An Analytical Model of Covid-19 Lockdowns*

Łukasz Rachel†

December 4, 2020

Abstract

This paper develops an analytical framework for studying risk mitigation behaviors and policies during an epidemic. The analytics uncover two novel insights. First, individual precautionary behavior dramatically flattens the epidemic curve, so much so that the infection externality leads to too much, not too little, social distancing in equilibrium. Second, the optimal policy does not flatten the curve per se; instead it avoids the second wave of infections and prevents the epidemic overshoot, minimizing cumulative deaths. The optimal policy is almost entirely independent of the economic parameters such as the value of statistical life.

JEL Classification: E1, I1, H0.

*I am grateful to Ricardo Reis, Ben Moll, Guido Lorenzoni, Martin Eichenbaum, Ron Milo, Matt Rognlie, Larry Christiano, Christopher Pissarides, Eran Yashiv and participants at workshops at the LSE, Northwestern University and the BoE for helpful comments. Views expressed here are solely of the author. *First Draft:* April 2020. *Disclaimer*: I have benefited hugely from discussions with colleagues in epidemiology, but ultimately I am an economist, not an epidemiologist. This paper uses the basic epidemiological model together with the tools of modern macroeconomics to better understand the current crisis.

†LSE and the Bank of England. *Email:* l.p.rachel@lse.ac.uk. *Web:* https://sites.google.com/site/lukaszrachel/

1 Introduction

This paper studies the risk-mitigation decisions of individuals and governments faced with a contagious and deadly disease. As is common in the recent macro-epidemiology literature, the model appends the baseline SIR framework with individuals and policy-makers who act to maximize their own or society's welfare. What is new is that the framework enables a sharp analytical characterization of these behaviors and their impact.

Crucially, the analytics help uncover important results that papers based on simulating macro-SIR models have missed. Rather than trying to write the most complete model, the approach in this paper is to develop an in-depth understanding of the benchmark which underlies richer, quantitative models in the literature and so makes these models more interpretable. Indeed, the core model that I work with here underlies most – if not all – of the macro-epidemiology models written since the Covid-19 crisis began.

I start with the analysis of the decentralized equilibrium in which perfectly informed individuals decide optimally on costly mitigation behavior.[1] Given the knowledge about own health status and no altruism, individuals who are currently infected or recovered never choose to lock down in equilibrium. I derive a closed-form expression that approximates the degree of equilibrium mitigation of the susceptibles. Their cautious behavior delivers a gently declining number of new infections over time, and thus leads to a dramatic flattening of the epidemic curve in the aggregate. The peak of infections occurs at the point when individuals first start to mitigate the risk, at which point the infection rate is way below the peak of the uncontrolled epidemic path. The effective reproduction

[1]The equilibrium analysis can easily be extended to a case of imperfect knowledge of individual health status. Uncertainty about individual immunity matters for individual behavior only in the latter stage of the epidemic, when the level of susceptibility is relatively low in the aggregate. At that point imperfect knowledge makes the susceptibles less cautious, since they operate under the assumption that with some probability they are already immune. The detailed study of this case is omitted for brevity.

Figure 1
$\mathcal{R}$ across countries remains very close to 1

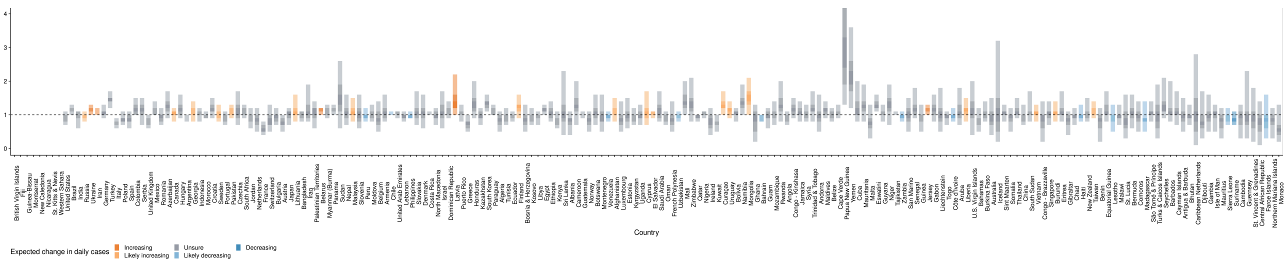

Source: https://epiforecasts.io/covid/posts/global/.

number $\mathcal{R}$ is slightly below one from that point onwards, broadly consistent with the data observed across many countries and regions after the first wave of the epidemic (Figure 1 shows the data for end-November 2020).[2]

I provide a closed-form expression that approximates the equilibrium infection rate, and show how to derive intuitive comparative statics. Furthermore, I show that the reduction in the infection rate relative to the mechanistic model comes at a cost of a much longer duration of the epidemic. All-in-all, the precautionary behavior significantly lowers cumulative deaths. But the epidemic persists for a long-time and so it is economically costly.

Next, I turn to the analysis of optimal lockdown policy. I consider four alternative instruments at the planner's disposal, from more to less information-hungry (equivalently,

[2]See https://rt.live/ and https://epiforecasts.io/covid/posts/global/ for real-time estimates of $\mathcal{R}(t)$ across US states and countries/regions, respectively. Since across the globe the level of susceptibility in all likelihood remains well above the herd immunity threshold, a mechanistic SIR model would have predicted a continued exponential rise. Containment policies played a part in preventing this outcome, but the fact that the epidemic trajectory is similar across countries with quite different policies suggests that individual precautionary response is a major factor in how the epidemic evolved (Cochrane (2020), Gans (2020)). Indeed, the individually responsive behavior predicted by the model is supported by the empirical analysis of actual behaviors during the Covid-19 epidemic (Goolsbee and Syverson, 2020). Of course the infection rate fluctuates over time – in many countries in Europe, the summer months of 2020 were associated with a lower infection rate, before the second wave of infection arrived in the autumn. These fluctuations could be driven by many factors, not least by the lags with which information can be processed by individuals translating into the corresponding mistakes in mitigation behavior. Nonetheless, these fluctuations are not a failure of the model – on the contrary, since they are relatively small in magnitude relative to the infection rates predicted by the mechanistic epidemiology model (see e.g. Ferguson *et al.* (2020)), they are broadly consistent with the behavioral SIR framework considered here and elsewhere in the macro-epi literature.

from more targeted to broader), with the isolation of the infected on one extreme, and mitigation measures applied to the entire population on the other.[3]

The best-possible policy instrument is the targeted isolation of the infected. If available, such isolation should start straight away and should be made permanent, because isolating the infected has a small cost (infection lasts for a short time, and the number of infected is low) but brings about first order health benefits: it can suppress the virus indefinitely.

The optimal strategy for the use of the broad lockdown policies is quite different, because locking down an economy is too costly to be in place indefinitely. A broad lockdown must thus be temporary, which means that the epidemic must run its course, with the level of susceptibility falling to at least the herd immunity threshold.[4] Given a constant infection fatality rate and assuming no cure or vaccine is forthcoming, the herd immunity threshold pins down the minimum feasible deaths. The planner seeks to avoid overshooting that minimum at smallest cost possible (aims to get there as quickly as possible). To achieve this the optimal strategy is to "hold fire" and only activate the lockdown relatively late, when the level of infections is high and the dynamics move fast – a stark illustration that optimal mitigation policy in the baseline macro-SIR model does not aim to flatten the curve and instead focuses on achieving the lowest cumulative death toll at minimum cost.

The optimal lockdown policy improves on the equilibrium mostly along the economic margin, with only a small improvement in terms of health outcomes.[5] This is because the cautious equilibrium behavior implies the long-run resting point is close to the herd

[3] In terms of broad lockdowns, I focus solely on mitigation (as opposed to suppression) measures. The reason for this is that time-limited suppression measures are not feasible, as I explain below. See Pollinger (2020) for a study of optimal suppression policy.

[4] The herd immunity threshold in this model is given by the inverse of the basic reproduction number $\mathcal{R}_0$.

[5] This is not to be confused with a very substantial improvement in terms of health outcomes relative to the naive mechanistic SIR model which assumes no policy nor behavioral response.

immunity threshold – the point associated with the minimum feasible death toll for *any* time-limited mitigation. The economic benefits of optimal policy on the other hand are substantial, as the planner cuts the duration of the crisis by an order of magnitude relative to the equilibrium.

Remarkably, the optimal lockdown policy is almost entirely independent of parameters that determine the value of statistical life (VSL). The lockdown starting point is fully pinned down by only two parameters: the basic reproduction number $\mathcal{R}_0$ and the lockdown effectiveness ε, where the latter is defined as a percentage reduction in $\mathcal{R}$ that is brought about by a lockdown. Both of these parameters sit on the epidemiological side of the model, while the VSL is determined by the economic parameters.[6] The end-point is does depend on parameters but the health consequences of changing the end-point are minuscule.[7] Contrary to much of the commentary in the literature, I conclude that it is *not* necessary for economists to "take a stand" on the value of statistical life in order to draw out normative implications in the context of a baseline epidemiological model.

Underlying the differences between the equilibrium and socially-optimal outcomes are two externalities: an *infection externality* and a *fiscal externality*.

My contribution is to point out that the infection externality present in the behavior of the susceptibles means there is *too much* social distancing in equilibrium relative to the social optimum, or more specifically, that the individually optimal mitigation might start earlier, and last longer, than the socially optimal one. The intuition for why the infection externality can work in this direction is that the (external) effects of today's decision to engage in social distancing lower the infection rate in the short-run but raise it further out (by the logic of flattening the curve – the curve is flatter but it is also fatter). This

[6]The underlying reason for why this near-independence occurs is that the optimal policy minimizes the number of deaths, and thus it is at the corner of the trade-off between deaths and economic costs. The corner solution is the optimal one for any reasonable parametrization of the model.

[7]This is true since at that point the level of susceptibility is close to the herd immunity threshold and the infection rate is low, making the system insensitive to policy changes.

intertemporal substitution of infection risk might occur at a very unfavorable trade-off from the social perspective. Another way to contrast the individuals' vis-a-vis the planner's mitigation strategies is that the individuals focus on the *infection risk at a point in time*, whereas the planner sees through to the end of the epidemic and cares about *cumulative infection risk*.

As to the fiscal externality, I impose that the government's role in the epidemic is to provide economic disaster relief in the form of income support.[8] In equilibrium, individuals take government transfers and future taxes as given, but their choices ultimately determine the level of debt incurred by the government during the epidemic and thus the level of future taxes – hence the externality.

In much of the paper I work with a simple extension of the workhorse epidemiology model of Kermack and McKendrick (1927). This baseline model omits several features of the environment which no doubt play a part in the Covid-19 crisis, most notably there are no ICU capacity constraints, nor is there a possibility that a treatment or a vaccine will be developed. Basing my analysis on the simple model is worthwhile because the mechanisms I explore lie at the center of many richer models and the hope is that lessons learnt here will be useful for understanding those more elaborate frameworks. In the final section of the paper I discuss how the aforementioned features change the analysis. In short, these elements put a premium on avoiding the strategies where the rate of infection is allowed to climb high. This discussion serves to highlight that the lesson from this paper ought not to be that the optimal policies prescribed by the analytical model should be applied directly. Instead, my model elucidates which considerations drive which characteristics of optimal policies, thus building a more thorough understanding of the trade-offs and through that enhancing the policy debate.

[8] specifically, I assume that the government covers a certain proportion of income of those who are in lockdown, financing this expenditure and the loss of tax revenues associated with lockdowns with borrowing.

Literature and contributions. This paper contributes to a fast growing literature on the macroeconomics of epidemics.[9] Several papers studied individual social distancing decisions in a decentralized equilibrium and socially optimal mitigation policies (Toxvaerd (2020), Farboodi *et al.* (2020), Jones *et al.* (2020), Alvarez *et al.* (2020), Piguillem and Shi (2020)).[10] Methodologically, I focus on the analytics of the baseline SIR model, which allow me to uncover several new results. I do not impose or externally calibrate the cost of death – instead this cost and all resulting value functions are consistent with the natural interpretation that the disutility of death as the foregone lifetime consumption utility, which changes endogenously over time.[11] My contribution is to show that optimal policy in the baseline model takes the form of a corner solution, and thus to demonstrate that optimal policy is independent of the parameter values such as the value of statistical life (VSL). This is an important result because the ability to draw normative implications has so far been hampered by the perceived high sensitivity of the results to the VSL – a parameter that is inherently difficult to calibrate definitively and convincingly.[12]

[9] Early papers by Atkeson (2020) and Stock (2020) provide an economist's perspective on the baseline SIR epidemiology models. Eichenbaum *et al.* (2020a) 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.

[10] Feng (2007) provides an epidemiological perspective on quarantine and isolation policies. Eichenbaum *et al.* (2020a) study the optimal tax on consumption during an epidemic. Acemoglu *et al.* (2020) 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; I consider different returns to scale in the infection matching technology in the Appendix.

[11] Unlike Toxvaerd (2020) I do not impose the simplifying steady-state assumptions over the transition and instead show how to characterize the solution to the full dynamic system.

[12] Beyond that, 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.* (2020) 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 *et al.* (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.* (2020) 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

To the best of my knowledge this paper is the first to highlight that the infection externality means that there is too much, not too little, social distancing in the decentralized equilibrium relative to the optimum. This stands in sharp contrast to some of the existing studies.[13]

Roadmap. The paper is structured as follows. Section 2 outlines the analytical model. 3 defines and studies the decentralized equilibrium. Section 4 studies the optimal lockdown policy, and Section 5 analyzes the externalities. Section 6 discusses how the results change when some additional features are incorporated in the environment. Section 7 concludes.

2 An analytical epi-macro model

The tractability obtains from a combination of three components. First, as mentioned above, the model abstracts from certain complicating factors such as endogenous infection fatality rates or the possibility of the vaccine. Once the results are obtained, it is then straightforward to study how the additional features of the environment change the conclusions. I take-up this task at the end of the paper. Second, I assume that the cost of social distancing (or lockdowns) is linear in its severity (whereas the literature tends to assume convex costs). This assumption is not unreasonable: it means that a short-but-strict lockdown can be just as costly as a milder but longer one. Indeed, anecdotal evidence and introspection suggests that many consumers and businesses might actually prefer a

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.

[13] The second contribution to the study of epidemic externalities is my analysis of the fiscal externality. Relative to the existing work my formulation – with the government financing its income support policies (and its tax shortfall) with emergency borrowing – leads to naturally adverse fiscal implications of mitigation behaviors and policies (whereas for example mitigation in Eichenbaum *et al.* (2020a) raise government revenue). In Rachel (2020) I show that this effect can be substantial quantitatively.

short-and-stark "circuit breaker" relative to a more drawn-out period of restrictions, casting some doubt on the validity of the convex cost assumption. The constant marginal cost of restrictions allows for solutions in the form of cut-off rules which can be characterized analytically. Again, once the sharp characterization of behavior is obtained, it is straightforward to understand the impact that convex costs have in this context. Third, I use the graphical apparatus of phase diagrams, which supplements the analytics and enhances the intuitive understanding of the dynamics of the disease.

2.1 Pre-epidemic environment

The economy is populated by a measure one of identical individuals whose instantaneous utility depends on consumption and labor supply $u(c, n)$ and whose time endowment is normalized to 1. Individuals solve the following utility maximization problem:

$$\max_{\lambda \in [0,1]} \int_0^\infty e^{-\rho t} \left(\lambda u^W + (1 - \lambda) u^L \right) dt \tag{1}$$

where λ is the probability of working and $u^W := u(w, 1)$ and $u^L := u(h, 0)$ denote the levels of utility[14] when individual is working (receiving wage w)[15] or staying at home (receiving income h).[16] Individuals choose the probability of working in any given period, so that the setup mirrors the Rogerson (1988) model of indivisible labor.[17] The purpose of this formulation is that it allows for a sharp analytical characterization of the social distancing behavior, since the marginal cost of social distancing is constant. An alternative assumption of a linear utility function (with a constant marginal utility of consumption)

[14]There is no capital or any other form of saving, so consumption is equal to income period-by-period.

[15]Government collects proportional taxes on gross labor income in normal times. The role of the government and its budget constraint are described in detail below.

[16]Income h is an exogenous parameter, and it is the sum of three components: income from market activities such as working from home (share ψ_{WFH}), home production (share ψ_{HPR}) and government transfer (share ψ_{GOV}). The shares ψ are constant and sum to one.

[17]This is equivalent to them choosing labor supply in a smooth fashion if marginal utility of consumption is constant (e.g. when the utility function is linear).

and a convex choice set for labor supply would deliver identical results.

The measure of individuals who work in the population is therefore given by λ. I assume that, absent any other considerations, individuals prefer to work: $u^W > u^L$. This implies that before the epidemic all individuals work: $\lambda = 1$.

Production technology is linear in labor, markets are competitive, and there is a government that taxes labor income at a rate τ_n and consumes the proceeds:

$$Y = AN \qquad w = A(1-\tau_n) \qquad G = \tau_n AN. \tag{2}$$

Market clearing conditions are:[18]

$$Y + (1-\lambda)h = C + G \qquad C = \lambda w + (1-\lambda)h \qquad N = \lambda + (1-\lambda)\psi_{WFH}\frac{h}{A(1-\tau_n)}. \tag{3}$$

2.2 Epidemic

To model infection I use a well-known SIR model with 4 population groups: susceptible, infected, recovered and dead. Once recovered, immunity is permanent.[19] The model is described by the initial value problem:

[18]The first equation says that aggregate income equals aggregate expenditure. The second says that household expenditure equals household income. The third says that the effective amount of labor employed in the market is equal to the fraction of the population going to work (λ) plus the labor input that is coming from working from home. Suppose that the productivity of working from home relative to going to work is Ψ. Since I assume that the (post-tax) income earned by working from home is equal to $\psi_{WFH} \cdot h$, we must have $\Psi A(1-\tau_n)\psi_{WFH} = \psi_{WFH} \cdot h$ and thus $\Psi = \frac{h}{A(1-\tau_n)}$. The second term on the right is thus the share of the population who work from home $(1-\lambda)\psi_{WFH}$ times the relative productivity.

[19]See Eichenbaum *et al.* (2020b) for analysis of reinfections.

$$\dot{S} = -\beta S I \tag{4}$$

$$\dot{I} = \beta S I - \gamma I \tag{5}$$

$$\dot{R} = \gamma_r I \tag{6}$$

$$\dot{D} = \gamma_d I \tag{7}$$

with initial conditions: $S_0 = 1 - \epsilon$, $I_0 = \epsilon$, $R_0 = D_0 = 0$ and with $\gamma = \gamma_r + \gamma_d$.[20]

A certain fraction of infections can be eliminated by costly mitigation behaviors and policies. To capture this idea in a general and flexible way, I assume that the infection rate β is a sum of two components:

$$\beta = \beta_n \lambda_S \lambda_I + \beta_o \tag{8}$$

where λ_S and λ_I are variables that denote the share of susceptible and infected individuals that are active in the labor market and β_n and β_o are parameters.[21] The first term on the right of (8) denotes infections that can be eliminated through behaviors and policies, and the second denotes infections which cannot be eliminated.[22]

The key object in this simple model is the herd immunity threshold defined as:

$$\bar{S} := \frac{\gamma}{\beta} = \frac{1}{\mathcal{R}_0}. \tag{9}$$

[20] I assume that ϵ is small but positive.

[21] Note that this model nests the standard SIR setting with no behavioral feedback under $\beta_n = 0$. Setting $\beta_o = 0$ takes the model to the other extreme, where all infections can be eliminated by mitigation behavior.

[22] Taken literally, the setup assumes that all infections that can be eliminated happen in the workplace. This appears to rule out infections through other economic channels, such as social consumption. But this is with no loss of generality, since, given the hand-to-mouth consumers populating this economy, labor supply and consumption are linked one-to-one. Indeed, a reader should have a broader interpretation in mind: what matters is that the term $\beta_n \lambda_S \lambda_I$ denotes the part of the infection rate that is reducible by costly changes in behavior / policies, and β_o is the part that is irreducible. In practice some infections may be reduced by changes in behavior that are not (significantly) costly – hand-washing or mask-wearing are two examples. Since these behaviors offer a free gain, they are always used fully both in the decentralized equilibrium and by the planner, effectively justifying their omission in (8).

It is a *threshold* because infections are on the rise as long as $I > 0$ and $S > \bar{S}$. It is easy to appreciate the central role played by that threshold by analyzing the model's dynamics with a phase diagram in $S - I$ space (Figure 2).[23] The solid vertical line in the Figure denotes the herd immunity threshold $\bar{S}$. The grey arrows show the direction and the velocity of the system from any initial point under the naive assumption of no behavioral or policy response. An important observation is that the dynamics of the system are fast at the top of the diagram and slow at the bottom (which is easy to see analytically from equation (4)). Finally, the arched line shows the dynamics of the disease in a mechanistic SIR model with no behavioral or policy response (i.e. for $\lambda_S = \lambda_I = 1$). The peak infections occur exactly at the herd immunity point. The dynamics feature an *epidemic overshoot* – a phenomenon whereby a significant portion of the population get infected *after* the herd immunity threshold has been reached.

2.3 Mitigation

Mitigation behaviors or policies reduce $\lambda^i,\ i \in \{S, I, R\}$ to below one. I use the term "lockdown" to mean an extreme version of social distancing that reduces at least one of the λ^i all the way to zero, so that:

$$\beta(t) = \begin{cases} \beta_n \lambda^S \lambda^I + \beta_o & \text{no lockdown at } t \\ \beta_o & \text{lockdown at } t \end{cases}.$$

Graphically a lockdown can be represented by a rightward shift of the vertical line to some $\bar{S}_L > \bar{S}$ and a corresponding change in the system's dynamics (shown in pink in Figure 2).

[23] When drawing the diagram I assume for concreteness that $\mathcal{R}_0$ equals 2.5, which is representative of the many epidemiological studies of the early stages of the Covid-19 pandemic before mitigation policies, but importantly none of the results rely on that specific number.

Figure 2
Dynamics of the epidemic in the analytical model

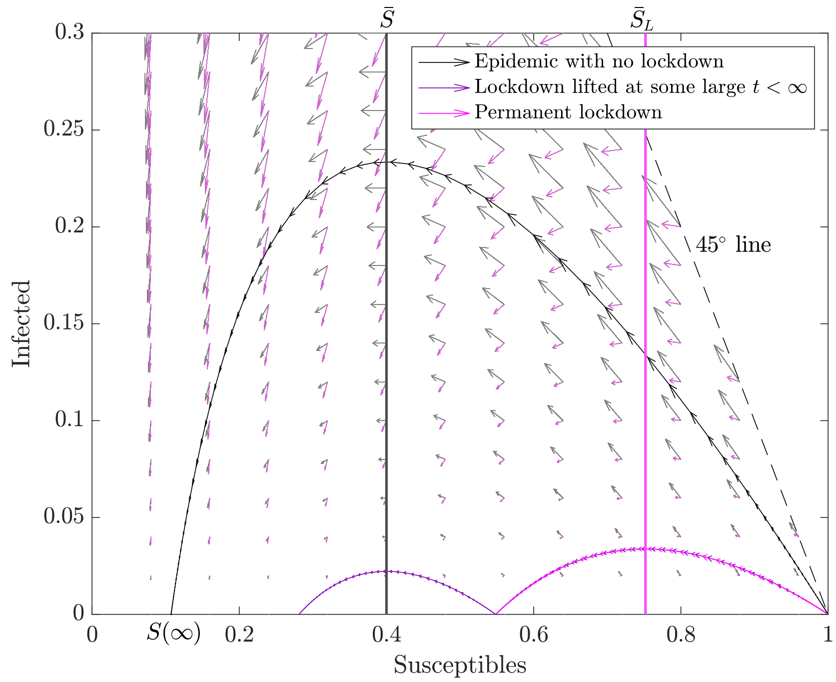

How far to the right the vertical line shifts depends on how effective a full lockdown is. In this paper I assume that mitigation is powerful enough to stop the spread of the disease:

ASSUMPTION 1: $\gamma > \beta_o$.

Assumption 1 implies that a lockdown reduces $\mathcal{R}_0$ to below unity.[24] The available evidence suggests that this is indeed true in the case of Covid: for example, Flaxman *et al.* (2020a) estimate that the lockdown measures implemented in Europe reduced $\mathcal{R}_0$ to below 1 with 99.9% confidence. It is thus fairly clear that Assumption 1 is strongly supported by the data.[25]

[24]By implication, a lockdown implemented at any time reduces the effective reproduction number $\mathcal{R}(t)$, defined as $\mathcal{R}_0 \cdot S$ to below unity.

[25]For analysis under alternative assumptions on lockdown effectiveness, see Rachel (2020).

2.4 The government's role in the epidemic

In normal times the government runs a balanced budget, taxing income and spending the proceeds on its consumption G. I assume that the government is committed to the level of consumption G whatever the circumstances, so that this level of spending must be maintained through the epidemic. During the epidemic the government continues to collect labor income taxes to fund these pre-determined purchases, although tax collections decline as activity falls. It also finances a certain part of income h that all workers who stay at home due to the crisis receive. Lower tax receipts and higher outgoings create a hole in the government budget. I assume that the government borrows the required funds and procures the required output directly in the international market, paying some constant interest rate $\bar{r}$.[26] The government then finances its debt by levying constant lump-sum Corona-tax τ every period on those who survived the virus, starting from some date $\hat{T}$ (after the epidemic has ran its course). Given these assumptions, the government's intertemporal budget constraint is:[27]

$$B_0 \leq e^{-\bar{r}\hat{T}} \int_{\hat{T}}^{\infty} e^{-\bar{r}(t-\hat{T})} \left(S(t) + R(t)\right) \tau \, dt \tag{10}$$

where B_0 is the amount (in units of output) borrowed by the government at time-0 and τ is the lump-sum tax paid by each surviving individual each period. Consequently, from date $\hat{T}$, the budget constraints for the susceptible and the recovered individuals are[28]

[26] A natural – and fair – objection to this assumption is that Covid-19 is a global shock that affects all countries. The assumption here does not need be taken too literally for that reason, and there are other ways to rationalize it. One is to think of the governments as borrowing from the rich (those towards the top of the income and wealth distributions) who are outside of my model (Mian *et al.* (2019a,b)). Another is to think of the central bank as providing the necessary liquidity to the government during the pandemic. Economically what matters is that the government is allowed to borrow today against taxes tomorrow.

[27] I assume that the infected are exempt from paying the tax – this is innocuous since $\hat{T}$ is assumed to be large so that $I(t)$ is minuscule for $t > \hat{T}$.

[28] I verify that in equilibrium $w - \tau > h$ so that all survivors work in the post-epidemic steady state.

$$c_W^R = c_W^S = w - \tau. \tag{11}$$

The amount of fiscal support B_0 responds endogenously to the macro- and epidemiological-environment and to the potential lockdown policies that are adopted during the epidemic:

$$B_0 = \left(\psi_{GOV} h + (1 - \psi_{WFH}) A \frac{\tau_n}{1 - \tau_n} \right) \cdot \\ \cdot \int_0^\infty \left((1 - \lambda_S(t)) S(t) + (1 - \lambda_I(t)) I(t) + (1 - \lambda_R(t)) R(t) \right) dt. \tag{12}$$

The first term in parentheses denotes income paid to those in lockdown, and the second term covers the shortfall in revenues associated with declining tax receipts. Clearly, a longer and broader lockdown will increase the burden on the fiscal authority.

This concludes the description of the environment. I now proceed to the analysis of optimal behavior of individuals that live through an epidemic such as Covid-19.

3 Mitigation in the decentralized equilibrium

I begin with the formal definition of the competitive equilibrium and then proceed to characterization.

Definition 1. A *perfect-foresight competitive equilibrium in mixed strategies* is a sequence of macro variables Y and C, sequence of epidemic variables S, I, R, D, a sequence of social distancing probabilities $\{\lambda^i(t)\}_{i \in \{S,I,R\}} \in [0, 1]$, the level of Corona-tax τ and the level of government borrowing at time-0 B_0 such that: (i) households maximize their expected lifetime utility at time-0 taking the trajectory of the epidemic, behavior of other individuals, wages, government transfers and taxes as given; (ii) firms maximize profits taking wages as given; (iii) government adjusts borrowing and Corona-tax τ to satisfy demand for transfers, meet its spending commitments and satisfy its intertemporal budget constraint; (iv) the trajectory of the epidemic is consistent with the individual lockdown decisions; (v) goods and labor markets clear; (vi) individuals know their health status: they know if they are or had previously been infected.

By the law of large numbers, the equilibrium λ^i is also the share of individuals in health group $i \in S, I, R$ who work. More importantly, λ^i should be interpreted as the index of mitigation efforts, with $\lambda^i = 0$ denoting maximum restrictions and $\lambda^i = 1$ standing for 'business as usual'.

3.1 Equilibrium behavior of the infected and the recovered

To begin note that in equilibrium the infected and the recovered individuals do not lock down: $\lambda^I = \lambda^R = 1$. Since individuals know their health and since there is no altruism, they care only about maximizing their expected utility. For infected or recovered individuals there is no risk of re-infection, and because $u^W > u^L$, zero mitigation is their optimal choice.

3.2 Equilibrium behavior of the susceptibles

Susceptible individuals choose $\lambda^S(t)$ to maximize their expected lifetime utility (I drop the ^{S} subscript for notational convenience):

$$\max_{\{\lambda(t)\}_{t\geq 0}\in[0,1]} \int_0^\infty e^{-\rho t} \left(p_s(t) \left(\lambda(t)u^W + (1-\lambda(t))u^L \right) + p_i(t)u^W + p_r(t)u^W \right) dt \quad \text{subject to} \tag{13}$$

$$\dot{p}_s(t) = -p_s(t)(\beta_n\lambda(t) + \beta_o)I(t)$$
$$\dot{p}_i(t) = p_s(t)(\beta_n\lambda(t) + \beta_o)I(t) - \gamma p_i(t)$$
$$\dot{p}_r(t) = \gamma_r p_i(t)$$
$$\lambda(t) \in [0,1]$$

where p_s, p_i, p_r are the probabilities of being susceptible, infected and recovered at time t, respectively. Their evolution over time mirrors that of the aggregate shares of the different groups in the population but it is driven by individual social distancing choices $\lambda(t)$. Individuals take the number of infected individuals $I(t)$ as given. Once dead, individuals generate zero utility (implicitly there is $(1 - p_s - p_i - p_r) \cdot 0$ term in the maximand), which implies that the value of life is equal to the discounted utility flow.

In the Appendix I setup the Hamiltonian and derive the necessary conditions for the maximum. The following proposition characterizes the solution: the equilibrium voluntary mitigation through the epidemic.

Proposition 1. Equilibrium mitigation. *In the decentralized equilibrium susceptible individuals' mitigation efforts start at $T_0 > 0$ and end at $T_1 < \infty$ so that $\lambda(t) = 1\ \forall t \notin [T_0, T_1]$: there is no mitigation at the beginning and at the end of the epidemic when the infection rate is low. Mitigation severity $\lambda(t)$ at $t \in [T_0, T_1]$ is well approximated by, but more severe than:*

$$\lambda(t) \lesssim \frac{\gamma - \beta_o S(t)}{\beta_n S(t)} \tag{14}$$

which ensures that $\mathcal{R}$ remains close to but below 1. Susceptible individuals dial down the intensity of social distancing over time. Mitigation ends around the time when the herd immunity threshold $\bar{S}$ is reached.

Figure 3
Decentralized equilibrium: phase diagram

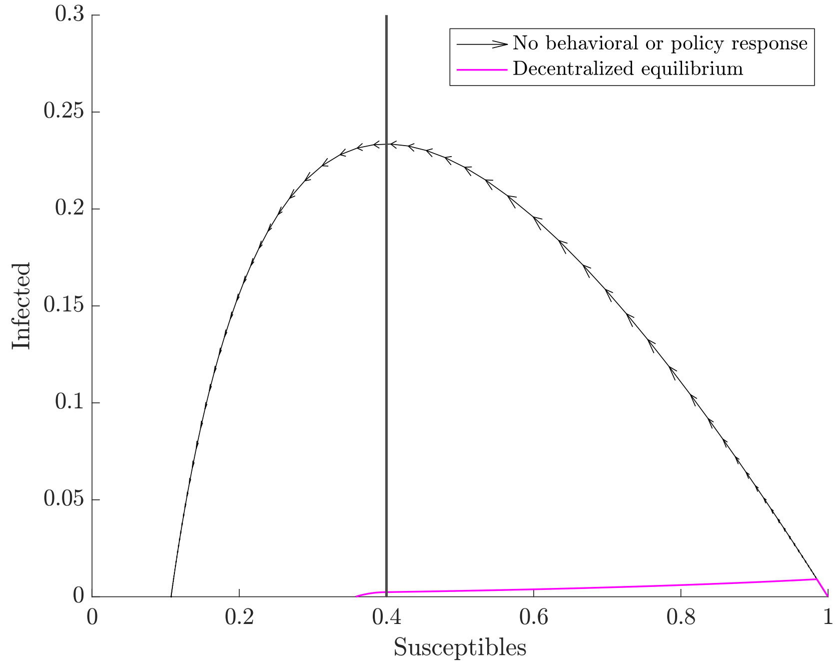

The equilibrium infection rate at any $t \in [T_0, T_1]$ equals

$$I(t) = \frac{u^W - u^L}{\beta_n \left(\eta^S(t) - \eta^I\right)} \tag{15}$$

where $\eta^S(t)$ and $\eta^I = \frac{u^W + \gamma_r \frac{u^W}{\rho}}{\rho + \gamma}$ denote the value functions of the susceptible and the infected, respectively. The equilibrium infection rate peaks at T_0 and decreases over time. If the utility cost of lockdown is not very large then the equilibrium infection rate is well approximated by

$$I(t) \approx S(t) \cdot \frac{\rho \mathcal{C}}{\beta_n \cdot \bar{S} \cdot IFR} \tag{16}$$

where $\mathcal{C} := \frac{u^W - u^L}{u^W}$ is the utility cost of lockdown and $IFR := \frac{\gamma_d}{\gamma}$ is the infection fatality rate.

Proof. See Appendix. □

The first result in Proposition 1 is that the individually-optimal mitigation efforts are discontinuous and non-monotonic in time. There is no social distancing at the beginning and at the end of the epidemic: with low infection rates, the costs of mitigation are greater than the expected health benefits. In the intermittent period mitigation is most aggressive early on. The precautions are gradually relaxed as the epidemic progresses and the pool of susceptible individuals decreases. Figures 3 and 4 provide the graphical illustration on the phase diagram and in the time series.[29]

These mitigation behaviors arrest the rise in infections predicted by the mechanical SIR model. The infection rate peaks exactly at the moment when social distancing measures are first introduced and then follows the downward sloping trajectory in equation (15). Consequently, the effective reproduction number $\mathcal{R}$ is close to but below 1.[30]

Susceptible individuals tolerate a higher infection rate in the early phases because the value of remaining healthy increases over time, making them more averse to the risk of contracting the disease. The reason for this is two-fold. First, the prospective cumulative likelihood of infection (which is approximately equal to $S(t) - \bar{S}$) is high at the start of the epidemic (when $S \approx 1$) and low towards the end (when $S = \bar{S}$). Intuitively, susceptibles at the start of the epidemic think that "they will probably get infected at some point anyway", and that discourages mitigation, making them tolerate a higher infection rate compared to at the end of the epidemic when "it would be a shame to get ill now, just before it is all over". Second, at the start of the epidemic individuals face the prospect of a long lockdown, while towards the end "the return to normality is nigh". Both forces mean that healthy individuals attach a higher value to remaining healthy at the end of the epidemic, and thus continue taking precautions even as the infection rate falls to low levels. This intuition is distinct from that in Farboodi *et al.* (2020), who find that imper-

[29] These Figures are for plotted under the calibration of the model that is described in the Appendix.

[30] In a concurrent paper, Toxvaerd (2020) analyzes a similar problem but arrives at somewhat different conclusions: he argues that the infection rate is constant through the epidemic. What drives the difference between his and my results is that he posits that $\eta^S(t)$ is constant over time. I show this is not the case and solve the model without imposing this simplifying restriction.

Figure 4
Equilibrium behavior flattens the epidemic curve

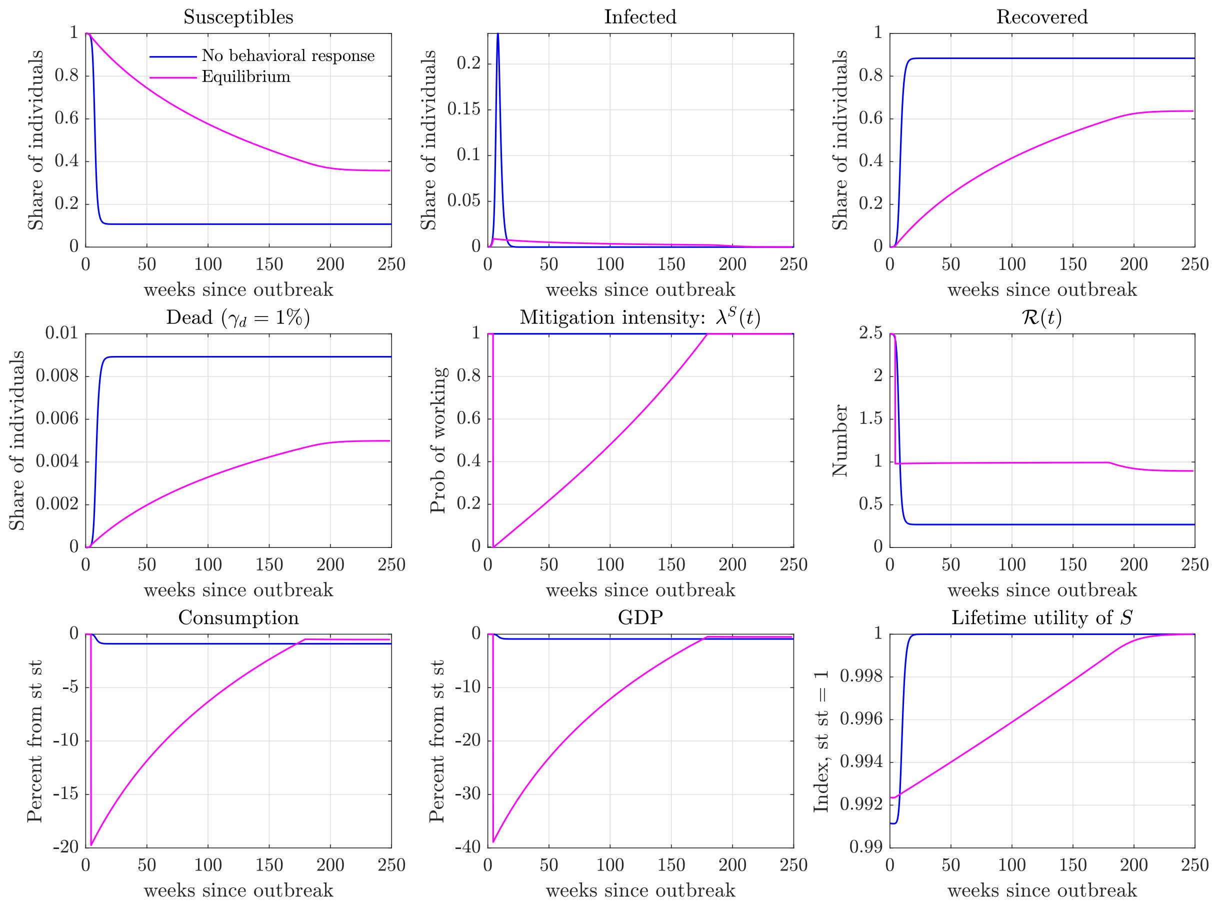

fect information about individual health status can drive a similar dynamic. It is closely related to the "fatalism effect" identified by Jones *et al.* (2020).

Not only the peak but also the long-run spread of the disease is significantly reduced by the voluntary precautionary behavior. Mitigation efforts stop around the time when the herd immunity threshold is reached, and the ultimate resting point is just below $\bar{S}$, meaning that the equilibrium behavior prevents a large proportion of the epidemic overshoot. Thus the often-repeated claim that mitigation measures only "kick the can down the road" is clearly incorrect: voluntary mitigation does save lives, even in the long-run. In fact voluntary social distancing leaves little room for improvement for *any* time-limited mitigation or lockdown in terms of the cumulative death toll in this setting. Note that it is possible to make this claim without computing the optimal mitigation strategy, and simply based on the observation that no such policy can reduce cumulative infections and deaths below the herd immunity threshold.

The approximation to the infection rate in (16) helps tease out comparative statics, which are very intuitive. Individuals tolerate higher infection rates if the cost of mitigation is high (for obvious reasons), if the discount rate is high (as they place higher weight on today's cost of mitigation relative to future benefit), or if the long-run reach of the epidemic is expected to be so widespread that there is little chance of avoiding infection (low $\bar{S}$). Conversely, the tolerable infection rate is low if mitigation is highly effective on the margin (high β_n) and if the disease is very deadly (high IFR).

An important corollary of the low equilibrium infection rate as compared to the mechanistic SIR model is that the transition to the post-epidemic steady state takes longer: if the average duration of the disease is 9 days, then instead of 20 weeks, the epidemic lasts around 200 weeks – a ten-fold increase. Consequently the impact on the economy is significantly more severe and long-lasting in equilibrium, as compared to the mechanistic model.[31]

4 Optimal lockdown

I now study the optimal lockdown strategy that a benevolent social planner would implement at the onset of the epidemic.

[31]The health benefits of individual precautions outweigh the economic cost however, with the value function of a susceptible individual at time-0 higher in the competitive equilibrium than under the "no-behavioral-response" scenario (the final panel in Figure 4).

4.1 Planner's objective and tools

I assume that the objective of the planner is to maximize lifetime utility of the susceptibles at time-0. This is a natural objective, not least because at the onset of the epidemic all people (except a vanishingly small fraction) are susceptible.

In terms of the tools at the planner's disposal I consider four possibilities:

Definition 2. I define four mitigation instruments as follows:

Type 1: *isolation of the infected:* planner sets $\lambda_I(t) \in [0,1]$. $\lambda_S(t) = \lambda_R(t) = 1 \forall t$.

Type 2: *susceptibles-only mitigation:* planner sets $\lambda_S(t) \in [0,1]$. $\lambda_I(t) = \lambda_R(t) = 1 \forall t$.

Type 3: *immunity passports:* planner sets $\lambda_S(t) = \lambda_I(t) \in [0,1]$. $\lambda_R(t) = 1 \forall t$.

Type 4: *all-in mitigation:* planner sets $\lambda_S(t) = \lambda_I(t) = \lambda_R(t) \in [0,1]$.

The tools are ordered from most to least demanding in terms of information available to the planner. Type-1 lockdown is the isolation of the infected: a limit case of a perfectly effective track-trace-isolate strategy. As we shall see, this is an ideal tool to deal with the crisis. The remaining three instruments are what I call the *broad* measures, in that they affect a significant proportion – perhaps a majority – of the population. The results turn out to be remarkably similar across these three broad instruments.[32]

4.2 Optimal isolation of the infected

With the ability to lock down only the infected individuals the planner solves:

$$\max_{\{\lambda(t)\} \in [0,1]} \int_0^{\hat{T}} e^{-\rho t} \left(S(t) u^W + I(t) \left(\lambda u^W + (1-\lambda) u^L \right) + R(t) u^W \right) dt + e^{-\rho \hat{T}} \left(S(\hat{T}) + R(\hat{T}) \right) \frac{u_\tau}{\rho}$$

[32] The susceptible-only mitigation serves as a useful comparison to the competitive equilibrium outcome, in which only the susceptibles isolate. A broad lockdown combined with immunity passports might become the realistic instrument of choice once large-scale antibody testing becomes available. The final instrument – the all-in lockdown – is a blunt one, but it has the crucial advantage that it does not require much information on the part of the planner. For that reason it is also the instrument that has been in use across much of the world in the early parts of the Covid-19 epidemic.

subject to (4), (5), (6), (7), (10) and (12), and taking S_0, I_0, R_0 as given. $u_\tau := u(w-\tau, 1)$ is the post-pandemic instantaneous utility flow. This problem is similar to the decentralized one, with three important differences. First, the individual probabilities p_s, p_i, p_r are replaced by aggregate population shares S, I, R. Second, λ appears only in the flow utility of the infected individuals, since the focus here is on the type-1 policy instrument. And third, the planner takes into account the fiscal implications of the pandemic (hence the final term).

The optimal isolation policy is simple:

Proposition 2. Optimal targeted isolation policy. *If the isolation of the infected is feasible at $t = 0$ and if the infection fatality rate is larger than ρ/β_n (a very small number), it is optimal to implement it immediately and permanently. Such policy prevents the epidemic: $\dot{I} < 0\ \forall t$ and with $S(0)$ vanishingly small, $S(\infty) \approx 1$.*

The intuition behind this result is that benefits of isolating the infected always outweigh the costs. Being infected is a short and temporary state: in the long-run the mass of infected individuals converges to zero. The utility cost of isolation is thus small. However, the health benefits of such policy are very large since this policy stops the epidemic in its tracks. In practice governments around the world might struggle to obtain sufficient information to implement such isolation strategy, at least early on in the epidemic. The result in Proposition 2 underscores the importance of testing programs that allow the policymakers to move in the direction of removing the infected from the general population.

4.3 Optimal broad lockdown policy

I now consider broad lockdown policy tools. To be concrete, consider the problem of a planner who has access to Type-2 S-only mitigation:

$$\max_{\{\lambda(t)\in[0,1]\}} \int_0^{\hat{T}} e^{-\rho t} \left(S(t) \left(\lambda(t)u^W + (1-\lambda(t))u^L \right) + I(t)u^W + R(t)u^W \right) dt + e^{-\rho \hat{T}} \left(S(\hat{T}) + R(\hat{T}) \right) \frac{u_\tau}{\rho}$$

subject to (4), (5), (6), (7), (10) and (12), and taking S_0, I_0, R_0 as given. The only difference from before is that the control λ now enters the flow utility of other groups (the susceptibles and/or the recovered, depending on the lockdown type). The following Proposition characterizes the optimal lockdown policy:

Proposition 3. Optimal broad policy. *If only the broad lockdown instruments are available then optimal policy is characterized by no restrictions at the start and at the end of the epidemic and full lockdown in between. Optimal broad lockdown – irrespective of the specific lockdown type – starts when the level of susceptibility reaches* $S^* = \exp \frac{1+\bar{S}_L \log \bar{S} - \bar{S}}{\bar{S}_L - \bar{S}}$. *The end of the lockdown is determined by the equality between the marginal cost and the marginal benefit of keeping restrictions in place: the broader the lockdown, the sooner it ends. Optimal policy achieves close to the lowest feasible death toll.*

Proof. See Appendix. □

Figure 5 provides a graphical illustration of the results in Proposition 3. The optimal strategy is bang-bang: at any point in time, the planner chooses either full lockdown or no lockdown. The optimal lockdown starts when the level of susceptibility falls to S^*. This ensures that the the epidemic trajectory under lockdown approaches $\bar{S}$ as a resting point. The lockdown ends when S is close to $\bar{S}$, resulting in few extra infections. Thus the optimal policy can be described as a *late, short and strong lockdown*: late because it begins at the relatively high level of infections; short because the dynamics of the epidemic towards

Figure 5
Optimal broad lockdown policy

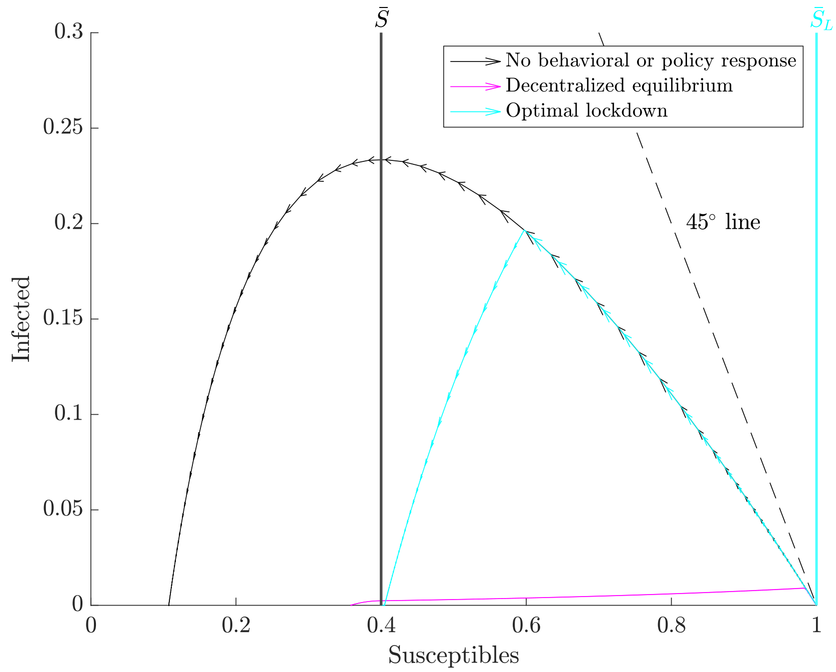

the top of the phase diagram are fast and so the lockdown is relatively short-lived; and strong in the sense that the planner uses the full extent of the lockdown over the entire lockdown episode.

The intuition for why this is an optimal strategy is as follows. Note first that the planner chooses a full-strength lockdown, since there is nothing to gain from gradualism: the minimum level of infections in steady state is $1 - \bar{S}$, and a more gradual approach cannot change that, even as it prolongs the epidemic. Therefore the optimal policy can be characterized by the start and end date of the lockdown episode. Starting the lockdown earlier than at S^* results in more time spent in lockdown and in higher deaths (due to the inevitable second wave when the lockdown is lifted), which cannot be optimal. Starting the lockdown later yields a significantly higher cumulative infections since the system is moving fast when the infection rate is high. Thus the planner will always adjust the end-date of the lockdown, and not the start date, to optimally control lockdown duration.

In fact the lockdown starts at S^* irrespective of what the values of other parameters of the model are. The lockdown end-date clearly does depend on the parameters: in particular, lockdown will be released earlier for broader (and thus more expensive) lockdown types, and if the value of statistical life is lower. But the differences in the resulting epidemic trajectories are small: the system's dynamics are slow around the herd immunity threshold, and the resting point is always close to $\bar{S}$. We thus have the following corollary:

Corollary 1. Independence of economic parameters. *The start date of the optimal broad lockdown is independent of the economic parameters that determine the value of statistical life. And while the end-date of lockdown depends on parameters, the marginal effect of parameter changes on the dynamic of the epidemic are small.*

Taken together, Proposition 3 and Corollary 1 show that the baseline epidemiology model has clear implications for optimal policy, irrespectively of the details of the economic environment. In particular, the optimal action of a policymaker faced with an epidemic is described by Proposition 3 no matter what precise value she attaches to a saved life or how painful economically the lockdown is.[33] The reason behind this striking result is that optimal policy already achieves the lowest possible deaths. In that sense policy is already "at the corner" of the trade-off between deaths and economic losses, and thus changes in the relative weights driven by various parameters have little impact on the optimal response. It is thus not necessary to "take a stand" on the difficult question of the value of statistical life in order to answer the normative questions posed by an epidemic.

A useful way to appreciate the power of the optimal lockdown policy relative to the equilibrium outcome is to look at the time series of the epidemic and macro variables (Figure 6). Optimal policy reduces the overall death toll to its feasible minimum and does so in a way that avoids lengthening the epidemic. In fact the duration of the epidemic

[33]This insensitivity to parameter values operates within the reasonable bounds. Clearly, in the limit, if the cost of the lockdown approaches the cost of death, it is clear that the planner might find it preferable to avoid lockdown altogether.

is comparable to the mechanistic “no behavioral or policy response” model scenario, and yet optimal policy reduces the prevalence of the disease and total number of deaths very significantly relative to that mechanistic outcome. As anticipated above, the reduction in the number of deaths relative to the equilibrium outcome is not very large, since equilibrium mitigation already avoids much of the epidemic overshoot. Macroeconomic losses under optimal policy are sharp, but they are smaller than those incurred in the decentralized equilibrium.

Figure 6

Optimal broad lockdown minimizes the cumulative death toll

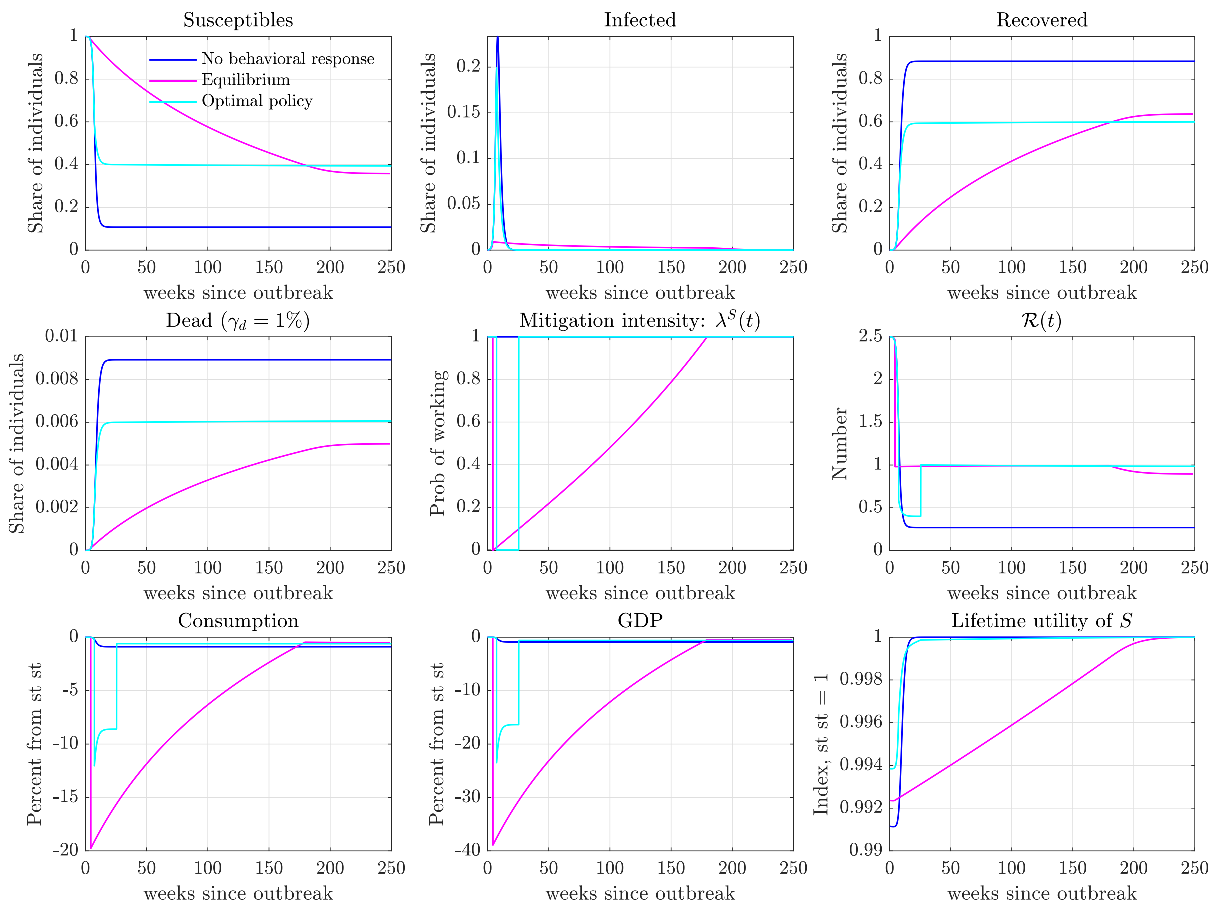

Note: this Figure uses the calibration of the model outlined in the Appendix.

Clearly then, the socially optimal and equilibrium trajectories are very different. In the next Section I analyze the externalities that are underlie those differences.

5 Externalities

There are two externalities present in the decentralized equilibrium. The *infection externality* comes from the fact that each individual takes the current and future economy-wide rates of infection as given, and yet individual decisions do, in equilibrium, drive the economy-wide infection rates. The *fiscal externality* arises because, when deciding on their mitigation strategy, individuals take government transfers and future taxes as given. But individual choices have an impact on government finances: longer or more widespread social distancing leads to higher borrowing and thus higher future taxes.

5.1 Infection Externality

It is useful to consider infection externality separately for individuals by their health status. Starting with the recovered, there is no externality present in their behavior, since they do not engage in any social distancing in equilibrium and the planner finds that optimal: whenever she can distinguish the recovered from others, the planner always chooses to let them roam freely.[34]

The situation is reversed in the case of the infected individuals. Recall that they never lock down in equilibrium, but if the planner can identify them, she always prescribes full isolation (recall Propositions 1 and 2). Thus there is always a negative externality in the behavior of the infected.

[34]In the present model the number of recovered individuals does not affect the number of new infections. This is because of the "quadratic matching" assumption guiding how new infections develop (see equations (4) and (5)), which is the standard assumption in the literature. Acemoglu *et al.* (2020) explore alternative formulations in which the number of the recovered in the population does matter for how many new infections there are. Note, however, that while in that case the "diluting impact" is external to the recovered individuals' choices, there is still no externality: the recovered individuals' choice not to engage in any social distancing is already optimal.

The more subtle and interesting part is the externality in the behavior of the susceptible part of the population. The comparison of the socially optimal susceptible-only lockdown with the equilibrium mitigation in Figure 5 reveals a stark result:

Proposition 4. Infection externality. *The infection externality can mean there is too much, not too little, social distancing in equilibrium. In particular the optimal lockdown starts later and is shorter than individually-optimal social distancing.*

That is, the infection externality is of the opposite sign to what is often asserted. What explains this finding?

Recall that the reason for the infection externality is that individual social distancing impacts on future infection probabilities, and this margin of influence is not taken into account in the decentralized decision problem. The key observation is that the sign of this effect changes depending on the horizon in question: more social distancing today *leads to lower infection rates tomorrow*, but to *higher infections further out*. Intuitively, social distancing efforts that "flatten the curve" also make the curve "fatter". Beyond the reduction in the long-run infection rate due to the epidemic overshoot, mitigation behaviors and policies can only substitute infection risk intertemporally: the cumulative infections must in any case reach at least $1 - \bar{S}$. The planner arrives at this smallest feasible death toll in a way that minimizes the economic disruption. In other words, the planner focuses on the cumulative infection risk over the duration of the epidemic, while individuals focus on limiting the infection risk at each point in time. This piecemeal approach means that individuals are overly cautious and consequently spend a much longer time living with an epidemic.

Proposition 4 thus clarifies the nature of the infection externality present in epidemiology models with endogenous behavior. It shows that the prevailing one-way view of the externalities in the context of an epidemic is incomplete. It is true that the behavior of the infected always generates negative externality in the form of higher infection rates. But externalities that emerge from the behavior of susceptibles are more subtle. In particular, social distancing in the decentralized equilibrium can start before the socially optimal lockdown. This result could be useful when trying to understand why in some countries private choices to mitigate risks were running ahead of the official guidelines.[35]

This result contrasts with some of the previous conclusions in the literature. For example, Jones *et al.* (2020) obtain numerical results which seem to suggest that the planner implements a starker and earlier lockdown compared to the laissez-faire equilibrium. These simulations are obtained in presence of multiple and interacting assumptions, such as the functional form of the costs of mitigation, the assumption about the accumulation of mitigation efforts, and the approximations to the value functions, which may be somewhat difficult to disentangle. The main advantage of the analytical framework developed here is that the economic forces shaping the trajectories of the disease in the decentralized equilibrium and under optimal policy are clear and explicit. This clarity allows for the novel insights into the precise nature of the infection externality.

5.2 Fiscal externality

Fiscal externality is macroeconomic in nature: given the government's commitment to finance a given share of lockdown income up to some fixed replacement rate, each additional day that any individual spends in lockdown carries fiscal consequences. Yet those are not reflected in the problem of atomistic individuals.

Proposition 5. Fiscal externality. *The fiscal externality unambiguously prolongs the equilibrium social distancing relative to the social optimum.*

Proof. Equation (12) implies that government debt B_0 is increasing in the duration of the

[35] Indeed, numerous surveys suggest that a very significant fraction of people would continue to refrain from certain activities even if the official lockdown was lifted. See, for example April 2020 YouGov Survey for the UK and April 2020 Pew Research Center Survey for the US.

lockdown. Government's intertemporal budget constraint (10) implies that future taxes τ are increasing in B_0. The effects of higher taxes are external to individuals' problem because individuals take taxes as given. But higher taxes unambiguously lower long-run continuation values, which reduce the optimal duration of the lockdown, all else equal. $\square$

The logic behind fiscal externality highlights that it can be thought in more general terms as capturing negative macroeconomic effects of lockdowns. For example, it is probable that a lockdown will lead to some deterioration in the stock of social, human, organizational and tangible capital. Such effects, not internalized by individual consumers, workers and firms, would lead the planner to apply a shorter lockdown, relative to the decentralized outcome.

6 Game-changers?

So far I have studied the workings of the baseline analytical model of the Covid-19 epidemic and its implications for the equilibrium and for optimal policy. The model has been kept deliberately simple, abstracting from many features of the real world by design. I now consider feasible suppression of the virus, healthcare capacity constraints, and the possibilities of vaccine or treatment discovery and of temporary immunity, and ask whether these features can have game-changing implications for the conclusions about optimal policy.

6.1 Feasible full suppression

There are two lessons from the baseline model with regards to suppression (defined as keeping the virus in check above the herd immunity threshold $\bar{S}$). The first is that suppression is made *possible* only by restrictive measures being made *permanent*.[36] The second is that suppression is *optimal* only if the isolation of infected individuals – the type-1 mitigation policy above – is feasible, since only that instrument is sufficiently cheap to be in place permanently. How would an alternative assumption with regards to suppression change the analysis?

Feasible full suppression means that there exists a positive threshold of infected $\underline{I}$ below which the virus is eliminated. A natural candidate for such threshold is 1 person: a number below 1 indicates that the last patient has recovered or died, and that no-one (in a given community or society) is no longer a carrier.

Under this alternative assumption a sufficiently long but finite lockdown can fully suppress the virus. This strategy will indeed be optimal as long as the utility cost of the lockdown is not prohibitively high and the required lockdown duration is not excessively long. The length of the required mitigation depends on the effectiveness of the lockdown (Figure 7). For example, if the infection lasts on average for 9 days, starting with a 1% infection rate in the UK – 660k infected – it takes 34 weeks to reduce the number of infections below 1 if the lockdown reduces $\mathcal{R}_0$ to $\frac{1}{2}$, and 62 weeks if the lockdown is less powerful and reduces $\mathcal{R}_0$ to $\frac{3}{4}$. Since the authorities may find it difficult to implement a lockdown of such effectiveness for so long, such full suppression strategy on its own may run into practical feasibility constraints. Lockdowns might need to be supplemented with other types of policies.

[36]There are several reasons why these implications are plausible. First, even the most sophisticated track-and-trace strategies may struggle to identify and isolate all existing cases, given the widespread community transmission already in place in case of Covid-19. Second, even if such strategies were successful domestically in a given country, virus is likely to make a re-entry from abroad once borders open, especially that, to date, 213 countries and territories recorded cases of Coronavirus. Finally, even in the unlikely event that the entire world managed to identify and suppress all human cases, the virus could re-emerge if and when it is transmitted, once again, from animals to humans. Nonetheless, examples of the past viruses, from MERS to Avian Flu, suggest that suppression is indeed possible. It is this possibility that I discuss in this Section.

Figure 7
Suppressing the virus from 1% infection rate

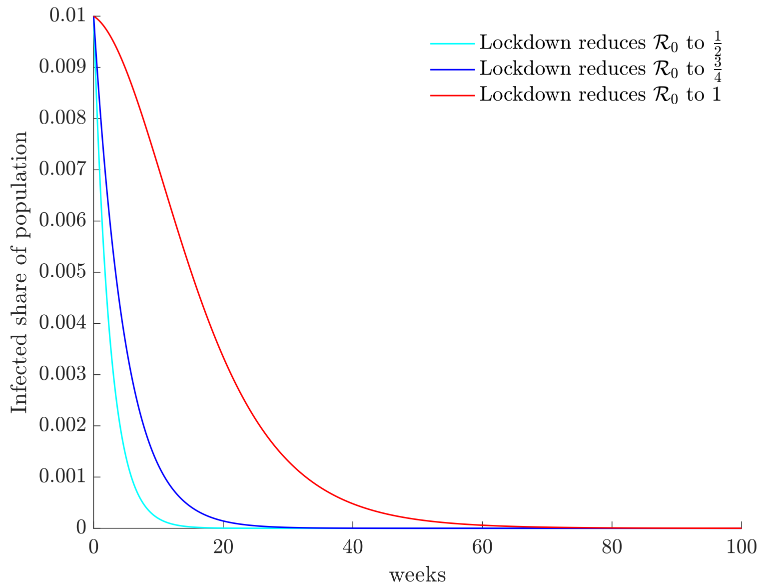

Figure 7 provides a hint of what these policies may be. It shows that the decline in the infection rate slows down over time. This highlights the potential usefulness of policies that can increase the suppression threshold above 1. For example, a test-trace-isolate policy may become feasible when the number of infected is below some threshold $I^{T\&T}$ which is larger than 1 (perhaps it may be equal to 0.1% of the population). Starting from a situation of a widespread transmission in the community, a reasonable strategy is to implement a lockdown until the $I^{T\&T}$ threshold is reached, and then switch to the test-trace-isolate (type-1) strategy. Such policy mix may in practice allow for a feasible suppression at levels of susceptibility well above $\bar{S}$. Whether that strategy is preferred to the one characterized in Proposition 3 will depend on the feasibility, cost and ultimately the effectiveness of the test-trace-isolate strategy, as well as on the herd immunity threshold and the infection fatality rate. Given the estimates of the herd immunity threshold of 40% or less for Covid-19 and the infection fatality rate of between 0.5% and 1%, it is reasonable to expect that the suppression strategy will indeed prove optimal for a wide range of parameter values.

6.2 Healthcare capacity constraints: ICU beds and PPE

One of the main concerns at the onset of the Covid epidemic was that the sudden increase in the number of infections will overwhelm healthcare systems, leading to higher death rate among patients and putting the medical and care staff at risk. The decisions to implement full lockdowns early have been driven, in large part, by this concern. Indeed, even given this endogenous policy response, there is a clear correlation between the infection fatality rate and intensive care unit (ICU) capacity across countries (Figure 8).

Figure 8
Coronavirus deaths and ICU capacity

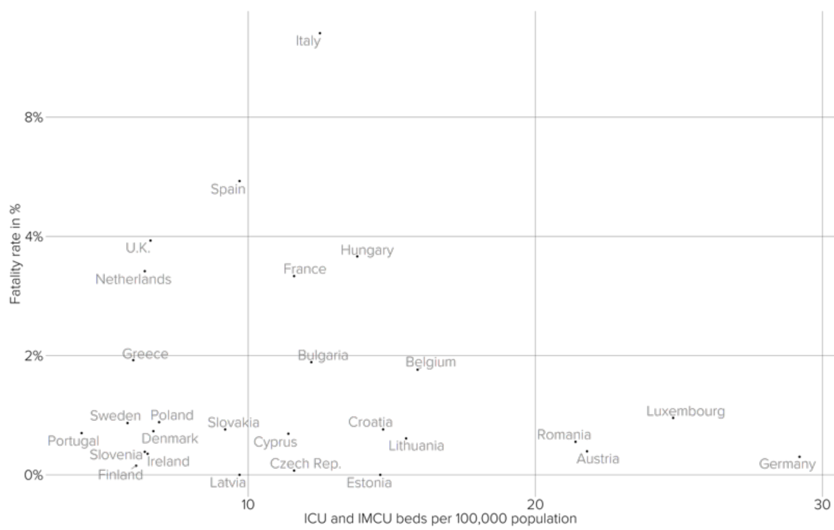

Source: Politico Research, John Hopkins University and Rhodes *et al.* (2012).

A natural way to capture this phenomenon is to depart from the assumption of an exogenous and constant death rate and instead model it a function of the currently infected (this is the approach followed by Eichenbaum *et al.* (2020a); Alvarez *et al.* (2020); Kaplan *et al.* (2020) and others). A particularly simple and appealing approach would be to consider a threshold $\bar{I}$ at which the infection fatality rate jumps up, reflecting lack of appropriate care for some severely ill patients. The analysis in this paper suggests that in this case the optimal policy would aim to use this capacity up to the limit, so that the herd immunity threshold can be reached in minimum time and deaths from the virus are again minimized. The intuition is that in the model with no capacity constraints, the planner strictly prefers infection paths with higher infection rates – the ones that lie further north on the phase diagram. This is because along such a path the disease progresses faster, implying a shorter duration of the epidemic and lower economic costs. The healthcare systems capacity constraints introduce a countervailing force: the planner now avoids high infection rates that are associated with rising mortality. This effect leads to an earlier and lengthier lockdown, aligning the socially optimal and the equilibrium outcomes much closer together. To the extent that only the planner takes these capacity constraints into account,[37] this introduces the *capacity externality* which works to offset the infection externality discussed in Section 5.

6.3 Possibility of a vaccine or a treatment

The baseline model rules out the possibility that a cure or a vaccine will be developed. An effective and widely available vaccine removes the susceptibility out of the population, shifting the system horizontally to the y-axis of the phase diagram (Figure 9). The infection is then eradicated as those with the virus recover or die. The new steady state A is at the origin. An effective treatment "removes" the infected part of the population, shifting the system vertically down to the x-axis of the diagram, making point B a stable steady state (since any newly infected patients can now be immediately treated and cured).

An arrival of a vaccine can modeled to be deterministic or stochastic (or a mixture

[37] It is not clear that this is the case, however. Introspection and casual observation strongly suggest that healthcare systems capacity appears to be a salient feature of the individual choices. Nonetheless, it should be clear from the previous discussion that the planner's optimal path will be affected more than the equilibrium path, because only the latter already achieves the flattening of the epidemic curve.

Figure 9
Dynamics of the epidemic when a vaccine or a treatment are found

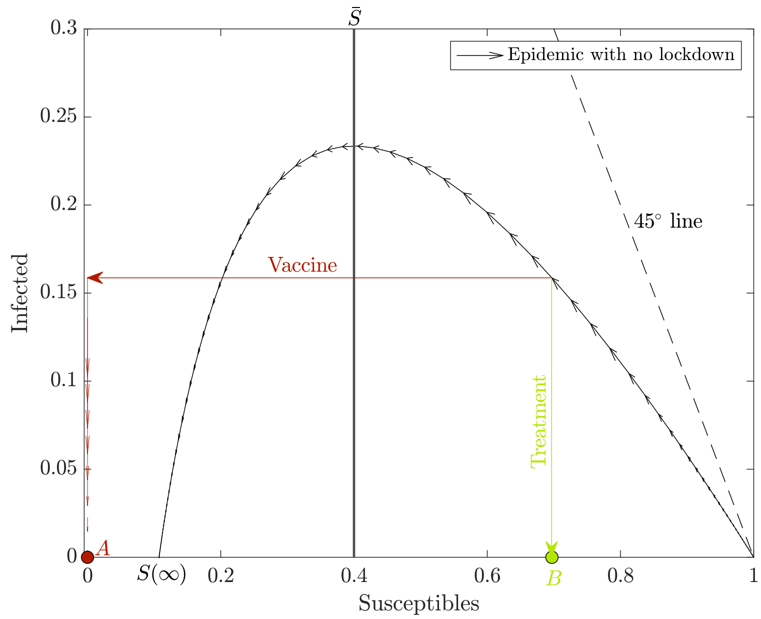

of the two). If a vaccine is certain to appear at a point $t = T_{vacc}$, then the planner's problem boils down to a choice between two strategies: suppression until T_{vacc} and the optimal policy of Proposition 3. There will be a threshold value of T^*_{vacc} at which the planner will be indifferent between the two, so that a vaccine arriving after T^*_{vacc} does not change the conclusions in Proposition 3, while an earlier vaccine leads to suppression being the optimal choice. If the arrival of a vaccine is uncertain, for example if it is a Poisson process, then the planner will front-load mitigation efforts. A positive chance of such discovery effectively shortens the horizon of the planner, making her dislike the paths of the epidemic that are in the upper part of the phase diagram where the increases in cases are particularly rapid. Intuitively, the possibility of a vaccine introduces a positive value to waiting or 'staying put', leading to more front-loaded lockdowns that last for longer.

Figure 10
Dynamics of a naive model with immunity lasting 2 years on average

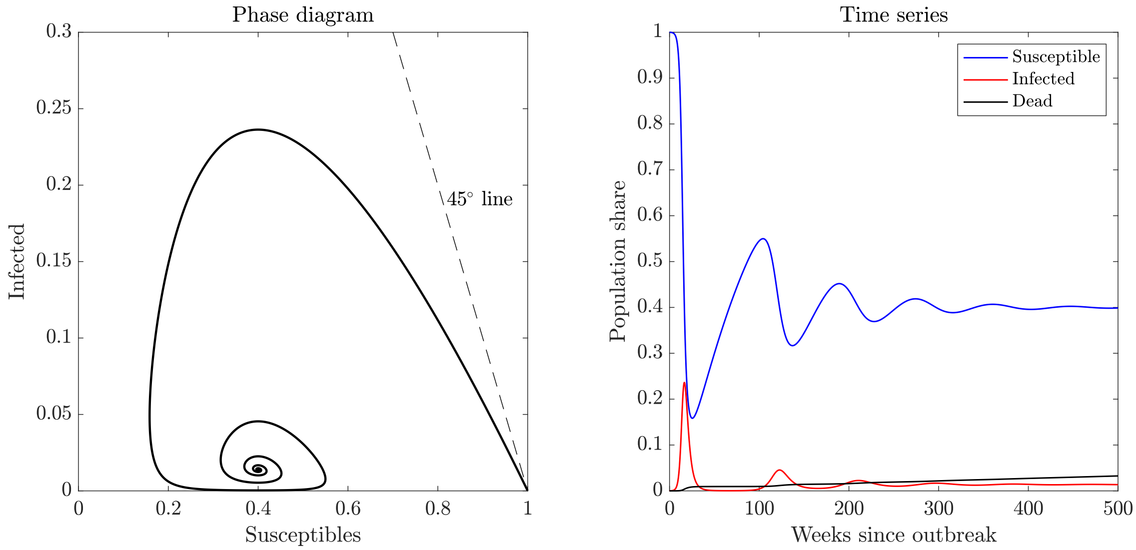

6.4 Reinfection risk

With temporary immunity, the path of the disease under a mechanistic model with no behavioral or policy responses takes the form of a spiral, with multiple waves of progressively smaller magnitudes arriving over time (Figure 10). The infection is never extinguished completely. Reinfection risk effectively introduces multiple and recurring epidemic overshoots. Based on the arguments put forward in this paper, it is easy to see that optimal policy will take a similar form to that characterized in Proposition 3: such policy will aim to avoid the overshoots and instead direct the system at the long-run resting point straight away. The gains from policy can be particularly large, since the cycling of the epidemic trajectory can potentially be avoided with appropriate policies in place.

7 Conclusion

In this paper I characterized analytically the equilibrium mitigation strategy and optimal lockdown policy for the Covid-19 pandemic. I showed that equilibrium mitigation dramatically reduces the peak of infections, but at the price of a significant lengthening of the pandemic. The optimal policy in the baseline model achieves lower overall deaths in a much shorter time frame by delivering a late, short and strong lockdown. The contrast between the optimal policy and the decentralized outcome highlights that the view on the externalities that has dominated the debate is incomplete: in particular, there may be too much social distancing in the decentralized equilibrium, driven by the here-and-now approach to infection risk by the individuals who are susceptible.

While the motivation and the context for this work is the Covid-19 crisis, the conclusions transcend this specific setting, providing economists with a set of analytical tools that can continue to be deployed in the context of the current crisis or any future epidemic event.

References

ACEMOGLU, D., CHERNOZHUKOV, V., WERNING, I. and WHINSTON, M. D. (2020). A Multi-Risk SIR Model with Optimally Targeted Lockdown.

ALVAREZ, F., ARGENTE, D. and LIPPI, F. (2020). A Simple Planning Problem for COVID-19 Lockdown *.

ATKESON, A. (2020). What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios. *NBER Working Paper Series*, p. 25.

BAHAJ, S. and REIS, R. (2020). Central bank swap lines during the Covid-19 pandemic. (2), 1–12.

BERGER, D., HERKENHOFF, K. and MONGEY, S. (2020). An SEIR Infectious Disease Model with Testing and Conditional Quarantine.

CHANG, R. and VELASCO, A. (2020). Economic Policy Incentives to Preserve Lives and Livelihoods.

COCHRANE, J. (2020). Dumb reopening might just work.

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

DINGEL, J. and NEIMAN, B. (2020). How Many Jobs Can be Done at Home? *Becker Friedman Institute White Paper*, (March).

EICHENBAUM, M. S., REBELO, S. and TRABANDT, M. (2020a). The Macroeconomics of Epidemics.

—, — and — (2020b). The Macroeconomics of Testing and Quarantining.

FARBOODI, M., JAROSCH, G. and SHIMER, R. J. (2020). Internal and External Effects of Social Distancing in a Pandemic. *SSRN Electronic Journal*.

FENG, Z. (2007). Final and Peak Epidemic Sizes for Seir Models. *Mathematical Biosciences and Engineering*, **4** (4), 675–686.

FERGUSON, N. M., LAYDON, D., NEDJATI-GILANI, G., IMAI, N., AINSLIE, K., BAGUELIN, M., BHATIA, S., BOONYASIRI, A., CUCUNUBÁ, Z., CUOMO-DANNENBURG, G., DIGHE, A., DORIGATTI, I., FU, H., GAYTHORPE, K., GREEN, W., HAMLET, A., HINSLEY, W., OKELL, L. C., VAN ELSLAND, S., THOMPSON, H., VERITY, R., VOLZ, E., WANG, H., WANG, Y., GT WALKER, P., WALTERS, C., WINSKILL, P., WHITTAKER, C., DONNELLY, C. A., RILEY, S. and GHANI, A. C. (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. *Imperial.Ac.Uk*, (March), 3–20.

FLAXMAN, S., MISHRA, S., GANDY, A., UNWIN, H. J. T., MELLAN, T. A., COUPLAND, H., WHITTAKER, C., ZHU, H., BERAH, T., EATON, J. W., MONOD, M., IMPERIAL COLLEGE COVID-19 RESPONSE TEAM, GHANI, A. C., DONNELLY, C. A., RILEY, S. M., VOLLMER, M. A. C., FERGUSON, N. M., OKELL, L. C. and BHATT, S. (2020a). Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*.

—, —, —, UNWIN, J. T., COUPLAND, H., MELLAN, T. A., ZHU, H., BERAH, T., EATON, J. W., GUZMAN, P. N. P., SCHMIT, N., CILLONI, L., AINSLIE, K. E. C., 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, C. A., RILEY, S., OKELL, L. C., VOLLMER, M. A. C., FERGUSON, N. M. and BHATT, S. (2020b). 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, J. S. (2020). The Economic Consequences of R=1: Towards a Workable Behavioural Epidemiological Model of Pandemics.

GARIBALDI, P., MOEN, E. and PISSARIDES, C. (2020). Modelling contacts and transitions in the SIR epidemics model. *CEPR Covid Economics*, (5), 1–20.

GLOVER, A., HEATHCOTE, J., KRUEGER, D. and RÍOS-RULL, J.-V. (2020). Health versus Wealth: On the Distributional Effects of Controlling a Pandemic *.

GOOLSBEE, A. and SYVERSON, C. (2020). *Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020*. Tech. Rep. 9, National Bureau of Economic Research, Cambridge, MA.

GUERRIERI, V., LORENZONI, G. and STRAUB, L. (2020). Macroeconomic Implications of COVID-19 : Can Negative Supply Shocks Cause Demand Shortages ?

JONES, C. J., PHILIPPON, T. and VENKATESWARAN, V. (2020). Optimal Mitigation Policies in a Pandemic: Social Distancing and Working From Home.

JORDÀ, Ò., SINGH, S. R. and TAYLOR, A. M. (2020). Longer-Run Economic Consequences of Pandemics.

KAPLAN, G., MOLL, B. and VIOLANTE, G. (2020). Pandemics According to HANK extremely preliminary-no stable results yet.

KERMACK, W. O. and MCKENDRICK, A. G. (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.

KNIESNER, T. J. and VISCUSI, W. K. (2019). Legal Studies Research Paper Series The Value of a Statistical Life. *Oxford Research Encyclopedia of Economics and Finance*, (19), 15–19.

LAUER, S. A., GRANTZ, K. H., BI, Q., JONES, F. K., ZHENG, Q., MEREDITH, H. R., AZMAN, A. S., REICH, N. G. and LESSLER, J. (2020). The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application. *Annals of internal medicine*, **2019**.

LAVEZZO, E., FRANCHIN, E., CIAVARELLA, C., CUOMO-DANNENBURG, G., BARZON, L., SCIRO, M., MERIGLIANO, S., DECANALE, E., VANUZZO, M. C., ONELIA, F., PACENTI, M., PARISI, S., CARRETTA, G., DONATO, D., GAYTHORPE, K. A. M. and ALESSANDRA, R. (2020). Suppression of COVID-19 outbreak in the municipality of Vo', Italy. (Ci), 1–23.

LIU, Y., GAYLE, A. A., WILDER-SMITH, A. and ROCKLÖV, J. (2020). The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine*, **27** (2), 1–4.

MIAN, A., STRAUB, L. and SUFI, A. (2019a). Indebted Demand.

—, — and — (2019b). The Saving Glut of the Rich and the Rise in Household Debt *.

PIGUILLEM, F. and SHI, L. (2020). Optimal COVID-19 Quarantine and Testing Policies.

POLLINGER, S. (2020). Optimal Case Detection and Social Distancing Policies to Suppress COVID-19. (23), 1–49.

RACHEL, Ł. (2020). The second wave.

RACHEL, L. and SUMMERS, L. H. (2019). On Falling Neutral Real Rates, Fiscal Policy, and the Risk of Secular Stagnation. *Brookings Papers on Economic Activity*.

RHODES, A., FERDINANDE, P., FLAATTEN, H., GUIDET, B., METNITZ, P. G. and MORENO, R. P. (2012). The variability of critical care bed numbers in Europe. *Intensive Care Medicine*, **38** (10), 1647–1653.

ROGERSON, R. (1988). Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics*, **21** (1), 3–16.

STOCK, J. (2020). Data Gaps and the Policy Response To the Novel Coronavirus. *Nber Working Paper Series*, pp. 1689–1699.

TOMER, A. and KANE, J. W. (2020). How to protect essential workers during COVID-19.

TOXVAERD, F. (2020). Equilibrium Social Distancing.

Appendix (for online publication only)

Proof of Proposition 1

Proof. The Hamiltonian associated with the decentralized problem is:

$$\mathcal{H} = p_s\left(\lambda u^W + (1-\lambda)u^L\right) + p_i u^I + p_r u^W - \eta_s\left(p_s(\beta_n\lambda + \beta_o)I\right) + \\ + \eta_i\left(p_s(\beta_n\lambda + \beta_o)I - \gamma p_i\right) + \eta_r\gamma_r p_i$$

where the costate variables η_s, η_i, η_r can be interpreted as shadow values of being in each of the three health states. Because the Hamiltonian is linear in λ, the optimal control takes the bang-bang or singular form:

$$\lambda(t) = \begin{cases} 0 & \text{if } \psi(t) < 0 \\ \in [0,1] & \text{if } \psi(t) = 0 \\ 1 & \text{if } \psi(t) > 0 \end{cases}$$

where $\psi(t) := u^W - u^L - (\eta_s - \eta_i)\beta_n I$ is the switching function of this problem. By Pontriyagin's Maximum Principle the laws of motion for the co-states and the transversality conditions are:

$$\rho\eta_s - \dot{\eta}_s = \left(\lambda u^W + (1-\lambda)u^L\right) + (\eta_i - \eta_s)(\beta_n\lambda + \beta_o)I \tag{17}$$

$$\rho\eta_i - \dot{\eta}_i = u^W - \gamma\eta_i + \gamma_r\eta_r \tag{18}$$

$$\rho\eta_r - \dot{\eta}_r = u^W \tag{19}$$

$$\lim_{t\to\infty} \eta_s = \eta_r \tag{20}$$

$$\lim_{t\to\infty} \eta_r = \frac{u^W}{\rho} \tag{21}$$

Note that the transversality conditions imply that individuals do not internalize the fiscal costs of mitigation. Conditions (18), (19) and (21) imply:

$$\eta_r = \frac{u^W}{\rho}$$
$$\eta_i = \frac{u^W + \gamma_r \frac{u^W}{\rho}}{\rho + \gamma}.$$

Note that this implies that η_s depends only on the level of infections I. In equilibrium the aggregate epidemic evolves as follows:

$$\dot{S} = (\beta_n\lambda + \beta_0)SI \tag{22}$$

$$\dot{I} = (\beta_n\lambda + \beta_0)SI - \gamma I. \tag{23}$$

This completes the set of conditions that pin down the equilibrium. To characterize the equilibrium path note first that $\psi(t) > 0$ when I is close to zero, implying that there is an interval of time at the start of the epidemic and an interval of time at the end of the epidemic with $\lambda(t) = 1$, that is, with no social distancing. This proves the existence of the equilibrium switching times $T_0 > 0$ and $T_1 < \infty$.

Consider what happens during $t \in [T_0, T_1]$. There are two possibilities: the control can be bang-bang ($\lambda \in \{0,1\}$) or singular ($\lambda \in [0,1]$). Which one it is depends on whether $\psi(t) = 0$ can be sustained over an interval of time. Suppose it can be; then over that interval:

$$(\eta_s - \eta_i) I = \frac{u^W - u^L}{\beta_n} \tag{24}$$

To find the value of singular control note that, since the right hand side of (24) is constant, differentiating both sides with respect to time we get:

$$[(\eta_s \dot{-} \eta_i) I] = 0.$$

We also have:

$$[(\eta_s \dot{-} \eta_i) I] = (\dot{\eta}_s - \dot{\eta}_i)\, I + (\eta_s - \eta_i) \dot{I}$$

Equations (17) and (18) yield

$$\dot{\eta}_s - \dot{\eta}_i = -\left(\lambda u^W + (1-\lambda) u^L\right) - (\eta_i - \eta_s)(\beta_n \lambda + \beta_o) I + \rho \eta_s$$

so that

$$[(\eta_s \dot{-} \eta_i) I] = \left[-\left(\lambda u^W + (1-\lambda) u^L\right) - (\eta_i - \eta_s)(\beta_n \lambda + \beta_o) I + \rho \eta_s\right] I + (\eta_s - \eta_i)\left[(\beta_n \lambda + \beta_0) S I - \gamma I\right].$$

This is equal to zero if

$$-\left(\lambda u^W + (1-\lambda) u^L\right) + \rho \eta_s + (\eta_s - \eta_i)\left[(\beta_n \lambda + \beta_0)(S + I) - \gamma\right] = 0.$$

Using (24):

$$-\left(\lambda u^W + (1-\lambda)u^L\right) + \rho\eta_s + \frac{u^W - u^L}{\beta_n I}\left[(\beta_n\lambda + \beta_0)(S+I) - \gamma\right] = 0.$$

We also have $\eta_s = \frac{u^W - u^L}{I\beta_n} + \frac{u^W + \gamma_r \frac{u^W}{\rho}}{\rho+\gamma}$. Plugging this into the equation above and performing a little algebra yields the expression for λ that is required for the interval over which the control is singular:

$$\lambda^{req} = -\frac{I}{S}\frac{\rho + \gamma_r - \gamma_d \frac{u^L}{u^W - u^L}}{\rho+\gamma} - \frac{\beta_o}{\beta_n}\left(1 + \frac{I}{S}\right) + \frac{\gamma}{S\beta_n} - \frac{\rho}{S\beta_n}. \tag{25}$$

Conjecture that in equilibrium $\frac{I}{S}$ is small. If the discount rate ρ is small, then the following is a good approximation to (25):

$$\lambda^{req} \approx \frac{\gamma}{S\beta_n} - \frac{\beta_o}{\beta_n} = \frac{\gamma - \beta_o S}{S\beta_n} = \frac{\frac{\gamma}{S\beta_0} - 1}{\frac{\beta_n}{\beta_0}} = \frac{\frac{1}{\mathcal{R}^e_{full}} - 1}{\frac{\beta_n}{\beta_0}}$$

where $\mathcal{R}^e_{full}$ is the effective reproduction number after a full lockdown ($\lambda = 0$) is implemented. The singular equilibrium can thus be sustained if $\mathcal{R}^e_{full} \leq 1$. This is indeed the case under Assumption 1, implying $\lambda^{req} > 0$ so that a singular equilibrium path of the control can be sustained.

The stringency of social distancing is decreasing over time since under Assumption 1:

$$\frac{\partial\lambda^{req}}{\partial S} = \frac{-\beta_o\beta_n S - (\gamma - \beta_o S)S\beta_n}{(S\beta_n)^2} < 0$$

and $\dot{S} \leq 0$. The fact that cumulative infection risk declines over time (since S declines towards $S(\infty)$) and the fact that the stringency of social distancing decreases over time imply that η_s increases over time in the interval $t \in [T_0, T_1]$.

To derive the approximation to the infection rate, start with

$$\eta_s \approx p\eta_i + (1-p)\eta_r \tag{26}$$

where p is the cumulative probability of getting infected in the future. This is an approximation since it ignores the expected utility cost of lockdown. Denoting with S_∞ the long-run level of susceptibility in equilibrium, we can approximate p as:

$$p = \frac{S - S_\infty}{S} \approx \frac{S - \bar{S}}{S} = 1 - \frac{\bar{S}}{S}$$

which conjectures that the equilibrium outcome avoids much of the epidemic overshoot (the approximation will be accurate if this is the case in equilibrium). Also, we can approximate $\eta_i = \eta_r \left(1 - \frac{\gamma_d}{\rho+\gamma}\right) \approx \eta_r(1 - IFR)$. Combining these three approximations we get

$$\eta_s - \eta_i \approx \frac{S}{\bar{S}}\eta_r IFR.$$

Since $\eta_r = \frac{u_W}{\rho}$ equation (24) implies

$$I(t) = \frac{u^W - u^L}{\beta_n(\eta_s - \eta_i)} \approx S \cdot \frac{\rho\mathcal{C}}{\beta_n \cdot \bar{S} \cdot IFR}.$$

□

Proof of Proposition 2

Proof. The Hamiltonian associated with this problem is:

$$\mathcal{H} = Su^W + I\left(\lambda u^W + (1-\lambda)u^L\right) + Ru^W - \eta_s\left((\beta_n\lambda + \beta_o)SI\right) + \eta_i\left((\beta_n\lambda + \beta_o)SI - \gamma I\right) + \eta_r\gamma_r I$$

Since the Hamiltonian is again linear in the control, the solution is bang-bang or singular, with the switching function $\psi(t) = u^W - u^L - (\eta_s - \eta_i)\beta_n S$. The evolution of the costates is:

$$\rho\eta_s - \dot{\eta}_s = u^W + (\eta_i - \eta_s)(\beta_n\lambda + \beta_o)I \tag{27}$$

$$\rho\eta_i - \dot{\eta}_i = \lambda u^W + (1-\lambda)u^L - (\eta_s - \eta_i)(\beta_n\lambda + \beta_o)S - \gamma\eta_i + \gamma_r\eta_r \tag{28}$$

$$\rho\eta_r - \dot{\eta}_r = u^W \tag{29}$$

$$\eta_s(\hat{T}) = \eta_r(\hat{T}) = \frac{u_\tau}{\rho} \tag{30}$$

Guess and verify that immediate and permanent lockdown is optimal: $\lambda = 0 \forall t$. Consequently, $\psi(t)$ is negative for all t. Given Assumption 1, an immediate suppression of the virus means that $S \approx 1 \forall t$, which means that $\eta_s \approx \frac{u^W}{\rho}$ and $\eta_i \approx \frac{u^W + \gamma_r \frac{u^W}{\rho}}{\rho+\gamma} = \frac{\frac{u^W}{\rho}(\rho+\gamma_r)}{\rho+\gamma} = \eta_s\left(1 - \frac{\gamma_d}{\rho+\gamma}\right)$. The term $\frac{\gamma_d}{\rho+\gamma}$ is approximately equal to, but is slightly smaller than, the infection fatality rate. Our guess implies:

$$u^W - u^L < \frac{u^W}{\rho} IFR\beta_n$$

which yields the following condition on the parameters for the immediate and full lockdown strategy to be optimal:

$$IFR > \left(1 - \frac{u_L}{u_W}\right)\frac{\rho}{\beta_n}$$

Therefore a sufficient (but not necessary) condition for full and immediate lockdown to be optimal is:

$$IFR > \frac{\rho}{\beta_n}.$$

Note that this is a very weak condition. For Covid-19 if $\mathcal{R}_0 = 2.5$ and lockdown reduces $\mathcal{R}_0$ to or below 1, then $\beta_n > 1.2$. With the annual discount rate of 4%, the right hand side is smaller than 0.00067 – several orders of magnitude smaller than the infection fatality rate of Covid, which is estimated to be around 0.5-1%. $\square$

Proof of Proposition 3

Proof. The Hamiltonian is:

$$\mathcal{H} = S\left(\lambda u^W + (1-\lambda)u^L\right) + Iu^W + Ru^W - \eta_s\left((\beta_n\lambda + \beta_o)SI\right) + \eta_i\left((\beta_n\lambda + \beta_o)SI - \gamma I\right) + \eta_r\gamma_r I$$

Once again, because the Hamiltonian is linear in λ, the optimal control will be bang-bang or singular. The switching function is now:

$$\psi(t) = u^W - u^L - (\eta_s - \eta_i)\beta_n I$$

which looks similar to that in the decentralized problem. However, η_i is now time-varying. Specifically the behavior of the costates is given by the following system of equations:

$$\rho\eta_s - \dot{\eta}_s = \left(\lambda u^W + (1-\lambda)u^L\right) + (\eta_i - \eta_s)(\beta_n\lambda + \beta_o)I \tag{31}$$

$$\rho\eta_i - \dot{\eta}_i = u^W + (\eta_i - \eta_s)(\beta_n\lambda + \beta_o)S + \gamma_r(\eta_r - \eta_i) - \gamma_d\eta_i \tag{32}$$

$$\rho\eta_r - \dot{\eta}_r = u^W \tag{33}$$

$$\eta_s(\hat{T}) = \eta_r(\hat{T}) = \frac{u_\tau}{\rho} \tag{34}$$

Can $\psi(t) = 0$ be sustained over an interval? Suppose yes. Then over that interval:

$$(\eta_s - \eta_i)I = \frac{u^W - u^L}{\beta_n}$$

From the two adjoint equations we have

$$\dot{\eta}_s - \dot{\eta}_i = (1-\lambda)(u^W - u^L) + (\eta_s - \eta_i)(\beta_n\lambda + \beta_o)(I - S) + \gamma_r(\eta_r - \eta_i) - \gamma_d\eta_i.$$

Combining this with the state equation we have:

$$(\dot{\eta}_s - \dot{\eta}_i)\, I + (\eta_s - \eta_i)\dot{I} = (1-\lambda)(u^W - u^L) + (\eta_s - \eta_i)(\beta_n\lambda + \beta_o)I + \gamma_r\eta_r - \gamma\eta_s.$$

which implies the following must hold along the interval:

$$(1-\lambda)(u^W - u^L) + \frac{u^W - u^L}{\beta_n}(\beta_n\lambda + \beta_o)I + \gamma_r\eta_r - \gamma\eta_s = 0.$$

Rearranging this equation gives the value of control that is required along the singular interval:

$$\lambda^{req} = \frac{1 + \frac{\beta_0}{\beta_n}I - \frac{\eta_s\gamma - \eta_r\gamma_r}{u^W - u^L}}{1 - I}. \tag{35}$$

The term $\frac{\eta_s\gamma - \eta_r\gamma_r}{u^W - u^L}$ is approximately equal to $\frac{\eta_s\gamma_d}{u^W - u^L}$, that is, it is the ratio between expected utility loss upon infection and death relative to the instantaneous utility cost of lockdown. As long as lockdown is not "worse than death", this term is always positive and large: for example, if the probability of death given disease is 1% and the discount rate is 4% per annum, this term is greater than 300 throughout the epidemic. Together with Assumption 1 this means that the numerator in (35) is negative, and $\lambda^{req} < 0$: a contradiction since only $\lambda \in [0, 1]$ is feasible. It follows that the singular control cannot be sustained over an interval, and thus the optimal policy is bang-bang.

To see that the optimal lockdown starts at $S^* = \exp \frac{1+\bar{S}_L \log \bar{S} - \bar{S}}{\bar{S}_L - \bar{S}}$ note first that such lockdown results in the epidemic trajectory that approaches $(\bar{S}, 0)$ (see Rachel (2020) for more on this result). The proof for why this must be the optimal start date of the lockdown is as follows: suppose the lockdown starts earlier. Then the resulting death toll is higher, and the lockdown is longer, which cannot be optimal. Suppose the lockdown starts later. Then for a given delay Δt, the resulting new infections would be higher than if the lockdown end date was adjusted by Δt instead (since $\dot{S} = -\beta S I$ which increases in I). Thus the adjustment to the length of the lockdown is always done through the lockdown end-date, rather than the start date. $\square$

Calibration of the model

The table below outlines the calibration of the model used to produce the figures in the main text. In addition, I specify $u(c, n) = \log(c)$.

Description		Value	Target / Source
Macro parameters			
ρ	Discount rate, annualized	4%	$VSL \approx$ \$10m, Kniesner and Viscusi (2019) ...
$\bar{r}$	Government borrowing cost	1%	Rachel and Summers (2019)
A	Wage rate	1153	Target \$60k income per capita
$\frac{h}{w}$	Income replacement rate	80%	Policies in several countries
ψ_{WFH}	Working from home share	1/2	Dingel and Neiman (2020)
ψ_{HPR}	Home production share	1/6	Tomer and Kane (2020)
ψ_{GOV}	Govt transfer share	1/3	Davies (2020)
τ_n	Tax rate in normal times	34%	OECD average
Epi parameters			
$\mathcal{R}_0$	Basic reproduction number	2.5	Flaxman *et al.* (2020b), Lavezzo *et al.* (2020) ...
γ_r	Recovery intensity	$0.99 \cdot \frac{7}{10}$	Avg duration: 10 days (Lauer *et al.* (2020), Liu *et al.* (2020))
γ_d	Death intensity	$0.01 \cdot \frac{7}{10}$	Mortality rate of 1%
ε	Lockdown effectiveness	0.6	$\mathcal{R}_0 \downarrow$ to 1, e.g. Flaxman *et al.* (2020b) ...