social distancing: Evidence on covid-19 exposure rates from chinese provinces and selected countries. Technical report.

Collard, F., C. Hellwig, T. Assenza, S. Kankanamge, M. Dupaigne, N. Werquin, and P. Fève (2020). The hammer and the dance: equilibrium and optimal policy during a pandemic crisis.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.

Gerlagh, R. (2020). Closed-form solutions for optimal social distancing in a sir model of covid-19 suppression.

Gollier, C. (2020). Pandemic economics: optimal dynamic confinement under uncertainty and learning. *The Geneva Risk and Insurance Review 45*(2), 80–93.

Horimoto, T. and Y. Kawaoka (2001). Pandemic threat posed by avian influenza a viruses. *Clinical microbiology reviews 14*(1), 129–149.

Jones, C. J., T. Philippon, and V. Venkateswaran (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical report, National Bureau of Economic Research.

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.

Koh, W. C., L. Naing, and J. Wong (2020). Estimating the impact of physical distancing measures in containing covid-19: an empirical analysis. *International Journal of Infectious Diseases 100*, 42–49.

Liu, X., J. Huang, C. Li, Y. Zhao, D. Wang, Z. Huang, and K. Yang (2021). The role of seasonality in the spread of covid-19 pandemic. *Environmental research 195*, 110874.

Lounis, M. and D. K. Bagal (2020). Estimation of sir model's parameters of covid-19 in algeria. *Bulletin of the National Research Centre 44*(1), 1–6.

Munoz, S. J. (2020). Era5-land hourly data from 1981 to present, accessed: 2021-04-01. *Copernicus Climate Change Service (C3S) Climate Data Store (CDS)*.

Phucharoen, C., N. Sangkaew, and K. Stosic (2020). The characteristics of covid-19 transmission from case to high-risk contact, a statistical analysis from contact tracing data. *EClinicalMedicine 27*, 100543.

Ritchie, H., E. Mathieu, L. Rodes-Guirao, C. Appel, C. Giattino, E. Ortiz-Ospina, J. Hasell, B. Macdonald, D. Beltekian, and M. Roser (2020). Coronavirus pandemic (covid-19). *Our World in Data*.

Rossman, H., T. Meir, J. Somer, S. Shilo, R. Gutman, A. B. Arie, E. Segal, U. Shalit, and M. Gorfine (2021). Hospital load and increased covid-19 related mortality in israel. *Nature communications 12*(1), 1–7.

Slovic, P. (1987). Perception of risk. *Science 236*(4799), 280–285.

Slovic, P. E. (2000). *The perception of risk.* Earthscan publications.

Toxvaerd, F. (2020, March). Equilibrium Social Distancing. Cambridge Working Papers in Economics 2021.

Tversky, A. and D. Kahneman (2013). Judgment under uncertainty: Heuristics and biases. In *HANDBOOK OF THE FUNDAMENTALS OF FINANCIAL DECISION MAKING: Part I*, pp. 261–268. World Scientific.

Appendix A: Algorithm

We use the following algorithm to compute equilibrium sequences of the epidemic. To initiate the algorithm, we

- choose simulation length T, a norm for convergence and a stopping rule;
- parametrize the model;
- set $\{\tau_t\} = \bar{\varepsilon}$;
- guess a sequence of $\{\varepsilon_t^S\}_{t\geq 0}$;
- compute steady values of V^S, V^I, V^R that apply starting from $T+1$.[12]

Then iterate over the following steps:

1. Simulate the epidemic using $H(\cdot)$ and the guess of $\{\varepsilon_t^S\}_{t\geq 0}$ as well as the initial values $\mathcal{P}_0$ to obtain $\{\mathcal{P}_t\}$.
2. Compute sequences of policies $\{\tau_t\}$ and other state-contingent values using $\{\mathcal{P}_t\}$.
3. Integrate backwards by starting from V_T^R, V_T^I and V_T^S and sequentially compute optimizers $\{\varepsilon_t^R, \varepsilon_t^I, \varepsilon_t^S\}$ given $\{\mathcal{P}_t\}$ and $\{\tau_t\}$.
4. Update $\{E_t\}$ and report results in case of convergence. Else, reiterate with the updated sequence $\{\varepsilon_t^S\}_{t\geq 0}$.

Table 2 describes the values chosen in our numerical exercise.

Table 2: **Algorithm details***

Definition	Notation	Value
Time spells in a year	ω	52
Duration of simulation	T	300
Updating weight	λ	0.99
Stopping rule		1e-4

*All simulations are ran in MATLAB under a standard INTEL engine.

[12]For this step, we set the st. st. infection probability at a value $\phi = 1e-6$, which cannot start the epidemic anew and is roughly the probability to get infected with influenza.

Appendix B: Lockdown Policies

Lockdown policies

Threshold trigger	Duration	Maximum exposure	Welfare
0.0050	5	1.2045	0.9820
0.0050	5	1.8449	0.9734
0.0050	10	1.2045	0.9825
0.0050	10	1.8449	0.9733
0.0050	20	1.2045	0.9828
0.0050	20	1.8449	0.9732
0.0088	5	1.2045	0.9808
0.0088	5	1.8449	0.9734
0.0088	10	1.2045	0.9814
0.0088	10	1.8449	0.9734
0.0088	20	1.2045	0.9802
0.0088	20	1.8449	0.9733
0.0125	5	1.2045	0.9791
0.0125	5	1.8449	0.9734
0.0125	10	1.2045	0.9796
0.0125	10	1.8449	0.9734
0.0125	20	1.2045	0.9814
0.0125	20	1.8449	0.9733
0.0163	5	1.2045	0.9774
0.0163	5	1.8449	0.9733
0.0163	10	1.2045	0.9796
0.0163	10	1.8449	0.9733
0.0163	20	1.2045	0.9809
0.0163	20	1.8449	0.9734
0.0200	5	1.2045	0.9753
0.0200	5	1.8449	0.9729
0.0200	10	1.2045	0.9765
0.0200	10	1.8449	0.9732
0.0200	20	1.2045	0.9793
0.0200	20	1.8449	0.9733

Appendix C: Gradual exposure restrictions

Gradual exposure restrictions

Threshold trigger	Duration	Maximum exposure	Welfare
0.0050	5	1.8449	0.9795
0.0054	5	1.8449	0.9795
0.0057	5	1.8449	0.9795
0.0061	5	1.8449	0.9795
0.0064	5	1.8449	0.9794
0.0068	5	1.8449	0.9794
0.0071	5	1.8449	0.9794
0.0075	5	1.8449	0.9794
0.0079	5	1.8449	0.9794
0.0082	5	1.8449	0.9794
0.0086	5	1.8449	0.9794
0.0089	5	1.8449	0.9793
0.0093	5	1.8449	0.9793
0.0096	5	1.8449	0.9793
0.0100	5	1.8449	0.9793

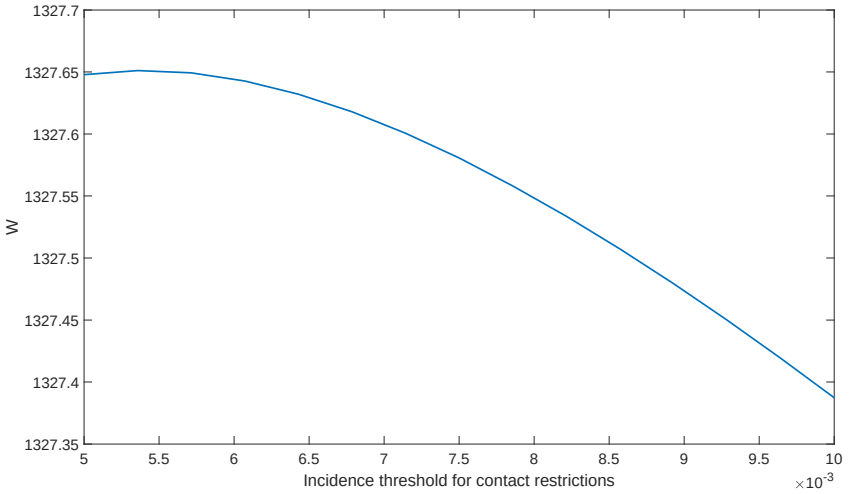

Figure 6: Welfare comparison across threshold triggers.