

WORKING PAPER · NO. 2020-34

A Simple Planning Problem for COVID-19 Lockdown

Fernando Alvarez, David Argente, and Francesco Lippi

APRIL 2020

A Simple Planning Problem for COVID-19 Lockdown *

Fernando Alvarez

University of Chicago and NBER

David Argente

Pennsylvania State University

Francesco Lippi

LUISS and Einaudi Institute for Economics and Finance

April 6, 2020

PRELIMINARY DRAFT

Abstract

We study the optimal lockdown policy for a planner who wants to control the fatalities of a pandemic while minimizing the output costs of the lockdown. We use the SIR epidemiology model and a linear economy to formalize the planners dynamic control problem. The optimal policy depends on the fraction of infected and susceptible in the population. We parametrize the model using data on the COVID19 pandemic and the economic breadth of the lockdown. The quantitative analysis identifies the features that shape the intensity and duration of the optimal lockdown policy. Our baseline parametrization is conditional on a 1% of infected agents at the outbreak, no cure for the disease, and the possibility of testing. The optimal policy prescribes a severe lockdown beginning two weeks after the outbreak, covers 60% of the population after a month, and is gradually withdrawn covering 20% of the population after 3 months. The intensity of the lockdown depends on the gradient of the fatality rate as a function of the infected, and on the assumed value of a statistical life. The absence of testing increases the economic costs of the lockdown, leads to worse welfare outcomes and shortens the duration of the optimal lockdown.

*First draft, March 23, 2020. We benefited from the comments of Andrew Atkeson, Gadi Barlevy, Mike Golosov, Fausto Gozzi, Francois Gourio, Lars Hansen, Kiminori Matsuyama, Magne Mogstad, Casey Mulligan, Tom Phelan, Filip Rozsypal, Fabiano Schivardi, Rob Shimer, Daniele Terlizzese, Fabrice Tourre, Marcelo Veracierto, Ivan Werning, and panelists and participants on the HELP! (Health and Pandemics Economics Group) seminar on March 27th, the World Bank's Development Policy and Covid-19 e-seminar on April 1st, and the Federal Reserve Bank of Chicago Virtual Macro Seminar on April 3rd. The authors declare to have no conflict of interest to disclose regarding the research on this paper.

1 Introduction and Overview

We present a simple optimal control model for the COVID19 epidemic. We adopt a variation in the SIR epidemiology model reviewed and proposed by [Atkeson \(2020\)](#) and [Neumeyer \(2020\)](#) to analyze the optimal lockdown policy. Our aim is to contribute to the ongoing discussion on the optimal policy response to the COVID shock, see [Barro, Ursua, and Weng \(2020\)](#); [Eichenbaum, Rebelo, and Trabandt \(2020\)](#); [Hall, Jones, and Klenow \(2020\)](#); [Dewatripont et al. \(2020\)](#); [Piguillem and Shi \(2020\)](#); [Jones, Philippon, and Venkateswaran \(2020\)](#) and the contributions in the volume by [Baldwin and Weder \(2020\)](#).

The typical approach in the epidemiology literature is to study the dynamics of the pandemic, for infected, deaths, recovered, as functions of some exogenously chosen diffusion parameters, which are in turn related to various policies, such as the partial lockdown of schools, universities, businesses, and other measures of diffusion mitigation, and where the diffusion parameters are stratified by age and individual covariates. This is the approach followed for instance by [Ferguson and Et.al \(2020\)](#). We use a simplified version of these models proposed by [Atkeson \(2020\)](#) to analyze how to optimally balance the fatality induced by the epidemic with the output costs of the lockdown policy. The novel aspect of our analysis is to explicitly formulate and solve a control problem, where the diffusion parameter is chosen to maximize a given social objective and taking into account the dynamic nature of the problem.¹ By computing the optimal policy, and the associated trajectories, we aim to gauge the intensity and duration of the optimal COVID lockdown, and to understand what are the key elements that determine it. A reason to write a planning problem directly is that with social interactions there is an externality to be corrected, as understood in much of the search literature and as carefully analyzed in [Eichenbaum, Rebelo, and Trabandt \(2020\)](#) and [Toxvaerd \(2020\)](#).

The planner's objective is to minimize the present discounted value of fatalities while

¹An optimal control problem based on a very similar epidemiological model can be found in [Hansen and Troy \(2011\)](#), but the objective function and the feasible policies are different.

also trying to minimize the output costs of the lockdown policy, since agents in lockdown are assumed not to produce. Underlying our setup is a SIR epidemiology model, describing how the virus propagates from the Infected agents to those Susceptible of falling ill, as well as the rates at which, once infected, agents either recover or die. We formulate and solve the problem of a planner who has access to a single instrument to deal with the epidemic: the lockdown of the citizens.² The planner's problem features a tradeoff between the output cost of lockdown, which are increasing in the number of susceptible and infected agents, and the fatality cost of the epidemic. The state of the problem is two dimensional and, in spite of its simplicity, it does not have an analytic solution. We study how the optimal *intensity* and *duration* of the lockdown depend on the cost of fatalities, as measured by the value of a statistical life, on the effectiveness of the lockdown (the reduction in the number of contacts once the citizens are asked to stay home), and on the possibility of testing, i.e. to identify those who acquired immunity to the disease. We show that if the fatality rate (probability of dying conditional on being infected) is increasing in the number of infected people, as is likely the case once the hospital capacity is reached, the policy maker motive for lockdown is strengthened. But, an optimal gradual lockdown may also arise in the absence of non-linear costs, since the costs and benefits of the policy depend on the state of the system.

We parametrize the model using a range of estimates about the COVID19 epidemic which are obviously uncertain. We explore several variations of our benchmark scenario, including the severity of the congestion effects on the fatality rate, a range of valuations for the value of lost lives, and the possibility of testing and releasing the recovered agents from lockdown. We illustrate these points using a quantitative analysis, accompanied by simulations, summary tables of the main results, and phase diagrams that allow us to study the problem under any initial condition.

The quantitative results are useful to gauge what parameters of the problem are important

²While the lockdown is not the only margin of action (other actions might involve reinforcing health treatment capacity and incentivizing the development of vaccines), in the short run this seems to be an important policy tool available and used by several countries.

in shaping the intensity and duration of the optimal lockdown policy. Several comparative statics analysis are developed. In our baseline parameterization, conditional on a 1% fraction of infected agents at the outbreak, the possibility of testing and no cure for the disease, the optimal policy prescribes a lockdown starting two weeks after the outbreak, covering 60% of the population after 1 month. The lockdown is kept tight for about a full month, and is subsequently gradually withdrawn, covering 20% of the population 3 months after the initial outbreak. The output cost of the lockdown is high, equivalent to losing 8% of one year's GDP (or, equivalently, a permanent reduction of 0.4% of output). The total welfare costs is almost three times bigger due to the cost of deaths (see [Figure 1](#) and [Table 2](#)).

These results are based on a relatively pessimistic parameterization of the fatality rate, and on the fraction of the population that would have been infected if there was no lockdown. We show that the intensity of the optimal lockdown depends critically on the gradient of the fatality rate as a function of the infected. If we assume that the fatality rate is constant, the intensity and duration of the lockdown are significantly reduced and, in some cases, completely eliminated, even though the welfare cost of the pandemic remain high. On the other hand, the value of the statistical life we use in our benchmark case (20 times annual GDP per capita) is on the low range of the estimates in the literature. Following [Hall, Jones, and Klenow \(2020\)](#), our benchmark value takes into account that the majority of the victims of the virus have a below average life expectancy. A higher value of statistical life (say 30 times annual GDP per capita), makes the abandonment of the lockdown more gradual, taking a bit more than six months to be totally abandoned. Considering a much larger value, in the order of 80 to 140 times the annual GDP per capita, implies a very strict lockdown that lasts for about 9 months, and a year after still has about 15% of the population in lockdown. Finally, our benchmark scenario assumes that there is a test that allows those that recover to go back to work, so that they are not subject to the lockdown. In the absence of such a test the optimal lockdown is shorter, but it involves roughly the same total number of hours lost due to the lockdown (see [Figure 6](#) and [Table 2](#)).

A byproduct of the calculations is the benefit of the lockdown policy, measured as a percentage permanent GDP flow of following the optimal policy vs the case of no lockdown (see [Table 2](#)). For instance, for the range of our preferred values, the total welfare cost of the virus is equivalent to a loss of 30% to 40% of one year GDP. From this loss, the part due to lockdown of workers is between 8% and 12% of one year GDP. We also use these measures to analyze the importance of different features. This confirms that the elasticity of the fatality rate to the number of infected is a key determinant of the optimal policy –we found that when the fatality rate is flat the optimal policy is to have no lockdown. As mentioned above, we also analyze the sensitivity of the optimal policy to a large range of values of the a statistical life, and to the unavailability to have a test to avoid locking out immune workers.

Needless to say the analysis has limitations: the underlying model has no heterogeneity in fatality rates nor in diffusion rates, the lockdown policy cannot be differentiated across agent’s type (e.g. young versus old, workers vs retirees). We also ignore direct health interventions that might be put in place to mitigate the consequences of the disease (building emergency hospitals, incentivizing the research and cure for the disease). Finally, we ignore the long-term economic consequences of a temporary lockdown, a potentially important effect, whose analysis requires adding one state to the problem.

Our objective is similar to that pursued by [Eichenbaum, Rebelo, and Trabandt \(2020\)](#). While they focus on a competitive equilibrium where a consumption tax is used to slow-down economic activity as well as the epidemic diffusion, we focus on a simpler planner problem that at each time directly mandates the level of lockdown to its citizen. In our setup, the interaction of the law of motion coming from the SIR model and the control of lockdown makes the problem non-convex, which requires the use global methods. Another recent contribution addressing the optimal control problem in the presence of contagion externalities can be found in [Jones, Philippon, and Venkateswaran \(2020\)](#).

The paper is organized as follows: the next section describes the planner’s problem and the epidemic model. [Section 3](#) uses some micro data to calibrate the key model parameters

drawing from [Atkeson \(2020\)](#); [Ferguson and Et.al \(2020\)](#). [Section 4](#) discusses the preliminary results of the optimal control problem under different scenarios.

2 A planner model of lockdown control

We start with a modified version of the, by now well known, SIR model as described in [Atkeson \(2020\)](#). Agents can be divided between those susceptible to be infected $S(t)$, those infected $I(t)$, and those recovered $R(t)$, i.e.

$$N(t) = S(t) + I(t) + R(t) \text{ for all } t \geq 0 \quad (1)$$

Typically, in the SIR model those “recovered” include those that have been infected and are now assumed to be immune and those have died. In our case, we will only include those that are alive, so $N(t)$ will be changing through time, by subtracting those that die. We normalize the initial population to $N(0) = 1$. The planner can decide to lockdown a fraction $L(t) \in [0, \bar{L}]$ of those susceptible and those infected, where $\bar{L} \leq 1$ allows us to realistically consider that even in a disaster scenario some economic activity such as energy and basic food production have to continue. Recovered agents are not in lockdown. A fraction of those that are in lockdown cannot be infected, nor can infect some other susceptible agents. We assume that the lockdown is only partially effective in eliminating the transmission of the virus. When L agents are in lockdown, then $(1 - \theta L)$ agents can transmit the virus, where $\theta \in (0, 1]$ is a measure of the lockdown effectiveness. If $\theta = 1$, the policy is fully effective in curbing the diffusion, but since some contacts will still happen in the population even under a full economic lockdown, we allow $\theta < 1$.

The law of motion of the susceptible agents is:

$$\dot{S}(t) = -\beta S(t)(1 - \theta L(t)) I(t)(1 - \theta L(t)) \quad (2)$$

where β is the number of susceptible agents per unit of time to whom an infected agent can transmit the virus, among those not in lockdown. All susceptible agents that get the virus become infected. For the infected, a fraction γ recovers, thus:

$$\dot{I}(t) = \beta S(t)(1 - \theta L(t)) I(t)(1 - \theta L(t)) - \gamma I(t) \quad (3)$$

A rate $0 < \phi(I) \leq \gamma$ per unit of time of those infected die. Thus, population decreases due to death among those infected as:

$$-\dot{N}(t) = D(t) = \phi(I(t)) I(t) \quad (4)$$

While we assume that the rate γ at which infected recover is exogenous, the fraction of those recovering that die is assumed to depend on the number of infected I , given by the function ϕ . Note that in this epidemiological model, locking down a part of the population, while economically costly, can be very powerful to reduce the rate at which susceptible agents become infected. This is because in the standard model, it is the product of the infected and susceptible that determine the new infections per unit of time. Hence, decreasing the number of contacts of each, decreases the new infections by its *square*. This completes the description of the epidemiological block of the model.

We assume that each agent alive produces w units of output, when she is not in lockdown. Agents are assumed to live forever, unless they die from the infection.³ The planner discounts all values at the rate $r > 0$. We assume that with probability ν per unit of time both a vaccine and a cure simultaneously appear, so that all infected are cured and all susceptible become immune, ie. all infected and susceptible become recovered. The planner seeks to maximize:

$$\int_0^\infty e^{-(r+\nu)t} \left((N(t) - [S(t) + I(t)] L(t)) w + \dot{N}(t) \chi + \frac{\nu}{r} N(t) w \right) dt \quad (5)$$

³We introduce a parameter to correct for this in the calculation of the value of a statistical life.

by choice of the control $L(t)$ at all times, i.e. choosing the fraction of agents to lockdown $L(t)$. In the baseline case we assume that those that have recovered from the virus can be identified, say with an antibody test, and don't need to be locked down; we also explore the case where such a test is not available. The objective function includes the present value of the output of those alive, as well as an extra cost χ , in units of output, for each agent that dies as a consequence of the virus above and beyond the loss output. The term $\frac{\nu}{r}N(t)w$ is the product of the probability that a cure takes place ν times the discounted value of output from there on, $wN(t)/r$. The planner's problem is subject to the law of motion of the susceptible individuals [equation \(2\)](#), the infected [equation \(3\)](#), the population, [equation \(4\)](#), and an initial condition $(N(0), I(0), S(0))$ with $I(0) > 0$ and $S(0) + I(0) \geq N(0)$. Below we give an equivalent, perhaps easier to interpret, formulation of the planner's objective function in terms of the output cost of lock down and the cost of fatalities evaluated using the value of a statistical life.

Discussion of modeling assumptions We now briefly review the key assumptions of the model:

1. We restrict the extent of the lockdown to $\bar{L} \leq 1$. This is to take into account, in a very rough manner, that some sectors cannot shut down even with the most severe lockdown, such as health, basic services, food production and distribution, etc.
2. If $\theta = 1$ the lockdown is able to completely stop the infections process, i.e. to achieve $\dot{S} = 0$ (at $L = 1$). If $\theta < 1$ the effectiveness of the lockdown policy is partial (people keep transmitting the virus) but at a lower rate. We will choose as a baseline $\theta = 0.5$, which we deem as reasonable, but we acknowledge we have no hard evidence to lean on.⁴

⁴This choice of θ has the interesting feature of making the marginal effect of the quadratic search effect of lockdown on the transmission of the virus the same as in the linear case. By this we mean that for $\theta = 0.5$ the term $(1 - \theta L)^2$, which gives the reduction in transmission of the virus in [equation \(2\)](#) and [equation \(3\)](#), has the same derivative with respect to L as the linear case of $(1 - L)$ at no lockdown ($L = 0$).

3. Those recovered are assumed to be immune and, in our benchmark analysis with a test ($\tau = 1$), we assume the recovered are not locked down. It seems like a property of any optimal policy, provided that previously infected agents cannot contract the infection again (see Dewatripont et al. (2020) for a related policy). The case with no test ($\tau = 0$), when everybody must be in lock down, is also discussed.
4. In the law of motion for S and I given by equation (2) and equation (3), we do not scale by the population $N(t)$ to compute the number of infected at time t . Instead we simply write, $\beta I(t)S(t)(1 - \theta L(t))^2$, instead of $\beta I(t)S(t)(1 - \theta L(t))^2/N(t)$. This seems to be the standard in the SIR literature, although it will be preferable to scale them by $N(t)$. Since some agents die, this cannot be done by simple setting $N(0) = 1$. We follow the literature, since this also saves one state variable.
5. Note that the assumption in the SIR model that new infections are proportional to the product of S and I , means that when one assume that it can lockdown those that have contacts, the term $(1 - \theta L)$, i.e the fraction not lockdown, enters in a quadratic form. Hence, if $(1 - \theta L)$ is one half, it has the effect to reduce to one-fourth the rate of transmission of infections in that period. This is similar to the feature in search theory aptly named “quadratic search”.
6. We assume that agents who are infected, but not in lockdown, can still produce as much as those susceptible or recovered not in lockdown. This can be easily changed, but a better model will have two types of infected, one with symptoms and one without. Those with symptoms should be lockdown first, and only after them, one should isolate the rest.⁵
7. We assume that agents in lockdown do not produce. A more realistic model will have at least some agent who can produce, even in lockdown (working remotely). This variation

⁵We do not include an intermediate state $E(t)$, as it is done in some of the SIR models –see Atkeson’s note–, where agents have no symptoms but can transmit the virus. This may be interesting to include.

is easy to explore by adding a relative productivity loss parameter in front of the $S + I$ term in the objective function.

8. The assumption that $\phi(I)$ is increasing in I , postulated in [equation \(8\)](#), captures that as I increases, the capacity to treat patients diminishes and the rate at which an infected agent dies increases.
9. We assume that all agent are infinitely lived, except for the risk of dying after contracting the virus. This simplification is acceptable given the short time horizon of the problem. On the other hand, not having an explicit age structure yields unrealistic effects on the impact of mortality risk.
10. The planning problem incorporates two types of cost. One is the economic activity lost during a lockdown among those susceptible and infected, and also due to the death of infected agents. Note that since we model an infinitely lived agent, the death of an infected person implies a cost equal to the present value of this person's income, which overestimates the economic cost of death, measured as the value of a statistical life. In the model parameterization discussed below, we use the extra term for the cost of each death (χ) to adjust the valuation of the years lost in term of typical consumption to the finite remaining years of life span, as done in the literature that uses the value of a statistical life (*v.s.l*) and the adaptations to this particular application.

An equivalent planner minimization problem. We can rewrite the problem in [equation \(5\)](#) into an equivalent one that conserves on the state space and has a perhaps simpler interpretation. So far we described the case where those that recover can be identified and hence not subject to lockdown. Below, we also consider the case that such a test is not available so that the lockdown must apply to the entire population. The parameter $\tau = 1$ will indicate that the test is available which is our baseline case; we will also discuss the case where $\tau = 0$, i.e. the test it is not available.

The problem of the planner consists in minimizing the present value, discounted at rate $r + \nu$, of the following flow cost:

$$wL \left[\tau(S + I) + 1 - \tau \right] + I\phi(I) \left[\frac{w}{r} + \chi \right] \quad (6)$$

so the cost of having a state (S, I) and a current control L is that an amount $wL(S + I)$ of output by the susceptible and infected is lost if the test is available ($\tau = 1$), or that an amount wL is lost in the case in which the test is not available ($\tau = 0$).

The cost of fatalities due to the infection is given by the present value of output that they would have produced, w/r , as well as by the extra cost of death, χ . Of course, what matters for the problem is the magnitude $\left[\frac{w}{r} + \chi \right]$, which has the same units as the value of a statistical life.

Thus, we can write the objective function as:

$$\int_0^\infty e^{-(r+\nu)t} \left(wL(t) \left[\tau(S(t) + I(t)) + 1 - \tau \right] + I(t) \phi(I(t)) \left[\frac{w}{r} + \chi \right] \right) dt \quad (7)$$

This problem has a two-dimensional state, (S, I) .⁶ Recall that by adjusting χ one controls the weight of fatalities relative to a unit of output. Indeed our benchmark case is very close to trade off a unit of output loss with the statistical value of a life, as commonly measured.

The per day rate at which infected people die, which we refer to as the fatality rate, is assumed to have the following functional form

$$\phi(I) = \varphi + \kappa I \quad (8)$$

which reflects that the capacity to treat patients diminishes and, as a result, the rate at

⁶The problem formulation assumes testing is available or not, through the parameter τ . It would be interesting to modify the setup so that τ would also be a control, incorporating its associated cost. Notice that in [equation \(6\)](#) when $\tau = 0$ the output cost is $wL(t)$, while the actual cost would be slightly smaller and equal to $wN(t)L(t)$, where $N(t) < N(0)$ since a small part of the population has died by time t . We use this simplification to maintain a parsimonious state space, as common practice in SIR models.

which infected agents die increases.

The planner solves the following Bellman-Hamilton-Jacobi equation:

$$\begin{aligned}
(r + \nu)V(S, I) = & \min_{L \in [0, \bar{L}]} \left[\tau(S + I) + 1 - \tau \right] + I\phi(I) \left[\frac{w}{r} + \chi \right] + \\
& - [\beta S(t)(1 - \theta L(t)) I(t)(1 - \theta L(t))] \partial_S V(S, I) \\
& + [\beta S(t)(1 - \theta L(t)) I(t)(1 - \theta L(t)) - \gamma I(t)] \partial_I V(S, I)
\end{aligned} \tag{9}$$

The domain of V is $(S, I) \in \mathbb{R}^2$ such that $S + I \leq 1$. Note that $V(S, I)$ can be interpreted as the minimum expected discounted cost of following the optimal policy in units of output loss. Recall that if $\phi(I) = 0$, then the value of the objective of the planner at time $t = 0$ will give $N(0)w$. Thus, the quantity $rV(S, I)/w$ converts the stock-value of the value function into a ratio of the flow cost relative to output at time $t = 0$, when there is no virus. We solve this problem by discretizing the model to weekly intervals, using value function iteration over a large non-linear grid for (S, I) . Finally, we notice that the value function has simple analytic expressions on the boundary of its domain, where the lockdown policy is not exercised. In particular, on the $I = 0$ axis we have $V(S, 0) = 0$, for all $S \in (0, 1)$, so that the cost is zero if nobody is infected. On the $S = 0$ axis we have $V(0, I) = \frac{(w+\chi)I}{r} \left(\frac{\varphi}{r+\nu+\gamma} + \frac{\kappa I}{r+\nu+2\gamma} \right)$ for all $I \in (0, 1)$.

3 Parametrization of the model

We parameterize the model using data from the World Health Organization (WHO) compiled by the Johns Hopkins University Center for Systems Science and Engineering (JHU CCSE) while acknowledging, like the rest of the recent literature, that at this point there is considerable uncertainty about infection, recovery, and mortality rates. The data includes the total cases, including separately those that have recovered and those that have died. We define active cases as the total number of cases minus those that either recovered or died.

We use daily observations of all the countries that have registered at least 100 active cases and include observations of the first 25 days after they first cross this threshold.

To calibrate β , the rate at which individuals who are infected bump into other people and shed virus onto those people, we use the daily increase in active cases and assume a value of 20 percent. The parameter γ governing the rate (per day) at which infected people either recover or die is considered a fixed parameter of the disease and is set to $\gamma=1/18$ reflecting an estimated duration of illness of 18 days as in [Atkeson \(2020\)](#) but also consistent with the fraction of infected agents that recovered or died according to the WHO as compiled in JHU CCSE.

We set the fatality rate $\varphi=0.01 \times \gamma$, which is consistent with the age-adjusted fatality rate estimated from the *Diamond Princess* cruise ship and with the lower bound mortality rate in the city of Vo' Euganeo – two cases where there has been extensive testing. We set $\kappa=0.05 \times \gamma$ so the fatality rate is 3 percent when 40 percent of the population is infected. There is considerable uncertainty on the fatality rate, mostly because the true rate of infected is not really known. For instance, [Eran, Jay, and Sood \(2020\)](#) argue that the number of infected is probably at least an order of magnitude larger, and thus the mortality rate much smaller.

We set the planner's discount factor to be consistent with a 5 percent annual interest rate and the per unit of time probability ν that a vaccine and a cure will appear so that it implies that it takes on average a year and a half for these medical discoveries to become available. We normalize output $w=1$ and set the extra cost of dying χ to zero. Note that in this case a unit of output produced by each agent, w can be interpreted as GDP per capita, let say 65,000 USD. On the other hand, using $r = 0.05$ (5% per year) and $\chi = 0$, then the shadow cost of each life lost used by the planner is $[w/r + \chi]$, which is 20 times annual GDP per capita, or about \$1.3 Million.

Our choice of the benchmark value for $\chi = 0$, and hence of the penalty deaths, is in line with [Hall, Jones, and Klenow \(2020\)](#). These authors use an utilitarian criterium to value the extra years of life lost among those likely to die due to the infection, obtaining a cost of about

30 times per capita annual consumption, which is very close to our benchmark value of 20 times annual GDP per capita.⁷ The value of 20 annual per capita GDP is much lower than the typical figures for statistical value of life, which are closer to \$ 10 million, see [Kniesner and Viscusi \(2020\)](#), or about 150 GDP per capita. We will report result for half this value, i.e. with a penalty of 10, and 30 annual GPD per capita, as well as for values closer to the typical value of statistical life, of 80 and 140 times annual GDP per capita.

Lastly, we assume that even in a disaster scenario, economic sectors such as health, government, retail, utilities, and food manufacturing will continue. These sectors combined account for 25-30% of GDP (2018). Thus, we set $\bar{L} = 0.7$.

Table 1: Parameter Values for Benchmark Case

Parameter	Value	Definition/Reason
β	0.20	Daily increase of active cases if unchecked
γ	1/18	Daily rate of infected recovery (includes those that die).
φ	$0.01 \times \gamma$	IFR: fatality per active case (per day).
κ	$0.05 \times \gamma$	Implies a 3 percent fatality rate with 40 percent infected.
r	0.05	Annual interest rate 5 percent.
ν	0.667	Prob rate vaccine + cure (exp. duration 1.5 years)
\bar{L}	0.70	1 - GPD share health, retail, government, utilities, and food mfg.
θ	0.50	Effectiveness of lockdown
χ	0	Value of Statistical Life $20 \times w$ (i.e. $v.s.l \approx \$1.3M$)

⁷Following [Hall, Jones, and Klenow \(2020\)](#) and the literature they refer to, one can use that a year of life lost is valued as three times annual consumption. Then, following their ideas, one can compute the expected number of years of lives lost to those that die as a consequence of the virus, conditional on being infected. They obtain a number between 10 and 15 years, with 10 being their headline figure. Thus, $3 \times 10 \text{ years} \times \text{annual consumption per capita} = 3 \times 10 \text{ years} \times 2/3 \times \text{annual GDP per capita} = 20 \times \text{annual GDP per capita}$.

4 Results

In this section, we present the result for our numerical calculations. We display the time path of the optimal policy starting at $I(0) = 0.01$, i.e. one percent of population infected at $t = 0$ for our benchmark parameter values.⁸ In particular, we display the time path of the optimal lockdown policy $L(t)$ as function of time, the fraction of the population for which lockdown applies $L(t)[\tau(S(t) + I(t)) + 1 - \tau]$, the path of infected $I(t)$, and the total accumulated fraction of dead up to time t . Recall that $N(0) = 1$, so both infected and the stock of dead can be all interpreted as fraction of the initial population. In this graphs, the horizontal axis is time, conditional on the cure-vaccine not occurring before that period. For comparison, we also plot the path if there is no lockdown policy, i.e. for $L(t) = 0$ for all $t \geq 0$. We complement these time paths with other type of figures: heat maps of the policy as a function of the state, including phase diagrams, and a 3-D representation of the value function, as well as with a table with the value of following the policy across different scenarios.

Benchmark case. We present the result for our benchmark case first, and then we do some sensitivity analysis. On one hand, our benchmark case favors a policy of lockdown due to the following features: (i) the chosen values for the parameters β and γ of the SIR model imply that a large fraction will be exposed to the virus if unchecked, $\lim S(t) \approx 0.97$ as $t \rightarrow \infty$, (ii) we assume that the fatality rate can increase from 1% to up to 3% of those infected when fraction of infected goes from 5% to 40%, and (iii) we assume that those recovered can be identified and hence they are not locked down. On the other hand, our benchmark case uses a value of statistical life of 20 times annual per capita GDP, which is in line with utilitarian values of life for those likely to be affected by COVID-19, see [Hall, Jones, and Klenow \(2020\)](#), but an order of magnitude smaller than the average value of a statistical life used in other public policy evaluations. We conduct sensitivity analysis for each of the these assumption

⁸We assume that the initial fraction of the population susceptible is 98%, or $S(0) = 0.98$.

below.

In what follows, we first present the result for our benchmark parameter case, in [Figure 1](#) and [Figure 3](#), as well as its associated value function and optimal policy [Figure 2](#). A phase diagram for the benchmark case is shown in [Figure 7](#). For the benchmark parameter values, the top panel of [Figure 1](#) shows that the lockdown starts only after two weeks. For these parameters, the fraction of the population in lockdown peaks at 60%, about 10 days after the lockdown starts, and decreases to be below 10% of the population by the sixteenth week after the lockdown started. The lockdown ends in about 4 months. This implies a considerable flattening of the curve of infected, as can be seen in the middle panel of the figure, by comparing the red (without no lockdown) vs blue lines (with optimal policy). In the long run, the total number of deaths is about 0.80% smaller with the optimal policy, as can be seen in the bottom panel of the figure.

Lower effectiveness of lockdown. We explore the sensitivity to parameter values by changing the effectiveness of the lockdown, i.e. by reducing it from $\theta = 0.5$ to $\theta = 0.3$. In the case of less effective lockdown, displayed in [Figure 3](#) i.e. lower θ , the duration and severity are lower. The fraction of population peaks in approximately 20 days, but it decreases at a faster rate, reaching zero lockdown two months after the lockdown started. Instead, if the lockdown were to be more effective, say $\theta = 0.7$, the duration will be even longer –case not displayed in figures.

Constant fatality rate function (no congestion of health care system). Importantly, the results change dramatically, in the sense that it is essentially optimal to have zero lockdown, if $\kappa = 0$. This is the case where the fatality rate is constant at 1%, so there is no congestion on the health care system.

Different values of statistical life. In [Figure 4](#) we explore the consequence of a smaller implied statistical value of life, half the value of the benchmark case. Unsurprisingly, the

lower value of $v.s.l$ considerably diminishes the optimal lockdown level and duration, peaking at a lower value than benchmark case, and with duration since start to finish of about 50 days.

In [Figure 5](#), we explore the effect of increasing the value of statistical life to 30 times annual GDP per capita –which is in the upper end of the values consider by [Hall, Jones, and Klenow \(2020\)](#).⁹ In this case, the lockdown starts in two weeks –a bit faster than the benchmark–, peaks in a month with about 60% of the population in lockdown, and decreases linearly and slowly, until is abandoned slightly more than six month after it started. The fraction of population in lockdown reaches 10% only in slightly more than about 4 month after the lockdown started. In this case, the more aggressive, and specially longer, lockdown policy implies that the fraction of death after the epidemic is over is reduced by 1%, about 0.20% more than in the benchmark.


For reference, we have also computed the cases where the value of statistical life is 80 and 140 times annual per capita GDP, even though we view these values as too high for this particular application. As expected the optimal lockdown rate is very high, for a very long time. In the case of $v.s.l. = 80$ ann. per capita GDP, the lockdown rate starts in two weeks, and $L(t) = \bar{L}$ for about 8 months. The fraction of population in lockdown reaches 50% slightly about 3 months after the lockdown started, and it is approximately 15% a year into the lockdown, when $L(t)$ is below its allowed maximum. The case of $v.s.l. = 140$ ann. per capita GDP is similar, since it is $L(t) = \bar{L}$ for about 10 month. Another small difference is that the lockdown starts almost right away in the case of the largest $v.s.l.$.

No testing of the recovered. [Figure 6](#) shows the results of the case with no test, so $\tau = 0$ in the model. In this case the lockdown applies to anybody in the population, including those that have recovered from the virus. Recall that in this case, it is less efficient to lockdown agents, because the recovered are also lockdown, which it has the cost of reducing output,

⁹For instance, this is consistent with a value of each year of extra life of 3 times annual consumption per capita, times 15 years of lost expected life expectancy conditional on dying after being infected.

without the benefit of reducing the transmission of the virus. In the case of no test, the lockdown declines much more sharply than in the benchmark case with testing. Interestingly, in total it lockdown about the same hours per capita as the benchmark case, but its it does not take advantage of the natural increase in those recovered.

Figure 1: Benchmark Case (medium effectiveness $\theta = 0.5$)



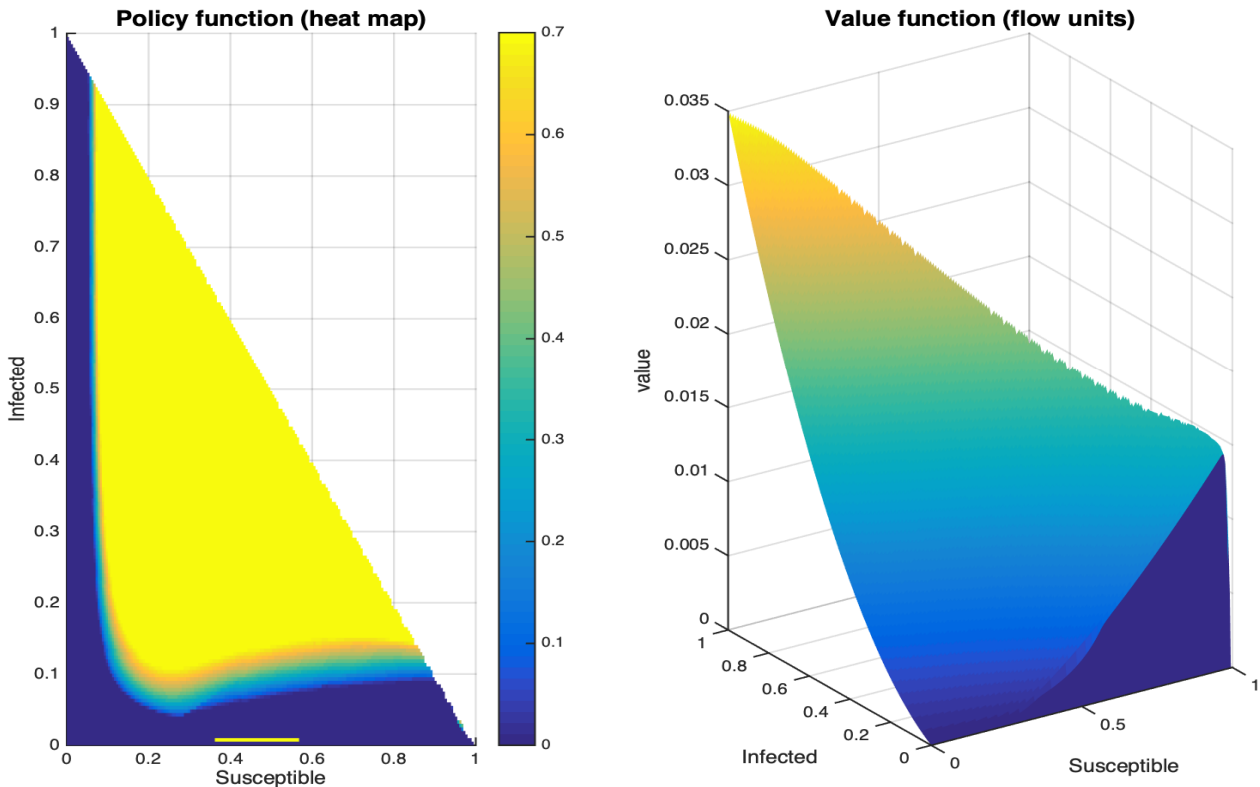
Note: The top panel shows the time paths for the lockdown rate (dashed blue line) and the share of the population under lockdown (solid blue line). The panel in the center shows the share of the population that is infected with control (blue) and without control (red). The bottom panel shows the share of the population that die under control (blue) and under no control (red). The paths shown are under medium effectiveness (i.e. $\theta=0.5$). The initial condition is $I(0) = 0.01$ and $S(0) = 0.97$.

Details on the benchmark case. Figure 2 displays the value function and the optimal policy for the benchmark parameter values. The value function is plotted in the right panel, for the relevant state space (S, I) , and normalized as explained above, so we display $rV(S, I)/w$ in the vertical axis. The units are permanent flow cost as a fraction of the total output previous to the virus. Thus, a value of 0.02, means a cost equivalent to a permanent

reduction of 2% percent in the value of output, relative to the flow of output before the virus. Note that the V in the boundary of the state space with either $S = 0$ or $I = 0$, has the properties described in [Section 2](#).

The left panel panel of [Figure 2](#) has a heat map of the optimal policy $L^*(S, I)$. Yellow indicates higher value of the lockdown rate L , and blue indicates lower values of L . Note that close to both boundaries, i.e. either $S = 0$ for $I = 0$, ie. it is optimal to have zero lockdown rate.

Figure 2: Value Function and Optimal Policy, benchmark case



Note: The figure on the left shows the optimal policy for the benchmark parameter values. The blue area indicates lower values of lockdown and the yellow color higher values. The figure on the left depicts the value function. The units for the value function are permanent flow cost as a fraction of the total output before the epidemic.

[Figure 7](#) shows the heat map from the optimal policy, as in the previous figure, and a two

phase diagrams, i.e. two paths, in the state space (S, I) that the system will take starting from an initial conditions. One path corresponds to the case of no policy, which moves to large values of the fraction of infected –given by the red line. The other path has the evolution of the state under the optimal policy –the dash light blue line–, where it can be seen that the fraction infected is much smaller, as a consequence of the lockdown policy. Also you can see how the two path coincide for a while, since the position of the seed of the virus is in the region of the state space where lockdown is not optimal. Then the optimal path is controlled by the lockdown, and shows a much lower fraction of infected in the vertical axis. Eventually, the path moves to the region with no lockdown, which occurs after the system has acquired herd immunity –note that at the end the trajectory is decreasing in I even if there is no lockdown. This phase diagram can be used to follow other alternative paths, such as what would happen if the optimal policy were to start after the virus has been unchecked for a longer period of time. Note that unless the susceptible have reached a very small number, starting the optimal policy later will involved immediate lockdown, i.e. the path will start in the yellow area.

Size of the welfare cost under optimal policy. [Table 2](#) summarizes the value of following the optimal policy vs. the value where there is no lockdown (No Policy), for different parameter values.¹⁰ As explained above, our favorite all summary measure is to report $rV(S(0), I(0))/w$. This number is the total expected discounted sum of future losses, both due to the lost GDP caused by the lockdown in future periods, as well as the values of the lost lives, where every life is evaluated at $[w/r + \chi]$. The multiplication by r in $rV(S(0), I(0))/w$, converts the expected discounted sum of future values into a permanent annual flow, and the division by w compares it with the level output before the virus was seeded. We also present the part of $rV(S(0), I(0))/w$ that is due to the output cost of the shutdown. This is just $r \int_0^\infty e^{-(r+\nu)t} [1 - \tau + \tau L(t)(S(t) + I(t))] dt$. Capitalizing the cost due to the lockdown, i.e.

¹⁰The no policy case should not be interpreted as what will occur in the absence of policy, since an equilibrium, even if constraint efficient, will feature individuals self-isolating.

dividing it by r , we can obtain also a simple statistic measuring the severity of the lockdown. For instance if this capitalized measure is 10%, then the lockdown is equivalent to stop 10% of the population to work for an entire year. Finally, the last column displays the present discounted value of leaving the virus unchecked, which we label as No Policy in column of [Table 2](#). In the three cases we express the losses in percentage. Importantly, if one wants to convert any of this figures to the equivalent into one time payment, we should just divide the them by r , i.e. multiplied them by 20 given our 5% annual interest rate.

For instance, for the benchmark case, the second row in top panel of [Table 2](#), following the optimal policy implies a permanent loss of approximately 1.5% of output. In other words, as a consequence of the outbreak of the virus, even following the optimal policy, welfare is comparable to an equivalent measure of being 1.5% permanently poorer. We can also recapitalize this loss and express it as fraction of one year GDP, obtaining, 28%. From the 1.5% total permanent loss, 0.4% is due to output losses due to lockdown, or an equivalent to 8% of one year GDP., While for the same parameters, if there is no lockdown (No Policy), the loss is equivalent to a permanent decrease in output of 1.9%. Thus, the optimal policy reduces the losses of an equivalent of permanent increase of 0.40% of GDP per year. The first three rows explore the different values of effectiveness of the lockdown.

The second panel, with four rows, corresponds to the case of different values of a statistical life. Recall that the benchmark case, $\chi = 0$ implies a value of a statistical of life (*v.s.l*) of about 20 annual GDP per capita. We consider a case with lower value, i.e. with a multiple of 10, and case with a higher value, 30 annual GDP per capita, and two cases with a much higher value, typical from uses of a statistical value of life in other contents, i.e. 80 and 140 annual GDP per capita. For a (*v.s.l*) of 30 annual GDP per capita, the part due to output loss is a permanent flow of 0.6% or equivalently a 12% fraction of one year GDP.

The next panel, with three rows corresponds to the case where the fatality rate $\phi(I)$ is constant at $\varphi = 0.01$, or equivalently $\kappa = 0$. In this case, as explained above, the optimal policy has no lockdown, so the loses of the optimal policy and the case of no policy are the

same, and also much smaller, since the death rate does not spike up. This highlights the importance of the assumption implied in our benchmark case that ϕ is increasing, which captures the extra fatalities due to the congestion in the health care system caused by a large number of infected.

The last panel, with four rows, corresponds to the case of no test, or $\tau = 0$ in the notation used above. For this case, we present five different values of χ , corresponding to five different values of a statistical life. Each row has *v.s.l* is a multiple of the annual GDP per capita, with multiples of 10, 20, 30, 80, and 140, and otherwise the parameters are the same as in the benchmark case. Comparing the benchmark case –i.e. the second row of the top panel– with the same case without test –i.e. the second row of the last panel– we find small differences. In particular the expected discounted cost under the optimal policy is, expressed as a permanent flow is 0.1% (10 basis point) higher without test than with test, i.e. 1.6% vs 1.5%, or in terms of a one time value approximately 2% of a year GDP. This difference gets larger for larger values of a statistical life. For a smaller value of *v.s.l.*, we find that the difference is smaller than 0.1%. Instead, as *v.s.l* increases to multiples of annual GDP per capita of 30, 80 and 140, this difference, measuring the value of the test as a permanent flow of GDP losses is 0.40%, 0.80% and 0.50% respectively. For our preferred parameter values, say *v.s.l.* between 20 and 30 times ann. per capita GDP, the value of the test is equivalent to between 2% and 8% of one year GDP.


Table 2: Welfare Losses $\left(\frac{rV(S,I)}{w}\right)$ with Optimal Policy vs. without Intervention

Case	Parameters	Welfare Loss	Output Loss	Welfare Loss
		w/Policy	w/Policy	No Policy
<i>Benchmark Case</i>				
Low effectiveness	$\theta=0.3$	1.7 %	0.3 %	1.9%
Medium effectiveness	$\theta=0.5$	1.5 %	0.4 %	1.9%
High effectiveness	$\theta =0.7$	1.4 %	0.5%	1.9 %
<i>Alternative Values of Statistical Life, different χ</i>				
$v.s.l = 10 \times$ GDP per capita	$\chi = -\frac{1}{2} \frac{w}{r}$	0.9 %	0.2 %	0.9 %
$v.s.l = 30 \times$ GDP per capita	$\chi = \frac{1}{2} \frac{w}{r}$	2.0 %	0.6 %	2.8 %
$v.s.l = 80 \times$ GDP per capita	$\chi = 3 \frac{w}{r}$	3.7 %	1.6 %	7.5 %
$v.s.l = 140 \times$ GDP per capita	$\chi = 6 \frac{w}{r}$	5.7 %	1.0 %	13.2 %
<i>Constant fatality rate $\kappa=0$</i>				
Low effectiveness	$\theta=0.3$	0.9 %	0.0 %	0.9 %
Medium effectiveness	$\theta=0.5$	0.9 %	0.0 %	0.9 %
High effectiveness	$\theta=0.7$	0.9 %	0.0 %	0.9 %
<i>No testing of the recovered $\tau = 0$</i>				
$v.s.l = 10 \times$ GDP per capita	$\chi = -\frac{1}{2} \frac{w}{r}$	0.9 %	0.1 %	0.9 %
$v.s.l = 20 \times$ GDP per capita	$\chi = 0$	1.6 %	0.4 %	1.9 %
$v.s.l = 30 \times$ GDP per capita	$\chi = \frac{1}{2} \frac{w}{r}$	2.2 %	0.6 %	2.8 %
$v.s.l = 80 \times$ GDP per capita	$\chi = 3 \frac{w}{r}$	4.5 %	2.5 %	7.5 %
$v.s.l = 140 \times$ GDP per capita	$\chi = 6 \frac{w}{r}$	6.2 %	2.7 %	13.2 %

Note: Welfare losses are measured by the permanent percent reduction in per capita GDP induced by the policy (or its absence) under various parametrizations. Output losses is the welfare cost component due to the reduced level of economic activity (i.e. excluding fatalities). The benchmark parameter values are from [Table 1](#). Multiplying any of the numbers in the last three columns by $1/r = 20$, converts the losses from permanent flow to a one time payment as a fraction of a year GDP. The initial condition for all scenarios is $I(0) = 0.01$ and $S(0) = 0.97$.

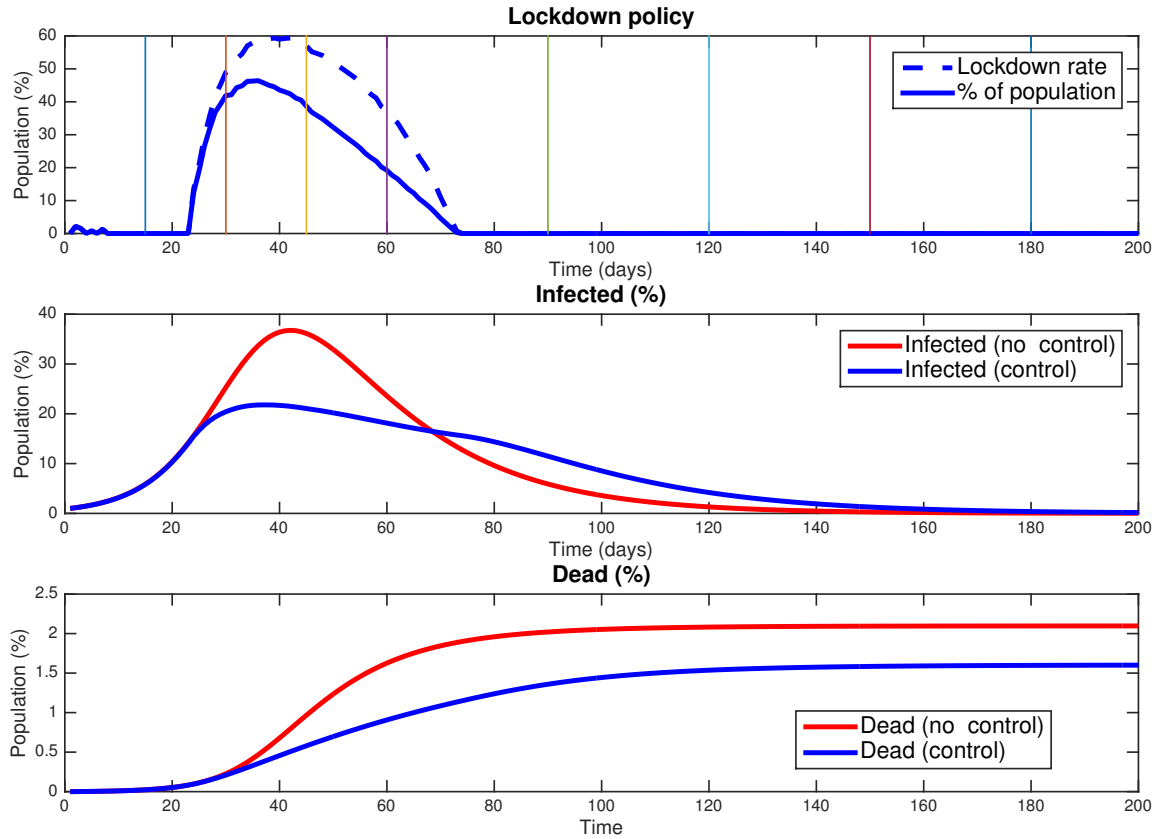
5 Other figures

Figure 3: Low Lockdown Effectiveness ($\theta = 0.3$)




Note: The top panel shows the time paths for the lockdown rate (dashed blue line) and the share of the population under lockdown (solid blue line). The panel in the center shows the share of the population that is infected with control (blue) and without control (red). The bottom panel shows the share of the population that die under control (blue) and under no control (red). The paths shown are under medium effectiveness (i.e. $\theta=0.3$). The initial condition is $I(0) = 0.01$ and $S(0) = 0.97$.

Figure 4: Lower Value of Statistical Life, $v.s.l. = 10$ ann. GDP p/capita)




Note: The top panel shows the time paths for the lockdown rate (dashed blue line) and the share of the population under lockdown (solid blue line). The panel in the center shows the share of the population that is infected with control (blue) and without control (red). The bottom panel shows the share of the population that die under control (blue) and under no control (red). The paths shown are under a lower value of statistical life (i.e. $\chi = -\frac{1}{2} \frac{w}{r}$, $v.s.l$ approximately \$0.65 Million). The initial condition is $I(0) = 0.01$ and $S(0) = 0.97$.

Figure 5: Higher Value of Statistical Life, ($v.s.l. = 30$ ann. GDP p/capita)



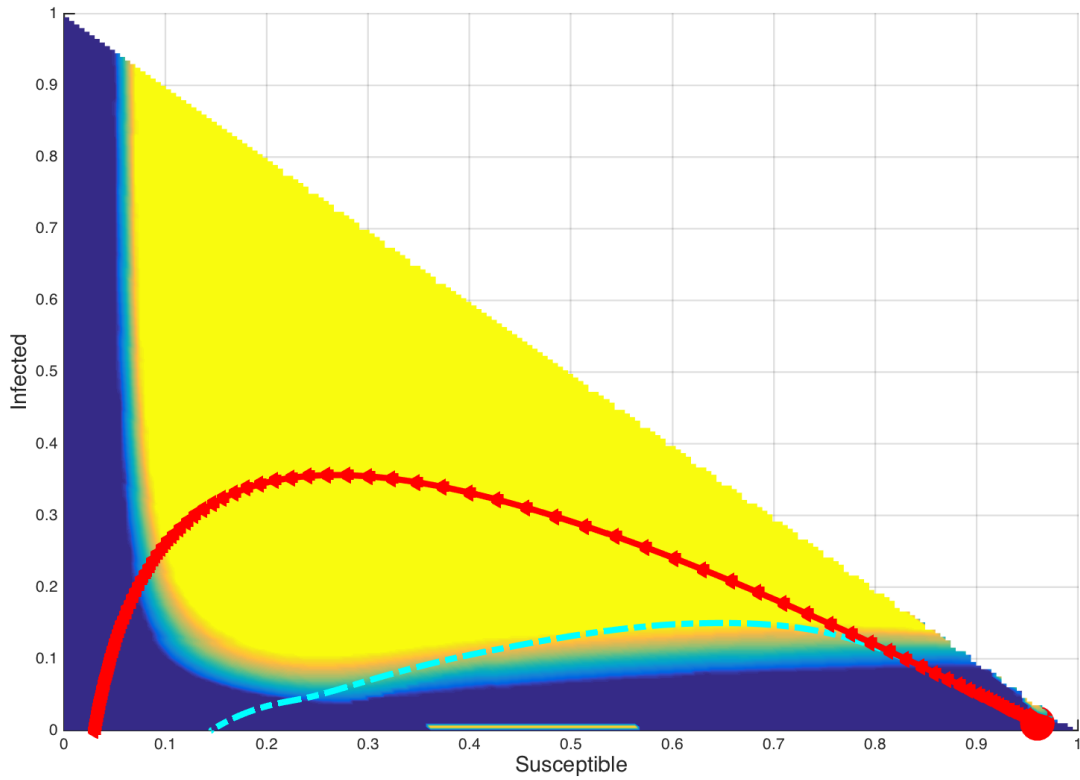
Note: The top panel shows the time paths for the lockdown rate (dashed blue line) and the share of the population under lockdown (solid blue line). The panel in the center shows the share of the population that is infected with control (blue) and without control (red). The bottom panel shows the share of the population that die under control (blue) and under no control (red). The paths shown are under a higher value of statistical life (i.e. $\chi = \frac{1}{2} \frac{w}{r}$, $v.s.l$ approximately \$1.9 Million). The initial condition is $I(0) = 0.01$ and $S(0) = 0.97$.

Figure 6: No testing of those recovered ($\tau = 0$)



Note: Except for the testing, which is set to $\tau = 0$, this figure uses the benchmark parameter values. The top panel shows the time paths for the lockdown rate and the share of the population under lockdown, which coincide in this case, since there is no testing. The panel in the center shows the share of the population that is infected with control (blue) and without control (red). The bottom panel shows the share of the population that die under control (blue) and under no control (red). The case shown is with no testing for those recovered. The initial condition is $I(0) = 0.01$ and $S(0) = 0.97$.

Figure 7: Phase Diagram, benchmark case with $I(0) = 0.01$



Note: The figure shows the optimal policy for the benchmark parameter values. The blue area indicates lower values of lockdown and the yellow color higher values. The figure also shows two phase diagrams in the state space (S, I) . The red dot indicates the initial condition $I(0) = 0.01$ and $S(0) = 0.97$. The red line corresponds to the case of no policy. The light blue line shows the dynamics of the state under the optimal policy.

References

- Atkeson, Andrew G. 2020. “What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios.” Staff Report 595, Federal Reserve bank of Minneapolis.
- Baldwin, Richard and Beatrice Weder. 2020. “Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes.” VoxEu.org, CEPR.
- Barro, Robert J., Jose F. Ursua, and Joanna Weng. 2020. “The coronavirus and the great influenza epidemic: Lessons from the spanish flu for the coronavirus potential effects on mortality and economic activity.” AEI Economics Working Paper 2020-02, American Enterprise Institute.
- Dewatripont, Mathias, Michel Goldman, Eric Muraille, and Jean-Philippe Platteau. 2020. “Rapid identification of workers immune to COVID-19 and virus-free: A priority to restart the economy.” Discussion paper, Universit Libre de Bruxelles.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2020. “The Macroeconomics of Epidemics.” Author’s website, Northwestern University.
- Eran, Bendavid, Bhattacharya Jay, and Neeraj Sood. 2020. “We cannot trust current COVID-19 models.” Author’s website, Stanford University.
- Ferguson, Neil M and Et.al. 2020. “Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand.” Author’s website, Imperial College London.
- Hall, Robert E., Charles I. Jones, and Peter J. Klenow. 2020. “Trading Off Consumption and COVID-19 Deaths.” Mimeo, Stanford University.
- Hansen, Elsa and Day Troy. 2011. “Optimal control of epidemics with limited resources.” *Journal of Mathematical Biology* volume 62:423–451.

Jones, Callum, Thomas Philippon, and Venky Venkateswaran. 2020. “Optimal Mitigation Policies in a Pandemic.” Tech. rep., NYU.

Kniesner, Thomas and Kip Viscusi. 2020. “The Value of a Statistical Life.” *Oxford Research Encyclopedia of Economics and Finance*, forthcoming .

Neumeyer, Pablo Andres. 2020. “Clase especial de epidemiologia.” Author’s website, Class notes, Universidad Di Tella.

Piguillem, Facundo and Liyan Shi. 2020. “The Optimal COVID-19 Quarantine and Testing Policies.” Working paper 20/04, EIEF.

Toxvaerd, Flavio. 2020. “Equilibrium Social Distancing.” Mimeo, Cambridge University.