Review of Economic Design
https://doi.org/10.1007/s10058-024-00374-w

ORIGINAL PAPER

# The second wave

Lukasz Rachel[1]

Received: 30 June 2023 / Accepted: 26 November 2024


## Abstract
What determines whether an epidemic unfolds in multiple waves? In the absence of a vaccine, populations remain vulnerable to future outbreaks as long as susceptibility levels stay above the herd immunity threshold. The effectiveness of mitigation policies is therefore critical: a highly effective lockdown can paradoxically increase the likelihood of a second wave. This paper uses a calibrated model to study both the decentralized equilibrium and the optimal policy in a scenario where mitigation is only moderately effective. The findings show that equilibrium and optimal mitigation strategies are qualitatively similar in this case. Fiscal costs decrease the optimal length of the lockdown, narrowing the gap between equilibrium and optimal policies. We also use the model to evaluate the welfare costs of deviating from the optimal policy.

**Keywords** Epidemic · Herd immunity · Equilibrium social distancing · Optimal containment policies

**JEL Classification** E1 · I1 · H0

## 1 Introduction

Covid-19 has hit countries around the world in several successive waves. Countries that have locked up more stringently at the start of the epidemic, such as New Zealand or China, experienced rebounds of the disease in the latter stages of the pandemic. The pandemic was truly over only once the vaccine roll-out had been completed and herd immunity has been reached. Drawing on these experiences, this paper presents a framework useful for understanding epidemic trade-offs. In particular, we study the

I am grateful to the editors and an anonymous referee for helpful comments.

✉ Lukasz Rachel
l.rachel@ucl.ac.uk

1 University College London, London, UK


Springer

important worst-case scenario, in which a vaccine is not available and reductions in economic activity do little to halt the spread of the disease.[1]

Central to this analysis is the concept of *lockdown effectiveness*, which we define as a percentage reduction in the basic reproduction number $\mathcal{R}_0$ that a lockdown achieves. It is a simple and intuitive measure of how powerful a lockdown is, and it can likely be estimated in practice.

We characterize the threshold of lockdown effectiveness $\bar{\varepsilon}$ such that if a lockdown is *more* powerful than this threshold, an immediate and arbitrarily long lockdown results in an *unstable* resting point: when a lockdown is lifted, the epidemic re-emerges and there is a second wave (even if there is no change in the epidemiological characteristics of the disease). Moreover, we show that the more effective the lockdown is, the greater the severity of the second wave: that is, there is a trade-off between the degree of the initial suppression of the virus and the return of the disease further down the line. It follows that, if in response to the second wave another round of restrictions is implemented, the epidemic can again be stopped in its tracks, albeit temporarily: a third and subsequent waves inevitably emerge once the successive rounds of restrictions are lifted. The paper uses a phase diagram to visualize the disease dynamics, which makes this logic particularly clear: if restrictions are powerful at stopping the spread of the disease but if they are costly and so cannot be enforced forever (but instead come in instalments), the epidemic will feature subsequent waves. In practice such a strategy (of intermittent lockdowns after which the disease appears again and again) might lead to epidemic fatigue (perhaps lowering lockdown effectiveness), and given the prolonged duration, might result in a heightened likelihood of the emergence of new variants. Thus, considering the potential for second and subsequent waves is crucial when designing policies.

A case that is of particular interest from a risk-management point of view is the pessimistic scenario - one in which maximum mitigation measures and behaviors such as strict lockdowns are not very effective: $\varepsilon < \bar{\varepsilon}$. This might arise if the disease spreads very easily, for example through breathing. In this case even drastically reducing economic activity might do little to stop the disease spread.We analyze this scenario in detail, solving for both a decentralized equilibrium and an optimal mitigation path. Both can be fully characterized by a start and end-date of lockdown, with the maximum mitigation measures in between these two dates and no mitigation outside that interval. Throughout we focus on the benchmark case where the only feasible way to contain the disease in the long-run is to reach the herd immunity threshold. Budish (2024) studies the complementary case: mitigation is sufficiently effective to bring the effective reproduction number to below 1, thus resulting in falling infections, and the society can wait for the development of a vaccine or a cure that ultimately ends the disease.

The main result is that, in the worst-case scenario of no vaccine and ineffective lockdown, the precautions taken by agents in the decentralized equilibrium and those mandated by the planner in the social optimum are similar. Consequently, the disease trajectories do not differ much. The restrictions are first activated at a similar time

[1] Atkeson (2023) studies the interplay between vaccine development and deployment and mitigation efforts early in the pandemic.

following the emergence of a disease. The planner's optimal lockdown tends to last longer than the precautions taken by the agents in the decentralized setting.

The paper studies optimal policies of different breadths, depending on the information set of the planner. At one extreme, the planner with perfect information would restrict the behavior of the infected agents only; at the other, she may rely on broad lockdown policies that restrict the behavior of everyone in the population, regardless of their health status. We show that the main result is not qualitatively affected.[2]

An important contribution of this paper is to include fiscal costs of lockdowns in the analysis. Fiscal costs arise since a decision to mitigate the disease is associated with higher spending (for support programs) and lower tax revenues (due to curtailed economic activity). The fiscal costs of lockdown matter for the optimal duration of lockdown policies, and they constitute an externality in the decentralized equilibrium, since an individual agent that draws on the pandemic emergency package is small relative to the macroeconomy, and thus does not internalize the fact that these outlays will ultimately need to be paid for. Assuming zero fiscal costs, the optimal lockdown is around 2 weeks longer than under the baseline calibration. Because the planner restricts activity for longer than the decentralized agents, the fiscal externality brings the optimum and the equilibrium closer together. In other words, the epidemic and fiscal externalities offset to some degree.

A final contribution of this paper is to characterize social welfare outside the optimum. This is a useful contribution because in practice the authorities work with imperfect information and data, and mistakes are inevitable. Analyzing the social welfare function allows us to better understand the welfare impact of these mistakes.

*Related literature*

The literature on the macroeconomics of epidemics is, by now, vast, and this paper makes no attempt to provide a comprehensive review. Early papers by Atkeson (2020), Avery et al. (2020) and Stock (2020) provide an economist's perspective on the baseline SIR epidemiology models. In an insightful and influential paper, Eichenbaum et al. (2021) study a competitive equilibrium of a discrete time economy populated by hand-to-mouth agents whose actions affect the rates of transmission of the disease and compare it to the socially optimal mitigation policies. That paper also analyzes the worst-case scenario, in that the calibration of the effectiveness is close to the one we focus on below.[3] Rachel (2023) complements this analysis by focusing on the case when lockdown is highly effective, and shows in particular that the equilibrium and the socially optimal epidemic trajectories look very different in that case. Mitigation in the decentralized equilibrium results in a much lower rate of infection and a much longer duration of the epidemic, as individuals try to avoid getting infected. On the other hand, the planner takes advantage of the effectiveness of its tool by allowing the infection rate to climb to high levels before implementing short and maximally-stringent lockdown that brings the epidemic trajectory to the herd immunity threshold.

2 For instance, when lockdown effectiveness is low, even in the perfect information case the planner's restrictions imposed on the infected portion of the population do not halt the spread of the disease.

3 Several other papers studied individual social distancing decisions in a decentralized equilibrium and socially optimal mitigation policies (Fenichel (2013), Toxvaerd (2020), Farboodi et al. (2021), Jones et al. (2021), Alvarez et al. (2021), McAdams (2021), Piguillem and Shi (2022), Phelan and Toda (2022), Antràs et al. (2023), McAdams et al. (2023)).

Thus, there can be too much distancing in the decentralized equilibrium, relative to what is socially optimal.[4]

There are three main contributions of the present paper to this important literature. First, the paper illustrates the importance of lockdown effectiveness for determining the possibility of the second wave, and uses graphical apparatus developed in Rachel (2023) to explain why in the worst-case scenario of no vaccine and moderately effective mitigation, the equilibrium and the planner's solutions are qualitatively and quantitatively similar. Second, the model is sufficiently tractable to evaluate the social welfare function for all lockdown policies, and to study not just the optimal but also the suboptimal policies, and the gradient of welfare loss as the policymaker moves away from the optimum. Finally, the third contribution is to study the role of the fiscal cost of a lockdown numerically and demonstrate that under a reasonable calibration of model parameters, the fiscal cost can play a substantial role quantitatively, shortening the duration of optimal lockdown.[5]

*Roadmap*

The paper is structured as follows. Section 2 describes the environment and introduces the concept of lockdown effectiveness. Section 3 characterizes the decentralized equilibrium and optimal policy. Section 4 uses the calibrated model to compute quantitative results. Section 5 concludes.

## 2 Setup

As is standard in the macro-epi literature on general equilibrium models of the epidemic, the model consists of two blocks: an economic block and an epidemiological block. We begin by describing the economic block, followed by the epidemiological block.

[4] Feng (2007) provides an epidemiological perspective on quarantine and isolation policies. Acemoglu et al. (2021) consider optimal policy in a model with multiple risk groups, highlighting that targeted mitigation policies improve the trade-off between economic activity and deaths. Garibaldi et al. (2020) use insights from equilibrium search theory to characterize the equilibrium and analyze externalities. A complementary paper by Miclo et al. (2022) studies the role of the healthcare system capacity constraints in shaping the optimal mitigation policy.

[5] The analysis of the fiscal footprint of lockdowns relates to the broader strand of work on policy implications of the Covid shock. Guerrieri et al. (2022) study whether the supply shock associated with the lockdown can lead to aggregate demand deficiency and thus warrant monetary and fiscal loosening. Chang and Velasco (2020) study the feedback loops between health outcomes and a range of policies, including fiscal policy. Jordà (2020) provide a long-term historical perspective and find that the natural rate is significantly lower in the years following a pandemic. Focusing on the most recent history, Bahaj and Reis (2020) describe how the swap lines arrangements by the Fed impacted the funding markets. Kaplan et al. (2020) build a HANK model of the pandemics and evaluate a range of policies to form a pandemic policy frontier. Glover et al. (2023) consider heterogeneity along the age and workplace dimensions to point out where the major disagreements on the severity and duration of mitigation policies lie. The paper also highlights the importance of testing: the result that track and trace policies bring about significant welfare benefits relative to other lockdown measures resonates with the findings of Berger et al. (2020) who consider conditional quarantine policies and show that a given reduction in death rates can be achieved with looser mitigation measures if more information is available.

## 2.1 The economic environment

Time is continuous. At time-0, the economy is populated by a unit-measure of individuals who solve the following optimization problem:

$$\max_{n(t)\in\{0,1\}} \mathbb{E}_0 \int_0^{\infty} e^{-\rho t} u(c(t))dt \;\; \text{s.t. } c = \begin{cases} w & \text{if } n = 1 \\ h & \text{if } n = 0 \end{cases} \tag{1}$$

with $u > 0$, $u' > 0$, $u'' \leq 0$. Each period, individuals decide whether to work and behave normally ($n = 1$) or to isolate and stay at home ($n = 0$). The choice is binary, meaning individuals either work or isolate. This is without loss of generality since in the equilibria considered in this paper, the optimal actions are discrete. If individuals work, they earn after-tax wage $w$. If they stay at home their income is equal to $h$, which is composed of the government payment (share $\psi_{GOV}$), working from home income (share $\psi_{WFH}$) and home production (share $1 - \psi_{GOV} - \psi_{WFH}$). There is no saving, so consumption is equal to income in each period: $c = w$ or $h$. We assume that $w > h$, so that households choose to work in normal times.

Production technology is linear in labor with productivity $A$, markets are competitive, and there is a government that taxes labor income at a rate $\tau_n$:

$$Y(t) = AN(t) \qquad w = A(1 - \tau_n).$$

The government spends the proceeds on an exogenous amount of public goods, $\bar{G}$. I assume that prior to the epidemic, the government runs a balanced budget, and so $\bar{G} = \tau_n A$. The epidemic creates a hole in the government budget, both because it lowers tax receipts and raises spending. For simplicity, we assume that the government borrows the required funds and procures the required output directly in the international market. The government faces a constant exogenous interest rate $\bar{r}$. The tax rate $\tau_n$ is fixed; instead, the government finances its debt by levying constant lump-sum post-pandemic-tax $\tau$ every period on those who survived the virus, starting from some date $\hat{T}$ (after the epidemic has ran its course).[6] Given these assumptions, the government's intertemporal budget constraint is:

$$\begin{aligned} (\psi_{GOV} h + \tau_n (A - \psi_{WFH} h)) \int_0^{\infty} e^{-\bar{r}t} \, & [(1 - \lambda_S(t))S(t) \\ & + (1 - \lambda_I(t)) \, I(t) + (1 - \lambda_R(t))R(t)] \, dt \\ \leq \int_{\hat{T}}^{\infty} e^{-\bar{r}t} \, (S(t) + R(t)) \, \tau dt & \end{aligned} \tag{2}$$

[6] To be precise, the epidemic disappears asymptotically. We thus assume that time $\hat{T}$ corresponds to the level of susceptibility that is below the herd immunity threshold, and the level of infectiousness that is sufficiently low. We clarify the meaning of all these concepts below.

where $\lambda_i,\ i \in \{S, I, R\}$ are the shares of *active* susceptible, infected and recovered agents (denoted by $S(t), I(t), R(t)$), respectively - i.e. the agents who work. The left-hand side of (2) is the net present value of the pandemic-driven losses. The first term in parentheses is the (instantaneous) amount government pays to each person who isolates. The second term is the loss of tax income, which is the loss of usual income adjusted for the fact that government taxes income that results from working from home.[7] The integral computes the measure of individuals that have been isolating and thus have received government payouts and pay less tax. The right-hand side is the net present value of post-pandemic tax receipts.

### 2.2 The epidemiological block

The dynamics of the disease are modelled using a canonical SIR model with endogenous behavior. Let $\beta(t)$ denote the potentially time-varying disease transmission coefficient. We assume that $\beta(t)$ is a sum of two components:

$$\beta(t) = \beta_n \lambda_I(t)\lambda_S(t) + \beta_o.$$

In that setting, $\beta_n$ is the parameter guiding reducible infections, and $\beta_o$ is the share of infections that remain no matter how strong mitigation is. The epidemic dynamics are then described by a canonical SIR model:

$$\dot{S}(t) = -\beta(t)S(t)I(t) \tag{3}$$

$$\dot{I}(t) = \beta(t)S(t)I(t) - \gamma I(t) \tag{4}$$

$$\dot{R}(t) = \gamma_r I(t) \tag{5}$$

$$\dot{D}(t) = \gamma_d I(t), \tag{6}$$

with given initial conditions[8] and $\gamma = \gamma_r + \gamma_d$.

With a slight abuse of notation, let $\beta$ represent the infection rate with no behavioral response: $\beta := \beta_n + \beta_o$, and define the herd immunity threshold as $\bar{S} := \frac{\gamma}{\beta}$. Two key properties of this system are worth noting. First, every point with $I(t) = 0$ is a steady state. Second, the number of infected will grow if and only if $I(t) > 0$ and $S(t) > \bar{S}$ where

$$\bar{S} := \frac{\gamma}{\beta} = \frac{1}{\mathcal{R}_0} \tag{7}$$

[7] For simplicity, we abstract from losses that result from the fact that some citizens die from the disease. This effectively assumes that $\bar{G}$ is expressed in per capita terms.

[8] Throughout the paper I assume that $S_0 = 1 - \epsilon$, $I_0 = \epsilon$ and $R_0 = D_0 = 0$.

is the herd immunity threshold (and $\mathcal{R}_0$ is the basic reproduction number).[9] Combining these two properties, it is clear that the herd immunity threshold divides the set of steady states into two segments:

**Lemma 1** *Steady states of the system with $I = 0$ and $S > \bar{S}$ are unstable, in the sense that a small perturbation to the number of infected triggers dynamics that take the system away from that steady state. Conversely, the steady states with $I = 0$ and $S < \bar{S}$ are stable.*

***Proof*** Follows directly from equation (4) and the definition of $\bar{S}$. □

These two segments and the herd immunity threshold will play a crucial role in determining the epidemic dynamics.

**Definition 1** The economy is in *lockdown* at $t$ if $\lambda_I(t) = 0$ or $\lambda_S(t) = 0$ or both.

The relative magnitudes of $\beta_n$ and $\beta_o$ parameters determine the effectiveness of mitigation behavior and policy. We define the effective reproduction number in the usual way: $\mathcal{R}_t := S(t)\mathcal{R}_0$. Denoting with a subscript $L$ the corresponding variables under lockdown, we have the following proposition, which defines formally the lockdown effectiveness $\varepsilon$.

**Proposition 1** *Lockdown effectiveness $\varepsilon$ satisfies:*

$$\varepsilon := \frac{\mathcal{R}_0 - \mathcal{R}_0^L}{\mathcal{R}_0} = \frac{\mathcal{R}(t) - \mathcal{R}^L(t)}{\mathcal{R}(t)} = \frac{\bar{S}_L - \bar{S}}{\bar{S}_L} = \frac{\beta_n}{\beta_n + \beta_o}. \tag{9}$$

*If $\varepsilon > \bar{\varepsilon}$, an immediate lockdown leads to unstable suppression of the disease: lifting the lockdown results in a second wave of infections. The threshold $\bar{\varepsilon}$ is given by:*

$$\bar{\varepsilon} = \frac{\bar{S}}{\bar{S} - 1}\left(1 - \log \bar{S} - \frac{1}{\bar{S}}\right).$$

*If $\varepsilon > \hat{\varepsilon}$, an immediate lockdown suppresses the virus completely, preventing the epidemic. But once the lockdown is lifted the epidemic starts over and follows the no-lockdown trajectory. Threshold $\hat{\varepsilon}$ is given by:*

$$\hat{\varepsilon} = 1 - \frac{1}{\mathcal{R}_0} = 1 - \bar{S}.$$

[9] Other properties can be derived noting that the epidemic trajectory admits an analytical solution. Dividing equation (4) by (3) we obtain a first-order ODE: $\frac{dI}{dS} = -1 + \frac{\bar{S}}{S}$, which, given the initial conditions $S_0$ and $I_0$, gives:

$$I(t) = -S(t) + \bar{S}\log S(t) + I_0 + S_0 - \bar{S}\log S_0. \tag{8}$$

Taking the limit as $t \to \infty$ and noting that $I(\infty) = 0$ shows that the initial basic reproduction number and the eventual share of the population that will have encountered the virus $S(\infty)$ are tightly linked: $\mathcal{R}_0 = \frac{\log S(\infty)}{S(\infty)-1}$. Since death rate is assumed to be an exogenous constant, total deaths that result from infection are given by $D(\infty) = \frac{\gamma_d}{\gamma}(1 - S(\infty))$. Moreover, condition (7) and equation (8) imply that peak infection rate implied by a simple SIR model with exogenous transmission rate is given by $I_{max} = -\bar{S} + \bar{S}\log \bar{S} + 1$.

***Proof*** See Appendix. □

The first part of Proposition 1 introduces and calculates the key feature of a lockdown: its effectiveness ($\varepsilon$). A more effective lockdown leads to a larger reduction in the initial reproduction number. The measure of effectiveness is defined as the percentage change in $\mathcal{R}_0$ (or equivalently $\mathcal{R}(t)$ and $\bar{S}$) due to the implementation of a lockdown. The Proposition demonstrates that this measure is equivalent to quantifying the share of infectious activities curtailed during the lockdown, making it a natural and intuitive indicator of effectiveness.

The second part of the Proposition establishes the threshold at which a lockdown becomes sufficiently effective for rapid virus suppression, with the system's resting point under a full and immediate lockdown surpassing the herd immunity level $\bar{S}$. However, lifting restrictions prematurely can trigger a second wave of infections.

The third part introduces a higher threshold, $\hat{\varepsilon}$, which represents the effectiveness level required to completely halt the spread of the virus. Although halting the epidemic seems advantageous, the proposition suggests that this only delays the problem: once the lockdown is lifted, the epidemic resumes its original trajectory as if no lockdown had occurred. This underscores a crucial aspect of the model- that a temporary lockdown cannot suppress the virus if the level of susceptibility remains above $\bar{S}$.

An important implication of Proposition 1 is the trade-off between the severity of the first and second waves of the epidemic:

**Proposition 2** *Consider an immediate and arbitrarily long but finite-duration lockdown policy lasting for a single interval of time, with* $\bar{\varepsilon} < \varepsilon < \hat{\varepsilon}$. *The peak of the second wave of the epidemic in terms of the highest attained infection rate and the proportion of people who at some point contract the virus,* $1 - S(\infty)$*, are both increasing in* $\varepsilon$.

***Proof*** See Appendix. □

### 2.3 Graphical representation

The results in the propositions have a useful graphical representation. Figure 1 shows the system dynamics, with and without behavioral or policy responses, in a phase diagram developed in Rachel (2023).[10] Both panels in Fig. 1 depict the $Susceptibles - Infected$ plane. The solid lines with arrows depict disease trajectories under different assumptions about policy and behaviour. The arrows show the direction of the evolution of the system. The solid vertical lines mark the herd immunity thresholds in the different lockdown scenarios. We assume that an epidemic begins with a small seed of infection and with near full susceptibility, so that the system starts in the bottom-right corner of the diagram, with $I$ close to but above zero and $S$ close to but below 1.

Consider first the trajectory of the disease as dictated by a standard SIR model with constant infection rate $\beta$. This trajectory, and the system dynamics associated

[10] Note that drawing the trajectory of the disease in a standard, constant-$\beta$ SIR model requires only the calibration of the herd immunity threshold $\bar{S}$, which is equivalent to calibrating $\mathcal{R}_0$.

with it, are identical in the two panels of the figure. In our parametrization, the herd immunity threshold is $\frac{\gamma}{\beta} = \frac{0.778}{1.94} = 0.4$. Since the epidemic starts with the level of susceptibility close to one, and because $\bar{S} < 1$, the infection rate initially increases and the susceptibility rate falls: the trajectory moves in the north-west direction in the phase diagram. Eventually, the level of susceptibility reaches $\bar{S}$. Due to the dynamics of the system, the long-run steady-state level of susceptibility $S(\infty)$ exceeds the herd immunity threshold. The difference between the two is the epidemic overshoot. In our parametrization, around 30% of the population get infected after the herd immunity threshold is reached.

The size of the arrows reflects the system's dynamics described in Eqs. (3) and (4), indicating that the system's speed is greater the larger are the infection rate and the susceptibility rate.

How does a lockdown affect the system's dynamics? During a temporary lockdown, the transmission rate $\beta$ decreases, which temporarily raises the herd immunity threshold. Thus, for the duration of the lockdown, the vertical herd immunity schedule shifts to the right (we denote the temporarily raised herd immunity threshold under lockdown with $\bar{S}_L$). The dynamics of the system obey the same logic as before, but with reference to the new, higher herd immunity threshold. For example, under lockdown, the infection rate rises as long as $S > \bar{S}_L$, and peaks at $S = \bar{S}_L$.

We can now better understand the notion of lockdown effectiveness and the thresholds characterized in Proposition 1. The key observation is that the size of the shift in the herd immunity threshold is directly proportional to lockdown effectiveness: the more effective the lockdown, the larger the rightward shift.

The Figure shows two key thresholds for lockdown effectiveness, as characterized in Proposition 1. The left panel shows the dynamics of the disease under lockdown if lockdown effectiveness is $\varepsilon = \bar{\varepsilon}$, where the threshold is characterized analytically in the proposition. This value of lockdown effectiveness means that if the lockdown is implemented for long enough, the epidemic dynamics reach the resting point $S(\infty)$ that is equal to the herd immunity threshold $\bar{S}$. The right panel shows what happens when lockdown is even more effective, with $\varepsilon = \hat{\varepsilon}$. In this case, the herd immunity threshold under lockdown increases to $\bar{S}_L = 1$. Consequently, the disease stops in its tracks: the infection rate is declining monotonically from the initial seed of infection.

The key observation is that the two lockdown trajectories have very different implications when the lockdown ends. Once the lockdown is lifted, the herd immunity threshold for the system's dynamics returns to $\bar{S} = \frac{\gamma}{\beta}$, 0.4 in our calibration. In the panel on the left, the dynamics under lockdown reach a stable steady state region, with the level of susceptibility smaller or (in this case) equal to $\bar{S}$. There is no second wave of infection when lockdown is lifted, and the epidemic ends. The situation is very different in the right panel, where the highly effective lockdown kept the system close to the initial starting point. Consequently, when the lockdown is lifted, the presence of infected individuals leads to a renewed increase in the infection rate. We observe a large second wave of infections as a result. Conditional on no further restrictions, the system's trajectory is similar to the path generated by the constant-$\beta$ SIR model.

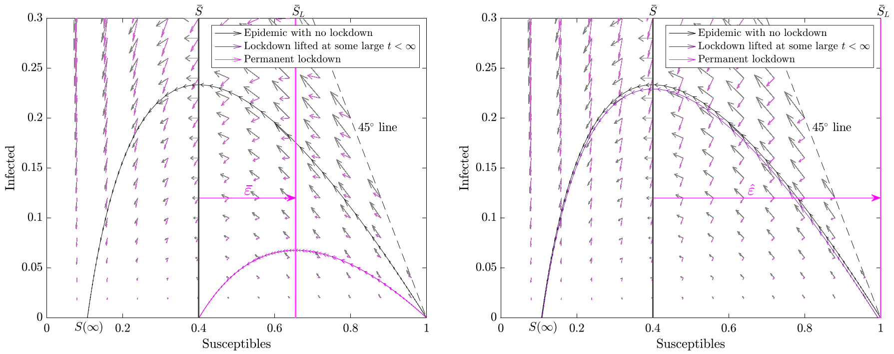


**Fig. 1** Lockdown effectiveness thresholds

Proposition 2 can now be visualized more easily. A more effective lockdown leaves susceptibility at a higher level, resulting in a larger rise in infections and cumulative deaths during the second wave after the lockdown is lifted.

### 2.4 Policy implications

The above discussion has clear implications for policy. In this simple benchmark model, the goal of policy is to manage the epidemic while the susceptible population remains above the herd immunity threshold. The objective is to reach herd immunity with the number of infectious individuals at or near zero. An example of such a trajectory is shown as the pink parabola in the left side of Fig. 1. Note that this trajectory hits the x-axis right at the herd immunity threshold, leading to no "overshooting". In contrast, as illustrated in the right panel of Fig. 1, an extremely tight lockdown that is eventually lifted does not significantly reduce cumulative deaths, relative to the naive model. This is because the pandemic simply restarts, following essentially the same trajectory as it would have done in the absence of any intervention.

In the numerical analysis below, we confirm this intuition and compare it to the decentralized equilibrium outcomes.

### 2.5 What is missing from the baseline framework

Before analyzing equilibrium and optimal lockdown in the next section, it is important to reflect on the key features of this baseline model.

First, it is important to note that the result that suppressing the virus with a finite-length lockdown is impossible is a consequence of the SIR model's mathematical structure. In reality, it may be possible to reduce the number of infections enough to control the disease until the very last patient recovers. Pollinger (2023) analyzes the case of optimal suppression. Second, the model assumes a constant, exogenous death rate, excluding any feedback between the number of currently infected individuals and the probability of death or recovery. In practice, death rate might depend on the number of infected because of healthcare systems' capacity constraints. Indeed, the rationale for ’flattening the curve’ was to avoid overloading ICU capacity and increasing deaths. For the analysis of this case, see ?. Third, the model ignores the possibility of vaccine development. Instead the only way to put an end to the virus is to achieve at least the level of herd immunity. Several papers in the literature tackle this issue (see. for example, Atkeson (2020), Phelan and Toda (2022), Gans (2023) and Avery et al. (2024)). Fourth, the model assumes that the immunity gained by the recovered is permanent. When immunity is temporary, the disease becomes endemic (Giannitsarou et al. 2021). In terms of the dynamics in the phase diagram, the trajectory of a naive model with no behavioral response can feature a spiral - several consecutive waves with ever smaller magnitudes, and asymptotic convergence to a steady state at the level of susceptibility equal to the herd immunity threshold. Avery et al. (2024) analyze the possibility of getting infected more than once, together with endogenous decisions to get a vaccine.

## 3 Equilibrium and optimal lockdown with $\varepsilon < \bar{\varepsilon}$

We now proceed to describe the laissez-faire equilibrium and optimal policy under the assumption that lockdown is not sufficiently effective to reduce the number of infected: $\varepsilon < \bar{\varepsilon}$.[11] This case is particularly interesting as it represents a worst-case scenario: we assume that not only there is no vaccine, but also the mitigation options available to the agents and to the planner are relatively ineffective.

### 3.1 Competitive equilibrium

The formal definition of the pure strategy equilibrium is:

**Definition 2** A *perfect-foresight competitive equilibrium in pure strategies* is a sequence of macro variables $Y$ and $C$, sequence of epidemic variables $S, I, R, D$, sequence of labor supply choices $\{n_i\}_{i\in\{S,I,R\}}$ and the associated sequence of lockdown indicators $\{\lambda^i(t)\}_{i\in\{S,I,R\}} \in \{0, 1\}$, the level of post-pandemic tax $\tau$ such that: (i) households maximize their expected lifetime utility at time-0 taking the trajectory of the epidemic, behavior of other individuals, wages, government spending and taxes as given; (ii) government adjusts post-pandemic tax rate $\tau$ to satisfy demand for transfers, meet its spending commitments and satisfy its intertemporal budget constraint; (iii) the trajectory of the epidemic is consistent with the individual lockdown decisions; (iv) markets clear.

[11] In Rachel (2023), I solve the optimal control problems of a susceptible individual and the planner, focusing on the opposite case of $\varepsilon > \hat{\varepsilon}$. In that scenario, the epidemic trajectory in the decentralized equilibrium features, at least for some interval of time, an interior level of mitigation, sufficient to keep the effective reproduction number below 1. As a result, the decentralized equilibrium results in a relatively low and declining rate of infections that is optimal given the private trade-off between health benefits and economic costs. This contrasts with the scenario that we focus on in this paper: when mitigation is not very effective, agents choose the maximum degree of mitigation (a corner solution) because the benefits of mitigation outweigh the costs. Socially optimal restrictions in the case of highly effective lockdown steer the epidemic towards the herd immunity threshold as fast as possible, which means that the epidemic trajectory reaches high infection rates before restrictions are activated. Such high-infection-rate strategy is not optimal when the lockdown is not effective, since it would result in a large epidemic overshoot.

Such equilibrium exists when $\varepsilon < \bar{\varepsilon}$.[12] To characterize the equilibrium, note first that the infected and the recovered individuals do not lock down voluntarily.[13] The value functions of the infected and the recovered individuals thus satisfy the following Bellman equations:

$$\rho W^I = u(w) + \pi_r \left( W^R - W^I \right) + \pi_d (0 - W^I) + \dot{W}^I \tag{10}$$

$$\rho W^R = u(w) + \dot{W}^R \tag{11}$$

with boundary conditions $W^I(\hat{T}) = W^R(\hat{T}) = \frac{u(w-\tau)}{\rho}$.[14]

A susceptible individual chooses the start and the end-date of the lockdown period, $\{T_0, T_1\}$, to maximize her expected utility:

$$\max_{\{T_0 \geq 0, T_1 \geq T_0\}} \left\{ \begin{array}{l} \int_0^{T_0} e^{-\rho t - \int_0^t \pi_W(s) ds} \left( u(w) + \pi_W(t) W^I \right) dt \\ + e^{-\int_0^{T_0} \pi_W(s) ds} \int_{T_0}^{T_1} e^{-\rho t - \int_{T_0}^t \pi_L(s) ds} \left( u(h) + \pi_H(t) W^I \right) dt \\ + e^{-\int_0^{T_0} \pi_W(s) ds - \int_{T_0}^{T_1} \pi_H(s) ds} \int_{T_1}^{\infty} e^{-\rho t - \int_{T_1}^t \pi_W(s) ds} \left( u(w) + \pi_W(t) W^I \right) dt \end{array} \right\} \tag{12}$$

subject to the time-varying infection probabilities:

$$\pi(t) = \begin{cases} \pi_W(t) = (\beta_n \lambda + \beta_o) I(t) & \text{no lockdown, } t \notin [T_0, T_1] \\ \pi_H(t) = \beta_o I(t) & \text{in lockdown, } t \in [T_0, T_1] \end{cases}, \tag{13}$$

where $\pi_W$ denotes the infection probability when working and $\pi_H$ denotes the infection rate when staying at home. The three lines in expression (12) correspond to the utility before, during and after mitigation, respectively.

The interior optimum satisfies the necessary conditions:

$$u(w) - u(h) = \beta_n I(T_0) \left( W^S(T_0, T_1, t = T_0) - W^I \right) \tag{14}$$

$$u(w) - u(h) = \beta_n I(T_1) \left( W^S(T_0, T_1, t = T_1) - W^I \right) \tag{15}$$

where $W^S(T_0, T_1, t)$ is the expected lifetime utility of a susceptible individual at time $t$, given the lockdown period starting at $T_0$ and ending at $T_1$. Note that this time-varying value function depends on the start- and end- date of the lockdown, since these dates

[12] See the Appendix for a proof.

[13] Since there is no altruism, individuals care only about maximizing their expected utility. For infected or recovered individuals there is no risk of re-infection, and because $w - \tau > h$ and there is no disutility of labor, the optimal choice for these individuals is to always work.

[14] Note that the cost of dying is the foregone utility of regaining health.

will determine the epidemic trajectory (and hence infection risk facing an individual) and the period over which a susceptible individual receives the reduced utility flow.

Conditions (14) and (15) emphasize the trade-offs present when deciding on their personal social distancing strategy: individuals balance the instantaneous utility cost of lockdown with the expected health benefit expressed in utility terms. This benefit is determined by the increased probability of infection when not in lockdown, multiplied by the difference in value between becoming infected and remaining healthy.

We now establish an important result regarding the timing of social distancing in equilibrium.

**Lemma 2** *In a competitive equilibrium social distancing never starts in the initial period* $t = 0$.

***Proof*** At $t = 0$, $\pi_W(0) - \pi_H(0) = \beta_n I(0) \approx 0$ but $u(w) > u(h)$. Therefore, it is optimal to postpone lockdown until a later date. □

The equilibrium is a fixed point between the decision rule (pinned down by (14) and (15)) and the epidemic trajectory (Eqs. (3) and (4)). This fixed point is found by iterating on the pair $(T_0, T_1)$ until the resulting epidemic trajectory satisfies conditions (14) and (15).

We further characterize the decentralized equilibrium numerically below. Before doing that, we first set out the planning problem.

### 3.2 Lockdown policies

Next, we consider the decision-making process of a planner whose objective is to maximize population-wide lifetime utility. We consider the following distinct tools:

**Definition 3** We define four types of lockdown instruments as follows:

Type 1: *isolation of infected only:* government sets $\{T_0 \geq 0, T_1 \geq T_0\}$. $\lambda_I(t) = 0$ for $t \in [T_0, T_1]$, and $\lambda_I(t) = 1$ otherwise. $\lambda_S(t) = \lambda_R(t) = 1 \forall t$.

Type 2: *lockdown of susceptibles only:* government sets $\{T_0 \geq 0, T_1 \geq T_0\}$. $\lambda_S(t) = 0$ for $t \in [T_0, T_1]$, and $\lambda_S(t) = 1$ otherwise. $\lambda_I(t) = \lambda_R(t) = 1 \forall t$.

Type 3: *lockdown with immunity passports for the recovered:* government sets $\{T_0 \geq 0, T_1 \geq T_0\}$. $\lambda_S(t) = \lambda_I(t) = 0$ for $t \in [T_0, T_1]$, and $\lambda_S(t) = \lambda_I(t) = 1$ otherwise. $\lambda_R(t) = 1 \forall t$.

Type 4: *broad lockdown:* government sets $\{T_0 \geq 0, T_1 \geq T_0\}$. $\lambda_S(t) = \lambda_I(t) = \lambda_R(t) = 0$ for $t \in [T_0, T_1]$, and $\lambda_S(t) = \lambda_I(t) = \lambda_R(t) = 1$ otherwise.

These tools are ranked from the most to the least information-intensive for the planner. Type 1 represents an idealized case of a perfectly effective track-and-trace strategy, where only the infected are isolated. The remaining three tools involve broader lockdowns that restrict the behavior of larger portions of the population.

Although these tools differ in terms of the population segments they target, they all lead to similar outcomes in terms of epidemic dynamics, as we explain below.

**Lemma 3** *Consider lockdown policies of the four types defined above and of fixed timing* $\{T_0, T_1\}$. *All four policies have identical effect on the epidemic dynamics.*

***Proof*** Under all four policies, $\beta(t)$ is given by

$$\beta(t) = \begin{cases} \beta_n + \beta_o & \text{if } t \notin [T_0, T_1] \\ \beta_o & \text{if } t \in [T_0, T_1] \end{cases}.$$

□

This result might appear surprising at first but it is intuitive once we note that, as long as *either* the infected or the susceptible are locked down, the reducible infections fall to zero.

Lockdown effectiveness plays a crucial role in determining the features of optimal policy. If lockdown is highly effective, it may create the risk of a second wave. Optimal policy avoids this by timing the lockdown to ensure that the epidemic's trajectory converges to $\bar{S}$ as a resting point (see the derivation and discussion of this result in Rachel (2023)). We instead now solve for the equilibrium and optimal policies when lockdown is only moderately effective.

## 4 Quantitative analysis

### 4.1 Calibration

**Macro Parameters**

I calibrate the model to a weekly frequency, with the discount rate $\rho = 0.96^{\frac{1}{52}} - 1$. The discount rate in this model is primarily important because it affects the continuation value of staying alive. The annual interest rate of 4% translates into value of statistical life equal to \$10 million, in line with the estimates in the literature (Andersson and Treich (2011), Kniesner and Viscusi (2019)). I assume that the government can borrow in the international markets at 1%, broadly matching the real borrowing costs observed across the industrial economies (Rachel and Summers 2019). In line with Rachel (2020), I calibrate $A$ to $\frac{1}{5} \cdot 24 \cdot 7 \cdot 34$, resulting in a per-capita annual income of \$60, 000. I set $\tau_n$ at 34%, reflecting the average tax rate in OECD economies.

The key parameter guiding the trade-off between the burden of death and the hardship of lockdown is $h$, which represents the value of home production.[15] We set $h$ to hit the replacement rate of 80%, so that:

$$\frac{h}{A(1-\tau_n)} = 80\%.$$

This choice is motivated by income support measures introduced by several countries, which provided furloughed workers with up to 80% of their pre-shock salary. Furthermore, we assume that $h$ is split equally between working from home, home production and government transfer: $\psi_{WFH} = \psi_{HP} = \psi_{GOV} = \frac{1}{3}$. This is motivated by the

[15] If $h$ is high and close to the wage income obtained while working, then lockdown is a relatively painless experience and we might expect the self-imposed lockdown to occur earlier and last for longer. Conversely, when $h$ is low, the trade-off between material well-being and possibility of illness and death is much steeper.

early estimates on the proportion of people that can work from home plus those who work in the essential sectors, as well as the initial estimates of the share of workers who have been furloughed (Dingel and Neiman (2020), Tomer and Kane (2020), Davies (2020)).

**Epidemiology Parameters**

Next, we calibrate the epidemiological parameters, broadly aiming to reflect the Covid-19 experience.

We assume that the initial seed of infection represents 0.1% of the total population. The recovery and death parameters $\gamma_r$ and $\gamma_d$ are determined by the mortality rate and the average time it takes to recover or die after contracting the virus. Since the model is weekly, a baseline case fatality rate of 1% and an average disease duration of 9 days,[16] we have that $\pi_d = 0.01 \cdot \frac{7}{9}$ and $\pi_r = \frac{7}{9} - \pi_d$.

In the baseline calibration we set the initial basic reproductive number to $\mathcal{R}_0 = 2.5$ (we explore robustness to doubling $\mathcal{R}_0$ below). This implies that the lockdown effectiveness threshold $\bar{\varepsilon}$ equals 0.39. We assume that lockdown effectiveness is less than the threshold, with $\varepsilon = 0.36$. Consequently, maximum restrictions reduce $\mathcal{R}_0$ by 36% to 1.6 (see e.g. Flaxman et al. (2020) or Lavezzo et al. (2020) for real-time studies of the infectiousness in the context of the Covid-19 pandemic). This calibration of lockdown effectiveness is qualitatively similar to the calibration in Eichenbaum et al. (2021).

### 4.2 Equilibrium and optimal lockdown

Figure 2 shows the trajectories of the epidemic in the decentralized equilibrium and under optimal lockdown policies. The main conclusion from this analysis is that the equilibrium trajectory differs very little from what is socially optimal. Both the planner and agents in the decentralized setting initially delay implementing restrictions (except if the planner knows exactly who is infected). The lockdown begins in week 4 or 5 of the epidemic, and lasts for about 20 to 24 weeks. The lockdown ends once the level of infectiousness is below 1% in all scenarios. Whether in the decentralized setting or in the socially optimal solution, the precautions save many lives: around a third of the population does not get sick, relative to the model with no behavioral response, and hence around $\frac{1}{3}\gamma_d$ of deaths, where $\gamma_d$ is the death rate, are avoided.

It is important to note that implementing optimal lockdown might present a challenge: agents begin isolation slightly before the planner's preferred lockdown starting date. This raises the question of implementability: whether in practice policymakers have any tools at their disposal that could maintain normal activity levels among the population during that time. One possible policy to address this timing discrepancy could involve subsidizing social interactions during that brief period. For example, a policy similar to the 'Eat Out To Help Out' program implemented in the United Kingdom in the summer of 2020, which subsidized 50% of restaurant bills, could

[16] Lauer et al. (2020) and Liu et al. (2020) estimated that the average latent and infectious periods in the case of Covid-19 were between 3-6 days, giving the overall duration of infection of around 9 days.

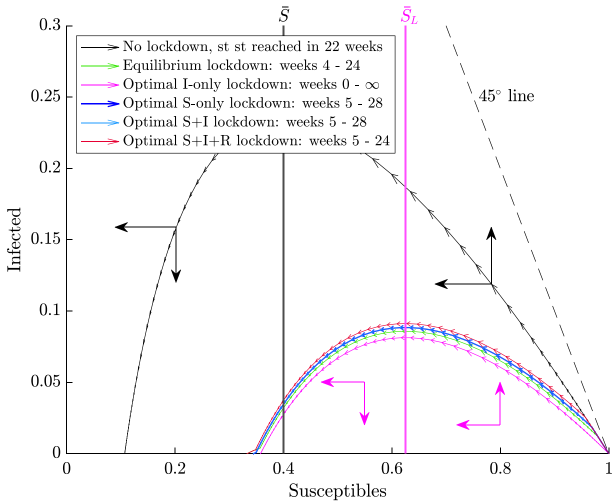


**Fig. 2** Equilibrium and optimal lockdown when $\varepsilon < \bar{\varepsilon}$

be considered.[17] See McAdams and Day (2024) for analysis of both "stay-at-home orders" as well as "go-out orders" in the context of a political economy model of an epidemic.[18]

### 4.3 Sensitivity to parameters

A striking result from the previous section is that the equilibrium and socially optimal mitigation strategies are similar, leading to nearly identical disease trajectories. As shown in Fig. 3, this result remains robust even under significant variations in the epidemiological parameters. The left-most panel shows the baseline calibration of the model. The middle panel shows the trajectories when the disease is much more transmissible, with $\mathcal{R}_0 = 5$, inspired by the Omicron variant of the Covid-19 virus. In this scenario, lockdown effectiveness and the infection fatality rate are held constant. The trajectories reach a higher infection rate and result in higher cumulative levels of infections. As a result, the epidemic is shorter in duration. However, as in the baseline scenario, the equilibrium and optimal trajectories remain very similar.

In the right panel, we consider both a high transmission rate and an infection fatality rate that is twice as high as the baseline. A more deadly disease results in a slightly longer lockdown, as both individuals and the planner balance fewer infections against more restrictions. However, the differences between the two strategies remain small. This outcome is unsurprising given that both strategies involve maximum lockdown for much of the epidemic, even in the baseline scenario. The potential to implement

[17] This, of course, does not suggest that the "Eat Out To Help Out" policy was socially optimal. For an empirical analysis of the 'Eat Out To Help Out' policy, see Fetzer (2022).

[18] A related issue of lockdown fatigue is analyzed in Carnehl et al. (2024).

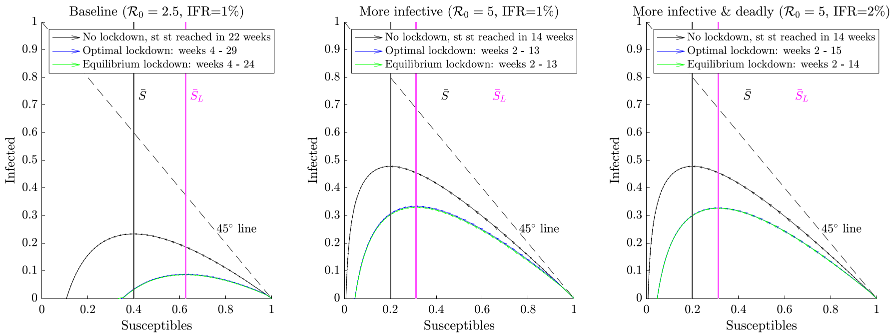


**Fig. 3** Lockdown effectiveness thresholds. Note: for clarity of exposition, the figure shows only the broad S+I+R lockdown

additional precautions due to the increased death risk is limited. This limitation applies to both agents in the decentralized equilibrium and the planner.

### 4.4 The effects of fiscal externality

Recall that the government covers one-third of the lockdown costs while also facing reduced tax revenues. How significant is the fiscal cost in shaping optimal lockdown policy? In the case where fiscal costs are considered, the optimal lockdown is approximately 12-14% (or 2 to 3 weeks) shorter compared to a scenario without fiscal costs (Fig. 4). As a result, fiscal externalities bring the timing of restrictions in the optimal policy closer to that in the decentralized equilibrium.[19]

### 4.5 Social welfare function

In practice, governments may lack complete information about the state of the epidemic and face operational constraints that limit their responsiveness. As a result, implemented policies may differ from the optimal strategies described above. This raises the question: how costly are such deviations?

To explore this, we can compute social welfare as a function of the lockdown's start and end dates. Figure 5 shows social welfare as a function of any permissible pair of start and end dates for lockdowns of the four types, when lockdown effectiveness is less than $\bar{\varepsilon}$. The vertical axis measures welfare loss in dollar terms, relative to a no-Covid counterfactual. The black 45-degree line represents the "no lockdown" scenario.

Implementing lockdown after week 22 is ineffective because, by that point, the epidemic has already run its course. In this region, the social welfare function slopes downward as lockdown duration increases, reflecting the rising cost of longer lockdowns.

While the shape of the social welfare function is similar for broad lockdowns, it differs significantly for isolation of the infected. In this case, welfare decreases as the

[19] This raises a natural question: what is the optimal level of income support during a lockdown, and how does it influence the design of lockdown policy? These important questions are left for future research.

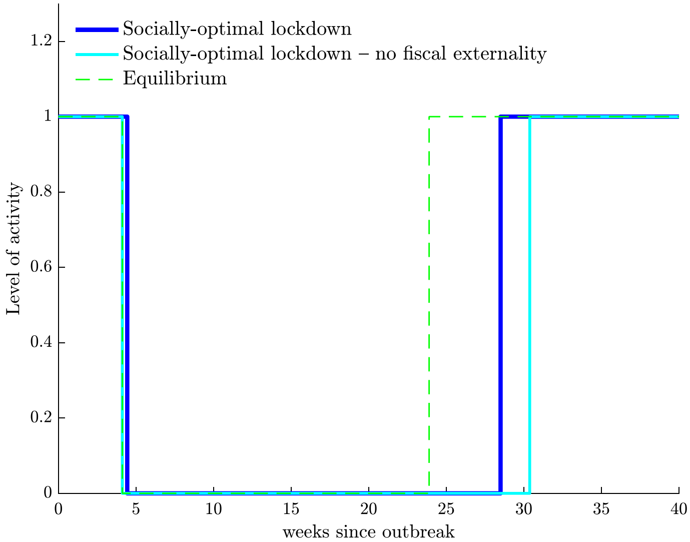


**Fig. 4** Activity level in the social optimum, with and without the fiscal externality. *Note*: for clarity of exposition, the figure shows only the broad S+I+R lockdown

start date is delayed and increases as the end date is extended. This result is intuitive, as the optimal strategy for isolating the infected involves immediate and permanent isolation.

For broader lockdowns, the figure shows that starting and ending the lockdown too early is not effective. An early and short lockdown leads to a large second wave later on, offering minimal health benefits while incurring significant economic costs.

### 4.6 Health and economic impacts of lockdown policies

Table 1 compares the four optimal lockdown strategies to the no lockdown scenario and the equilibrium lockdown, in terms of their epidemic, macro and welfare implications. The lockdown instruments are again ordered from the most targeted ($I$-only lockdown) to the broadest (the broad $S + I + R$ lockdown).

The first three rows review the timing and the duration of lockdowns across the scenarios. Equilibrium lockdown lasts for 20 weeks. The optimal lockdowns tend to start around the same time, and last for longer. The duration of the optimal lockdown falls with the breadth of the instrument, since the broader the lockdown, the higher the economic cost.

The next three rows show how the epidemic develops under these different policies. Immediate and permanent restrictions of $I-$only lockdown achieve the feasible minimum of cumulative infections. But the broader policies get quite close. Overall, the

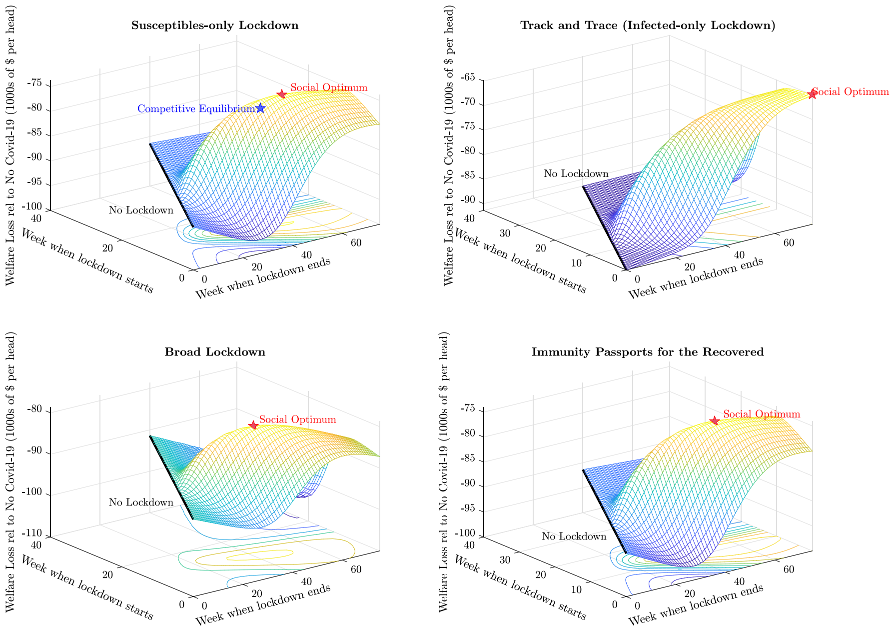


**Fig. 5** The Social Welfare Function for the four lockdown instruments when lockdown lowers $\mathcal{R}_0$ to 1.6

different optimal lockdown policies yield a similar picture in terms of the development of the disease, in line with the result in Proposition 3.[20]

The next three rows focus on the macroeconomic impact in the first year following the outbreak. The macro impact varies substantially across the scenarios. As highlighted above, the hypothetical no lockdown scenario - in which neither behavioral nor policy response to the virus are allowed - is associated with only a minor recession. This outcome is highly unlikely in human populations, however, since decentralized equilibrium mitigation leads to a large contraction in GDP of 15% in the first year; aggregate consumption declines by less as home production and government transfers provide a cushion. The required direct fiscal transfer is of the order of 7 percentage points of GDP. These numbers rise further in magnitude in the case of the optimal susceptible-only lockdown, simply due to its longer duration,[21] and in case of $S+I+R$ lockdown due to its breadth. The infected-only lockdown emerges as an effective and cheap way to limit the spread of the virus, highlighting the huge economic benefit that

[20] Relative to no lockdown, policies can reduce the share of people who contract the disease by over 20 percentage points by reducing the overshooting. Peak infections are reduced by a factor of two-and-a-half, even though peak infections are not in the planner's objective function in the baseline model (recall that the death rate is exogenous and constant). Given the constant death rate, deaths decline in line with the share of the population that contracted the virus.

[21] Note that the macroeconomic costs rise less-than-proportionately with lockdown duration. This is because the additional weeks happen at the end of the epidemic, when the number of susceptibles (who are in lockdown) is much lower than at the beginning. In other words in the case of $S$-only lockdown the first week of lockdown is always the most expensive.

can be brought by massive testing programs that might enable such instrument to be used in practice.

The final two rows of the table provide a summary measure that weighs the benefits and costs of the lockdown strategies in terms of saved lives and lost economic activity. When the virus first emerges and no mitigation strategies are possible, representative consumer's *lifetime* welfare drops by around $-0.9\%$, equivalent to \$92,000 *per capita* in terms of current income and consumption (in other words, one year and a half worth of income). This large welfare loss is reduced by a fifth with the equilibrium lockdown. Optimal lockdown policies can improve on the equilibrium only marginally, except if the authorities can identify and isolate the infected. If the government can only lock up the susceptibles it will do so for longer than in equilibrium, but the gains from fewer deaths are rather small (recall that at this point the epidemic moves slowly, so the extra 3 weeks of lockdown does not change the overall infection path all that much). If the only tool at the government disposal is the $S + I + R$ lockdown, then its optimal use performs worse than the equilibrium outcome.[22]

## 5 Conclusion

This paper presents two sets of results. The first set examines lockdown effectiveness and the potential for a second wave of infections. We demonstrated that the extent to which policy or individual behavior can reduce the infection rate plays a critical role in shaping both equilibrium dynamics and optimal policy outcomes.

The second set of results is quantitative and focuses on the case of limited lockdown effectiveness. Under a baseline parametrization, the optimal broad $S+I+R$ lockdown lasts approximately 22 weeks, and the resulting epidemic trajectory is similar to that observed in the decentralized equilibrium. The numerical analysis indicates that fiscal externalities significantly shorten the optimal lockdown by approximately 2-3 weeks. Additionally, I documented the shape of the social welfare function, which enabled an analysis of the welfare costs associated with mistimed policy interventions.

## Appendix

### Proof of Proposition 1

***Proof*** Equation (9) follows from the definitions of $\mathcal{R}_0$ and $\bar{S}$ and simple algebra. Immediate lockdown (i.e one with $T_0 = 0$) results in stable suppression if and only if

$$S_L(\infty) < \bar{S}.$$

Because $\bar{S}$ is fixed, to prove the result it suffices to show that $S_L(\infty)$ is increasing in $\varepsilon$. We have:

[22] It is important to be cautious when making this comparison, as the planner in this case has a large informational disadvantage relative to the individuals who know their health status perfectly.

**Table 1** Social and economic effects of the decentralized equilibrium lockdown and optimal lockdown policies, $\varepsilon < \bar{\varepsilon}$

| | | | Optimal lockdown policies | | | |
|---|---|---|---|---|---|---|
| | No lockdown | Equilibrium | I-only | S-only | S+I | S+I+R |
| *Lockdown* | | | | | | |
| Week lockdown begins | – | 4 | 0 | 4 | 4 | 4 |
| Week lockdown ends | – | 24 | $\infty$ | 29 | 29 | 26 |
| Duration (weeks) | – | 20 | $\infty$ | 25 | 25 | 21 |
| *Epidemic (% of population)* | | | | | | |
| Contracted virus | 89.3 | 66.7 | 64.2 | 65.2 | 65.2 | 65.9 |
| Peak infected | 23.4 | 8.6 | 8.1 | 8.7 | 8.7 | 8.7 |
| Total deaths | 0.89 | 0.67 | 0.64 | 0.65 | 0.65 | 0.66 |
| *Economic Impact* | | | | | | |
| Agg C loss (%) | −1 | −6 | −1 | −6 | −6 | −10 |
| GDP loss (%) | −2 | −15 | −2 | −16 | −17 | −26 |
| Fiscal support (rise in Debt/GDP, pp) | 0 | 7 | 1 | 8 | 8 | 13 |
| *Lifetime utility rel to no virus* | | | | | | |
| Welfare Loss (% of life utility) | −0.89 | −0.70 | −0.63 | −0.69 | −0.69 | −0.72 |
| Welfare Loss ($ per capita) | −92000 | −72000 | −66000 | −71000 | −71000 | −75000 |

$$S_L(\infty) - \bar{S}_L \log S_L(\infty) - 1 = 0.$$

Combining this with the result that $\bar{S}_L = \frac{\bar{S}}{1-\varepsilon}$ (implied by Eq. (9)) yields:

$$S_L(\infty) - \frac{\bar{S}}{1-\varepsilon} \log S_L(\infty) - 1 = 0. \tag{16}$$

Differentiating equation (16) with respect to $\varepsilon$ and rearranging we get:

$$\frac{\partial S_L(\infty)}{\partial \varepsilon} = \frac{\bar{S}(1-\varepsilon)^{-2} \log S_L(\infty)}{1 - \frac{\bar{S}_L}{S_L(\infty)}} \geq 0,$$

where the inequality follows from the fact that $S_L(\infty) \leq 1$ and $\bar{S}_L > S_L(\infty)$.

To find the value of the threshold note that it is pinned down by the condition $S_L(\infty) = \bar{S}$. Equation (9) implies that $\varepsilon = 1 - \frac{\bar{S}}{\bar{S}_L}$. And so:

$$-\bar{S} + \bar{S}_L \log \bar{S} + 1 = 0.$$

Dividing by $\bar{S}_L$ yields

$$-\frac{\bar{S}}{\bar{S}_L} + \log \bar{S} + \frac{1}{\bar{S}_L} = 0.$$

Therefore:

$$\bar{\varepsilon} - 1 + \log \bar{S} + \frac{1-\bar{\varepsilon}}{\bar{S}} = 0.$$

Rearranging:

$$\bar{\varepsilon} = \frac{\bar{S}}{\bar{S}-1}\left(1 - \log \bar{S} - \frac{1}{\bar{S}}\right).$$

The final part of the Proposition follows from the fact that $\hat{\varepsilon}$ is defined by $\mathcal{R}_0^L = \bar{S}_L = 1$. The trajectory post-lockdown is the same as in the no lockdown scenario because the initial seed of infection is assumed to be infinitesimally small. □

#### Proof of Proposition 2

***Proof*** Consider the trajectory of the epidemic from the end of the lockdown $T_1$ onwards, which is characterized by:

$$I(t) = -S(t) + \bar{S} \log S(t) + I_R + S_R - \bar{S} \log S_R,$$

where $(S_R, I_R)$ is the state of the system at $T_1$, and so $I_R = -S_R + \bar{S}_L \log S_R + 1$. For $T_1$ sufficiently large, $I_R \approx 0$. Thus:

$$I(t) = -S(t) + \bar{S} \log S(t) + S_R - \bar{S} \log S_R.$$

The peak occurs at $S(t) = \bar{S}$:

$$I_{max} = -\bar{S} + \bar{S} \log \bar{S} + S_R - \bar{S} \log S_R.$$

Differentiating with respect to $\varepsilon$ :

$$\frac{\partial I_{max}}{\partial \varepsilon} = \left(1 - \frac{\bar{S}}{S_R}\right) \frac{\partial S_R}{\partial \varepsilon} > 0,$$

where the inequality follows from the fact that $\bar{\varepsilon} < \varepsilon < \hat{\varepsilon}$ and $\frac{\partial S_R}{\partial \varepsilon} > 0$.

The final resting point is given by:

$$0 = -S(\infty) + \bar{S} \log S(\infty) + S_R - \bar{S} \log S_R.$$

Differentiating:

$$\frac{\partial S(\infty)}{\partial \varepsilon} = \frac{S(\infty)}{S(\infty) - \bar{S}} \left(1 - \frac{\bar{S}}{S_R}\right) \frac{\partial S_R}{\partial \varepsilon} > 0,$$

which proves the result. □

### Derivation of the first order conditions (14) and (15)

The first order necessary condition for $T_0$ to be optimal is:

$$\begin{aligned}
& e^{-\rho T_0 - \int_0^{T_0} \pi_W(s) ds} \left(u(w) + \pi_W(T_0) V(T_0)\right) \\
& -\pi_W(T_0) e^{-\int_0^{T_0} \pi_W(s) ds} \int_{T_0}^{T_1} e^{-\rho t - \int_{T_0}^{t} \pi_H(s) ds} \left(u(h) + \pi_H(t) V(t)\right) dt \\
& -e^{-\int_0^{T_0} \pi_W(s) ds} \left(e^{-\rho T_0} \left(u(h) + \pi_H(T_0) V(T_0)\right)\right) \\
& +e^{-\int_0^{T_0} \pi_W(s) ds} \int_{T_0}^{T_1} \pi_H(T_0) e^{-\rho t - \int_{T_0}^{t} \pi_H(s) ds} \left(u(h) + \pi_H(t) \left(V(t)\right)\right) \\
& -(\pi_W(T_0) - \pi_H(T_0)) e^{-\rho T_1 - \int_0^{T_0} \pi_W(s) ds - \int_{T_0}^{T_1} \pi_H(s) ds} W^S(T_1) = 0.
\end{aligned}$$

Re-arranging:

$$u(w) - u(h)$$

$$
\begin{aligned}
&+ (\pi_W(T_0) - \pi_H(T_0)) \cdot (W^I(T_0) \\
&- \int_{T_0}^{T_1} e^{-\rho(t-T_0)-\int_{T_0}^{t} \pi_H(s)ds} \left(u(h) + \pi_H(t) W^I(t)\right) dt \\
&- e^{-\rho(T_1-T_0)-\int_{T_0}^{T_1} \pi_H(s)ds} W^S(T_1)) = 0. \qquad (17)
\end{aligned}
$$

Since $W^S(T_0) = \int_{T_0}^{T_1} e^{-\rho(t-T_0)-\int_{T_0}^{t} \pi_H(s)ds} \left(u(h) + \pi_H(t)W^I(t)\right) dt + e^{-\rho(T_1-T_0)-\int_{T_0}^{T_1} \pi_H(s)ds} W^S(T_1))$ we obtain:

$$
u(w) - u(h) + (\pi_W(T_0) - \pi_H(T_0)) \cdot (W^I(T_0) - W^S(T_0)) = 0. \qquad (18)
$$

The first order condition for $T_1$ is:

$$
\begin{aligned}
&-e^{-\rho T_1 - \int_0^{T_0} \pi_W(s)ds - \int_{T_0}^{T_1} \pi_H(s)ds} \left(u(h) + \pi_H(T_1) W^I(T_1)\right) \\
&\quad -\pi_L(T_1) e^{-\int_0^{T_0} \pi_W(s)ds - \int_{T_0}^{T_1} \pi_H(s)ds} \int_{T_1}^{\infty} e^{-\rho t - \int_{T_1}^{t} \pi_W(s)ds} \left(u(w) + \pi_W(t) W^I(t)\right) dt) \\
&\quad + e^{-\int_0^{T_0} \pi_W(s)ds - \int_{T_0}^{T_1} \pi_H(s)ds} \\
&\quad \times \left\{ -e^{-\rho T_1} \left(u(w) + \pi_W(T_1) W^I(T_1)\right) \right. \\
&\quad \left. + \int_{T_1}^{\infty} \pi_W(T_1) e^{-\rho t - \int_{T_1}^{t} \pi_W(s)ds} \left(u(w) + \pi_W(t) W^I(t)\right) dt \right\} = 0.
\end{aligned}
$$

Re-arranging:

$$
\begin{aligned}
&\left(u(h) + \pi_H(T_1) W^I(T_1)\right) \\
&-\pi_H(T_1) \int_{T_1}^{\infty} e^{-\rho(t-T_1) - \int_{T_1}^{t} \pi_W(s)ds} \left(u(w) + \pi_W(t) W^I(t)\right) dt) \\
&- \left(u(w) + \pi_W(T_1) W^I(T_1)\right) \\
&+ \int_{T_1}^{\infty} \pi_W(T_1) e^{-\rho(t-T_1) - \int_{T_1}^{t} \pi_W(s)ds} \left(u(w) + \pi_W(t) W^I(t)\right) dt = 0.
\end{aligned}
$$

which finally yields:

$$
u(h) - u(w) + (\pi_H(T_1) - \pi_W(T_1)) \left[ W^I(T_1) - W^S(T_1) \right] = 0,
$$

which is the same as in the text.

### Existence of the pure strategy equilibrium

Consider first the existence of the pure strategy equilibrium. Let $\mathcal{L} \subset \mathbb{R}^2_+$ denote the set of feasible lockdown strategies of a given effectiveness $\varepsilon$. Each element of the set is a vector of two positive real numbers, $\{T_0, T_1\}$, with $T_0 \leq T_1$. This set is closed. As long as lockdown is not costless it is also bounded, since lockdown that lasts forever cannot be optimal. This imposes a condition on the end of the lockdown: $T_1 \leq T^V$ where $T^V$ is arbitrarily large but finite. Bounded in this way $\mathcal{L}$ is a compact set. Let $\mathcal{E}$ denote the set of sequences fully characterizing the epidemic trajectory: $\mathcal{E} := \{I, \pi_W, \pi_U\}_{t=0}^{\infty}$. Note that $\mathcal{E}$ contains all the epidemic information that is relevant for individual problem in (12). Let $f : \mathcal{L} \rightarrow \mathcal{E}$ be the correspondence that maps a given lockdown strategy to the resulting epidemic trajectory, and let $g : \mathcal{E} \rightarrow \mathcal{L}$ be the correspondence that maps the epidemic to optimal lockdown strategy. Finally, let $\mathcal{F} := f \circ g$, $\mathcal{F} : \mathcal{L} \rightarrow \mathcal{L}$ be the correspondence that maps a given lockdown strategy to optimal lockdown strategies. An equilibrium is a fixed point of $\mathcal{F}$. Kakutani's Fixed Point Theorem guarantees that $\mathcal{F}$ has a fixed point if set $\mathcal{L}$ is nonempty, compact and convex and $\mathcal{F}$ is upper hemicontinuous and convex-valued, that is the set $\mathcal{F}(x) \subset \mathcal{L}$ is nonempty and convex for every $x \in \mathcal{L}$. Clearly, the first four requirements are satisfied. But convex-valuedness of $\mathcal{F}$ fails when $\varepsilon > \bar{\varepsilon}$ because $g$ is not convex-valued when the infection curve is double-peaked. Loosely, a susceptible individual who faces a double peak of infections can use the lockdown to avoid the risk of one or the other. A convex combination of such strategies will result in a lockdown that happens in between the two peaks, which is clearly suboptimal. This proves that $g$ is not convex-valued and conditions for the existence of a fixed point are not satisfied.



## References

Acemoglu D, Chernozhukov V, Werning I, Whinston MD (2021) Optimal targeted lockdowns in a multi-group SIR model. Am Econ Rev Insights 3(4):487–502

Alvarez F, Argente D, Lippi F (2021) A simple planning problem for COVID-19 lock-down, testing, and tracing. Am Econ Rev Insights 3(3):367–382

Andersson H, Treich N (2011) The Value of a Statistical Life. In: *A Handbook of Transport Economics*, chapter 17. Edward Elgar Publishing

Antràs P, Redding SJ, Rossi-Hansberg E (2023) Globalization and Pandemics. Am Econ Rev 113(4):939–981

Atkeson A (2020) What will be the economic impact of COVID-19 in the US? Rough Estimates of Disease Scenarios, Technical report, National Bureau of Economic Research, Cambridge, MA

Atkeson A (2023) The impact of vaccines and behavior on US cumulative deaths from COVID-19. Rev Econ Des 2004:67

Avery C, Bossert W, Clark A, Ellison G, Ellison SF (2020) An economist's guide to epidemiology models of infectious disease. J Econ Perspect 34(4):79–104

Avery C, Chen F, McAdams D (2024) Steady-state social distancing and vaccination. Am Econ Rev Insights 6(1):1–19

Bahaj S, Reis R (2020) Central bank swap lines during the Covid-19 pandemic. Covid Econ 2:1–12

Berger D, Herkenhoff K, Mongey S (2020) An SEIR infectious disease model with testing and conditional quarantine

Budish E (2024) R<1 as an Economic Constraint. *Review of Economic Design*

Carnehl C, Fukuda S, Kos N (2024) Time-varying cost of distancing : Distancing fatigue and lockdowns

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

Davies R (2020) Coronavirus and the social impacts on Great Britain: 23 April 2020. *Office for National Statistics*, (April):1–12

Dingel JI, Neiman B (2020) How many jobs can be done at home? J Public Econ 189(March):104235

Eichenbaum MS, Rebelo S, Trabandt M (2021) The macroeconomics of epidemics. Rev Financ Stud 34(11):5149–5187

Farboodi M, Jarosch G, Shimer R (2021) Internal and external effects of social distancing in a pandemic. J Econ Theory 196:105293

Feng Z (2007) Final and peak epidemic sizes for Seir models. Math Biosci Eng 4(4):675–686

Fenichel EP (2013) Economic considerations for social distancing and behavioral based policies during an epidemic. J Health Econ 32(2):440–451

Fetzer T (2022) Subsidising the spread of COVID-19: evidence from the UK'S eat-out-to-help-out scheme*. Econ J 132(643):1200–1217

Flaxman S, Mishra S, Gandy A, Unwin JT, Coupland H, Mellan TA, Zhu H, Berah T, Eaton JW, Guzman PNP, Schmit N, Cilloni L, Ainslie KEC, Baguelin M, Blake I, Boonyasiri A, Boyd O, Cattarino L, Ciavarella C, Cooper L, Cucunubá Z, Cuomo-Dannenburg G, Dighe A, Djaafara B, Dorigatti I, Van Elsland S, Fitzjohn R, Fu H, Gaythorpe K, Geidelberg L, Grassly N, Green W, Hallett T, Hamlet, A, Hinsley W, Jeffrey B, Jorgensen D, Knock E, Laydon D, Nedjati-Gilani G, Nouvellet P, Parag K, Siveroni I, Thompson H, Verity R, Volz E, Gt Walker P, Walters C, Wang H, Wang Y, Watson O, Xi X, Winskill P, Whittaker C, Ghani A, Donnelly CA, Riley S, Okell LC, Vollmer MAC, Ferguson NM, Bhatt S (2020) Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries. *Imperial College London*, (March):1–35

Gans JS (2023) Vaccine hesitancy, passports, and the demand for vaccination. Int Econ Rev 64(2):641–652

Garibaldi P, Moen E, Pissarides C (2020) Modelling contacts and transitions in the SIR epidemics model. CEPR Covid Econ 5:1–20

Giannitsarou C, Kissler S, Toxvaerd F (2021) Waning immunity and the second wave: some projections for SARS-CoV-2. Am Econ Rev Insights 3(3):321–338

Glover A, Heathcote J, Krueger D, Ríos-Rull J-V (2023) Health versus wealth: on the distributional effects of controlling a pandemic. J Monet Econ 140:34–59

Guerrieri V, Lorenzoni G, Straub L, Werning I (2022) Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages? Am Econ Rev 112(5):1437–1474

Jones C, Philippon T, Venkateswaran V (2021) Optimal mitigation policies in a pandemic: social distancing and working from home. Rev Financ Stud 34(11):5188–5223

Jordà Ò, Singh SR, Taylor AM (2020) Longer-run economic consequences of pandemics. Rev Econ Stat 104:166

Kaplan G, Moll B, Violante G (2020) The great lockdown and the big stimulus: tracing the pandemic possibility frontier for the U.S. SSRN Electron J

Kniesner TJ, Viscusi WK (2019) The value of a statistical life. Oxf Res Encycl Econ Financ 19:15–19

Lauer SA, Grantz KH, Bi Q, Jones FK, Zheng Q, Meredith HR, Azman AS, Reich NG, Lessler J (2020) The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. Ann Intern Med, 2019

Lavezzo E, Franchin E, Ciavarella C, Cuomo-dannenburg G, Barzon L, Sciro M, Merigliano S, Decanale E, Vanuzzo MC, Onelia F, Pacenti M, Parisi S, Carretta G, Donato D, Gaythorpe KAM, Alessandra R (2020) Suppression of COVID-19 outbreak in the municipality of Vo', Italy. (Ci):1–23

Liu Y, Gayle AA, Wilder-Smith A, Rocklöv J (2020) The reproductive number of COVID-19 is higher compared to SARS coronavirus. J Travel Med 27(2):1–4

McAdams D (2021) The blossoming of economic epidemiology. Ann Rev Econ 13:539–570

McAdams D, Day T (2024) The political economy of epidemic management. Rev Econ Des

McAdams D, Song Y, Zou D (2023) Equilibrium social activity during an epidemic. J Econ Theory 207:105591
Miclo L, Spiro D, Weibull J (2022) Optimal epidemic suppression under an ICU constraint: an analytical solution. J Math Econ 101:102669
Phelan T, Toda AA (2022) Optimal epidemic control in equilibrium with imperfect testing and enforcement. J Econ Theory 206:105570
Piguillem F, Shi L (2022) Optimal Covid-19 quarantine and testing policies. Econ J 132(647):2534–2562
Pollinger S (2023) Optimal contact tracing and social distancing policies to suppress a new infectious disease. Econ J 133(654):2483–2503
Rachel L (2020) Leisure-enhancing technological change
Rachel L (2023) An analytical model of behavior and policy in an epidemic
Rachel L, Summers LH (2019) On secular stagnation in the industrialized world. Brook Pap Econ Act 2019(Spring):1–76
Stock J (2020) Data gaps and the policy response to the novel coronavirus. Technical report, National Bureau of Economic Research, Cambridge, MA
Tomer A, Kane JW (2020) How to protect essential workers during COVID-19
Toxvaerd F (2020) Equilibrium social distancing