DISCUSSION PAPER SERIES

DP14658

A SIMPLE PLANNING PROBLEM FOR COVID-19 LOCKDOWN

Francesco Lippi, Fernando Alvarez and David Argente

INTERNATIONAL MACROECONOMICS AND FINANCE

ISSN 0265-8003

A SIMPLE PLANNING PROBLEM FOR COVID-19 LOCKDOWN

Francesco Lippi, Fernando Alvarez and David Argente

Discussion Paper DP14658
Published 24 April 2020
Submitted 22 April 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

A SIMPLE PLANNING PROBLEM FOR COVID-19 LOCKDOWN

Abstract

We study the optimal lockdown policy for a planner who controls the fatalities of a pandemic while minimizing the output costs of the lockdown. The policy depends on the fraction of infected and susceptible in the population, prescribing a severe lockdown beginning two weeks after the outbreak, covering 60% of the population after a month, and gradually withdrawing to 20% of the population after 3 months. The intensity of the optimal lockdown depends on the gradient of the fatality rate with respect to the infected, and the availability of antibody testing that yields a welfare gain of 2% of GDP.

JEL Classification: I10, I18, C61

Keywords: lockdown, Quarantine, epidemic control, Dynamic programming

Francesco Lippi - flippi@luiss.it
LUISS university and CEPR

Fernando Alvarez - f-alvarez1@uchicago.edu
University of Chicago

David Argente - dargente05@gmail.com
Penn State university

Acknowledgements
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, Steve Mohr, 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 Pan- demics Economics Group) seminar, the World Bank's Development Policy and Covid-19 e-seminar, the Federal Reserve Bank of Chicago Virtual Macro Seminar. The authors declare to have no conflict of interest to disclose regarding the research on this paper.

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 17, 2020

Abstract

We study the optimal lockdown policy for a planner who controls the fatalities of a pandemic while minimizing the output costs of the lockdown. The policy depends on the fraction of infected and susceptible in the population, prescribing a severe lockdown beginning two weeks after the outbreak, covering 60% of the population after a month, and gradually withdrawing to 20% of the population after 3 months. The intensity of the optimal lockdown depends on the gradient of the fatality rate with respect to the infected, and the availability of antibody testing that yields a welfare gain of 2% of GDP.

*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, Steve Mohr, 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, the World Bank's Development Policy and Covid-19 e-seminar, the Federal Reserve Bank of Chicago Virtual Macro Seminar. The authors declare to have no conflict of interest to disclose regarding the research on this paper.

1 Introduction and Overview

We adopt a variation of the SIR epidemiology model reviewed by Atkeson (2020) and Neumeyer (2020) to analyze the optimal control model for the COVID19 epidemic. Our aim is to contribute to the ongoing discussion on the optimal policy response to the COVID19 shock, see Barro, Ursua, and Weng (2020); Eichenbaum, Rebelo, and Trabandtz (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, 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 et al. (2020). We use a simplified version of these models to analyze how to optimally balance the fatality induced by the epidemic with the output costs of the lockdown policy.[1] The novel aspect of our analysis is to explicitly formulate and solve a control problem, where the diffusion parameter is affected by the lockdown, that is chosen to maximize a social objective while taking into account the dynamic evolution of the system.[2] 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 Trabandtz (2020) and Toxvaerd (2020). The state of the problem is two dimensional and, in spite of its simplicity, it does not have an analytic solution. By computing the optimal policy and the associated trajectories, we aim to gauge the key elements that determine the optimal *intensity* and *duration* of the lockdown.

[1]While 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.

[2]An 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.

We solve the problem under different scenarios, that include congestion effects in the health care system, the effectiveness of the lockdown in reducing the diffusion of the virus and the possibility of testing for antibodies.

We parametrize the model using a range of estimates about the COVID19 epidemic. Since we recognize that several parameters are highly uncertain we explore a range of variations concerning 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. The quantitative results are useful to gauge what parameters of the problem are important in shaping the intensity and duration of the optimal lockdown policy. 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 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 Panel A in 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. In the less pessimistic cases, yet in our view still realistic, obtained by assuming a lower fatality rate and/or a lower speed of spread of the virus, the optimal lockdown is shortened by more than one month. The intensity of the optimal lockdown depends critically on the gradient of the fatality rate as a function of the infected. If we consider a constant fatality rate 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 times the annual GDP per capita, implies a very strict lockdown that lasts for about 9 months and maintains about 15% of the population in lockdown a year after.

Finally, our benchmark scenario assumes that there is an antibody test that allows those that recover to be issued an immunity card and 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 1 and Table 2). The most salient feature of the case where a test is not available is that the lockdown ends up sooner, more abruptly. The dynamics of the epidemiological model give the insight why this is optimal: as time goes by, the fraction of those recovered increases, and thus the lockdown becomes progressively less efficient to stop the transmission of the virus by locking down a progressively larger fraction of those that do not transmit it. The availability of such test has large welfare gains, in the order of 2% of one year's GDP.

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). Under our preferred values, the total welfare cost of the virus is equivalent to a loss of 30% to 40% of one year's GDP. From this loss, the part due to lockdown of workers is between 8% and 12% of one year's GDP.

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 (e.g. building emergency hospitals).

Our objective is similar to that of Eichenbaum, Rebelo, and Trabandtz (2020). While they focus on a competitive equilibrium where a consumption tax is used to slow-down economic activity and the epidemic diffusion, we focus on a simple planner's problem. In our setup, the interaction of the law of motion coming of the SIR model and the lockdown policy 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 outline is as follows: the next section describes the planner's problem and the epidemic model. Section 3 discusses the key model parameters. Section 4 reports the results of the optimal control problem under different scenarios. Section 4.1 quantifies the welfare costs of the lockdown policy under alternative scenarios, and compares them with the costs produced in a scenario without intervention. Section 5 discusses future extensions.

2 A planner model of lockdown control

We start with a modified version of the SIR model as described in Atkeson (2020). Agents are divided between those susceptible to be infected $S(t)$, those infected $I(t)$, and those recovered $R(t)$, i.e.

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

The "recovered" include those that have been infected, survived the disease, and are now assumed to be immune. Since we only include those that are alive $N(t)$ is changing through time. We normalize the initial population to $N(0) = 1$. The planner can lockdown a fraction $L(t) \in [0, \bar{L}]$ of the population, where $\bar{L} \leq 1$ allows us to consider that even in a disaster scenario some economic activity such as energy and basic food production will continue. 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:

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

where β is the number of susceptible agents per unit of time to whom an infected agent can transmit the virus after contact. All susceptible agents that get the virus become infected. For the infected, a fraction γ recovers, thus:

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

Note that locking down a part of the population, while economically costly, can be very powerful in reducing the rate at which susceptible agents become infected. This is because it is the *product* of the infected and susceptible that determines the new infections per unit of time. Hence, decreasing the number of contacts of each, decreases the new infections by its *square*. In search theory Diamond and Maskin (1979, 1981) aptly named this feature "quadratic search".

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

$$-\dot{N}(t) = \phi(I(t))\, I(t) \tag{4}$$

While we assume that the rate γ at which infected recover is constant, the rate at which the infected die varies with the number of infected I, according to

$$\phi(I) = [\varphi + \kappa I]\,\gamma \quad . \tag{5}$$

The term $[\varphi + \kappa I] \in (0, 1)$ is the "case fatality rate" (CFR), namely the proportion of infected persons that will die. It appears that the CFR is increasing with I, an assumption that reflects congestion effects in the health care system. The multiplication of the CFR by γ, the reciprocal of the infection expected duration, gives the fatality rate per unit of time.

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. The planner discounts all values at the rate $r > 0$. We also assume that with probability ν per unit of time both a vaccine and a cure appear, so that all infected are cured and all susceptible become immune. The problem consists in minimizing the following (discounted at rate $r + \nu$) present value:

$$\int_0^\infty e^{-(r+\nu)t} \left(wL\left(t\right) \Big[\tau(S(t) + I(t)) + 1 - \tau \Big] + \phi\left(I\left(t\right)\right) I\left(t\right) \cdot vsl \right) dt \tag{6}$$

The flow cost for the planner of having state (S, I) at t and selecting control L has two components. The first one is $wL\left[1 - \tau + \tau(S + I)\right]$, the output lost due to the lockdown. In the case where an antibody test is available ($\tau = 1$), so that an immunity card is released to the recovered, this equals w times the lockdown rate L times the population to which it applies, i.e. the sum of susceptible and infected. In the case where the test is not available ($\tau = 0$), the cost due the lockdown is wL.

The second component of the flow cost is the product of the number of deaths per period times the shadow value assigned to each death, or the value of a statistical life (vsl). In particular, if there are I infected, the deaths per unit of times are given by $\phi(I)\, I$. The cost of each fatality is given by the assumed value of a statistical life, which is discussed below. The planner's problem is subject to the law of motion of the susceptible 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 $S(0) + I(0) \geq N(0)$. Finally, when the vaccine and cure arrives there is no more cost, and thus the continuation value is zero.

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

$$(r+\nu)V(S,I) = \min_{L\in[0,\bar{L}]} wL\Big[\tau(S+I)+1-\tau\Big] + I\phi(I)\cdot vsl + \\ - \left[\beta\, S\, I(1-\theta L)^2\right]\partial_S V(S,I) \;+\; \left[\beta\, S\, I(1-\theta L)^2 - \gamma I\right]\partial_I V(S,I) \tag{7}$$

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 forgone output. We solve this problem by discretizing the model to daily intervals, using value function iteration over a dense grid for (S, I). Finally, note that the value function has analytic expressions on the boundary of its domain, where the lockdown policy is not exercised: on the $I = 0$ axis we have $V(S, 0) = 0$, for all $S \in (0, 1)$. On the $S = 0$ axis we have $V(0, I) = vsl \cdot\; I\left(\frac{\varphi}{r+\nu+\gamma} + \frac{\kappa I}{r+\nu+2\gamma}\right)$ for all $I \in (0, 1)$.

Discussion of modeling assumptions

1. We restrict the extent of the lockdown to $\bar{L} \leq 1$. This takes into account that some sectors cannot shut down (health, basic services, food production, 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.

3. In the law of motion for S and I, given by equation (2) and equation (3), we write $\beta IS(1 - \theta L)^2$, instead of $\beta IS(1 - \theta L)^2/N$. This seems standard in the SIR literature, although it would be preferable to scale them by N. But since the dead are a small fraction of N we follow the literature, and thus shrink the state space.

4. Infected not in lockdown are assumed to produce as much as those susceptible or recovered not in lockdown. Conversely, agents in lockdown produce zero. Both assumptions can be easily changed by rewriting the flow values of the objective function.

5. Agents are infinitely lived, except for the risk of dying of the virus. This simplification is acceptable given the short time horizon of the problem. We do correct for the age distribution of fatalities in our choice of the value of a statistical life.

3 Parameterization 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 γ=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 $\varphi = 0.01$, which is consistent with the age-adjusted fatality rate estimated form 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 $\kappa = 0.05$ 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 adopt a baseline value of a statistical life of 20 times w. 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, and the shadow cost of each life lost used by the planner is 20 times annual GDP per capita, or about \$1.3 Million USD.

Our choice of the benchmark value for $vsl = 20$, and hence of the penalty deaths, is in line with Hall, Jones, and Klenow (2020). These authors use an utilitarian criterion 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.[3] 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 results for a range of alternative values, considering vsl of 10, 30 and 80 annual GPD 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 $\bar{L} = 0.7$.

It goes without saying that the values for several parameter are speculative. We will conduct some sensitivity analysis to illustrate their importance.

[3]Following Hall, Jones, and Klenow (2020), one can use that a year of life lost is valued as three times annual consumption. Then, 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, $3 \times 10\, years \times$ annual consumption per capita $= 3 \times 10\, years\ \times 2/3\times$ annual GDP per capita $= 20 \times$ annual GDP per capita.

4 Results

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.[4] 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 these 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 $t \geq 0$.

Benchmark case. We present the results for our benchmark case first, and then implement a sensitivity analysis. 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, $\lim S(t) \approx 0.03$ as $t \to \infty$, *(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. However 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, 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 assumptions below.

Benchmark case with testing. Panel A of Figure 1 presents the result for our benchmark parameter case with testing. The first box in the panel shows the timeline for the optimal lockdown. The lockdown starts two weeks after the epidemic outbreak. The fraction of the population in lockdown peaks at 60%, about 1 month after the outbreak, and gradually

[4] We assume that the initial fraction of the population susceptible is 97%, or $S(0) = 0.97$.

decreases to reach about 10% of the population by the 14th week of lockdown. The lockdown ends in about 4 months. The policy yields a considerable flattening of the curve of infected, as shown in the middle panel of the figure, by comparing the red (no lockdown) vs the blue line (optimal policy). In the long run, the total number of deaths is about 0.80% smaller with the optimal policy, as shown in the bottom panel of the figure.

Benchmark case without testing. Panel B of Figure 1 shows the results of the case with no test, $\tau = 0$. 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 in lockdown, which has the cost of reducing output 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 both cases the lockdown involves similar costs in terms of forgone output, since in the absence of testing the lockdown duration is shorter but it applies to a larger fraction of people (recovered agents are also in lockdown). In spite of this similarity, we show below that welfare under the optimal policy with testing is higher, in the order of a permanent 0.1% GDP flow, which is equivalent to a one-time payment of 2% of GDP.

Details on the benchmark case with testing. 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 so that $rV(S, I)/w$ is on the vertical axis. The units are permanent flow costs as a fraction of the total output before the virus. Thus, a value of 0.02, means a cost equivalent to a permanent reduction of 2% percent in the value of output (measured before the virus). On the boundary of the state space, where $S = 0$ or $I = 0$, the function $V(S, I)$ has the properties described in Section 2.

The left panel of Figure 2 plots 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 a zero lockdown rate.

The left panel also plots two paths, using the phase diagram over the state space (S, I). These are the trajectories that the system will follow starting from the initial condition $I(0) = 0.01$ and $S(0) = 0.97$. The red path corresponds to the case of no intervention. The other path –the dashed light blue line– gives the evolution of the state under the optimal policy. It can be seen that the fraction of infected is much smaller under the lockdown policy. The two paths coincide for a while, since the initial condition lies in the region of the state space where lockdown is not optimal. Then the optimal path is controlled, and produces a much lower fraction of infected. 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 I is decreasing even if there is no lockdown. This phase diagram can be used to follow any 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.

Depending on parameter values and initial conditions the optimal policy may imply an early lockdown (for S, I pairs close to the $I = 0$ axis), followed by a relaxation of the policy and then by another lockdown.[5] We leave such cases for future investigations.

Lower effectiveness of lockdown. We explored the sensitivity to parameter values by changing the effectiveness of the lockdown, i.e. reducing it from $\theta = 0.5$ to $\theta = 0.3$. In the case of less effective lockdown the duration and severity are both smaller. The fraction of population peaks in approximately 20 days, but it decreases at a faster rate, reaching zero lockdown two months after the lockdown start. Instead, if the lockdown where to be more effective, say $\theta = 0.7$, the duration will be even longer.

Constant fatality rate function (no congestion of health care system). In this case the results change dramatically, in the sense that if $\kappa = 0$ under the benchmark parameters

[5] This can even be seen in the benchmark case, where there is a small isolated area of lockdown for very small positive value of I and value of S approximately between 0.4 and 0.5.

it is essentially optimal to have a zero lockdown. 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. Next we explore the consequences of a smaller implied statistical value of life, half the value of the benchmark case. Unsurprisingly, the lower value of vsl diminishes considerably the optimal lockdown level and duration, peaking at a lower value than the benchmark case, and with duration from start to finish of about 50 days. We also explored cases where the vsl is higher than our benchmark scenario. If we increase 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)[6] 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 months after it stared. The fraction of population in lockdown reaches 10% only after about 4 months 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.

Finally we consider the value of statistical life of 80 annual per capita GDP where, as expected, the optimal lockdown rate is very high and it last for a very long time. 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.

Two cases with less pessimistic parameter values. In the first case the value of β is half of the benchmark value, $\beta = 0.10$, so the virus spreads a slower pace and it reaches a lower fraction of the population even if $L(t) = 0$ for all t. In the second case, the baseline mortality rate is half of that in the benchmark case, i.e. $\varphi = 0.005$. Otherwise all the

[6]This is consistent with a value of each year of extra life of 3 times annual consumption per capita, times 15 years of lost life expectancy conditional on dying after being infected.

parameters are as in the benchmark case. In both cases the lockdown is at least one month shorter.

4.1 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, for different parameter values.

Our preferred 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 all future periods, as well as the values of the lost lives, where every life is evaluated using vsl. The multiplication by r in $rV(S(0), I(0))/w$, converts the expected present values into a permanent annual flow, and the division by w relates it to the output flow before the virus outbreak. We report separately the part of the flow cost $rV(S(0), I(0))/w$ that is purely due to the output cost of the shutdown.[7] The last column displays the present discounted values of the cost if $L(t) = 0$ for all t, which we label as "No Policy" in Table 2. In the three cases we express the losses in percentage.

Importantly, any of the three cost measures in Table 2 can be converted into the equivalent of one year's GDP by dividing them by r, or multiplying them by 20 given our 5% annual interest rate. For instance, dividing by r the flow measure of the cost due to the output lost gives a simple statistic measuring the severity and length of the lockdown. For instance, if this measure is 10%, then the lockdown is equivalent to lose 10% of a year's GDP, or equivalently to lockdown 10% of the population for a year.

The first three rows of Table 2 explore the different values of effectiveness of the lockdown. For the benchmark case, second row in the top panel, 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

[7] The part of the cost due to output is given by $rw \int_0^\infty e^{-(r+\nu)t} \left[1 - \tau + \tau L(t)(S(t) + I(t))\right] dt$.

express it as fraction of one year's GDP, obtaining 28%. From the 1.5% total permanent loss, 0.4% is output loss due to the lockdown, or equivalent to 8% of one year's GDP. For the same parameters, if there is no lockdown (No Policy), the loss is equivalent to a permanent decrease in output of 1.9%.

The second panel corresponds to the case of different values of a statistical life. Recall that the benchmark case assumes a value of a statistical of life (vsl) of 20 annual GDP per capita. We also consider cases of 10, 30, and 80 annual GDP per capita. For a vsl of 30 annual GDP per capita, the part due to output loss is a permanent flow of 0.6% or equivalent to 12% of one year's GDP.

The third panel corresponds to the case where the case fatality rate $\phi(I)$ is constant at $\varphi = 0.01$, or equivalently $\kappa = 0$. In this case, 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 heath care system caused by a large number of infected.

The fourth panel corresponds to the case of no antibody test ($\tau = 0$). For this case, we present different values of a statistical life. Each row expresses vsl as a multiple of the annual GDP per capita and otherwise the same parameters 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 the value of the test. In particular, the expected discounted cost under the optimal policy is, expressed as a permanent flow, 0.1% (10 basis points) 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's GDP. For a smaller value of vsl, we find that the difference is smaller than 0.1%. Instead, as vsl increases to multiples of annual GDP per capita of 30 and 80, the difference measuring the value of the test as a permanent flow of GDP loses is 0.20% and 0.80% respectively. For our preferred parameter values (i.e. vsl between 20 and 30 times annual GDP per capita), the value of the test is equivalent to between 2% and 4%

of one year's GDP.

The last panel of the table contains two cases in which makes the outbreak is less serious, as discussed above. In these two cases, the loses are considerable smaller.

5 Future work

There are several extensions of interest. Our benchmark analysis with linear costs to social activity implies a gradual lockdown. Allowing for non-linear output costs will affect the implementation of the lockdown and shorten its duration. We also overlooked the fact that a long lockdown could have "scarring" effects on the economy that could delay its restart (e.g. it could trigger a cascade of bankruptcies, with long unemployment spells affecting the workers' skills). Second, the quadratic search effects we assumed are a natural starting point under the SIR framework. Alternative matching technologies delivering different speed of transmission seem worth exploring. It would also be interesting to explore the optimal lockdown policy in a setup where social distancing is endogenous (i.e. individuals choose to engage in distancing), since behavioral changes often taken place before governments enact the lockdown. Third, the role of testing should be analyzed in more detail: the planner could choose both the optimal timing and the amount of testing to be delivered. Fourth, the considerable uncertainty surrounding key parameters of the SIR model suggests that a robust control approach is valuable. Lastly, it might be interesting to bring geographic elements and population migration into the picture. We believe all these are important topics for future research which are feasible to incorporate to this setup.

6 Figures and Tables

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	Case fatality rate of 1%
κ	0.05	Implies a 3 percent case 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
vsl	20	Value of Statistical Life 20 $\times$ w (i.e. $vsl \approx \$1.3M$)

Figure 1: **Time paths under baseline parameters**

Panel A – Case w / testing ($\tau = 1$)

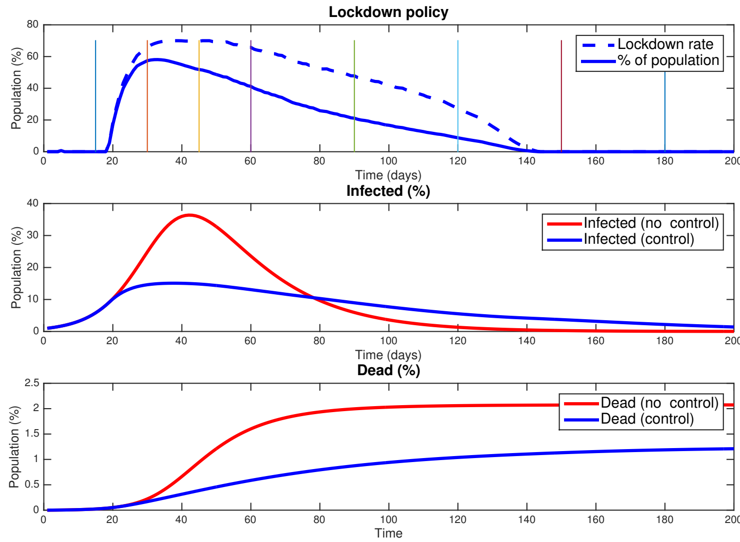

Panel B – Case w/o testing ($\tau = 0$)

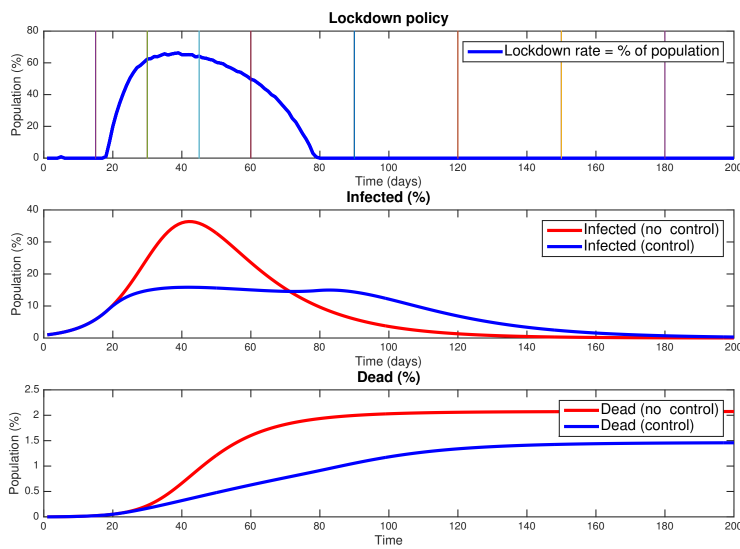

Figure Note: This figure uses the benchmark parameter values of Table 1. Panel A considers the case where a test is available. Panel B considers the case where the test is not available. The red lines describe the uncontrolled system, where no Lockdown is exercised. The blue lines correspond to the optimal control case. The initial condition is $I(0) = 0.01$ and $S(0) = 0.97$.

Figure 2: Value Function and Optimal Policy, benchmark case

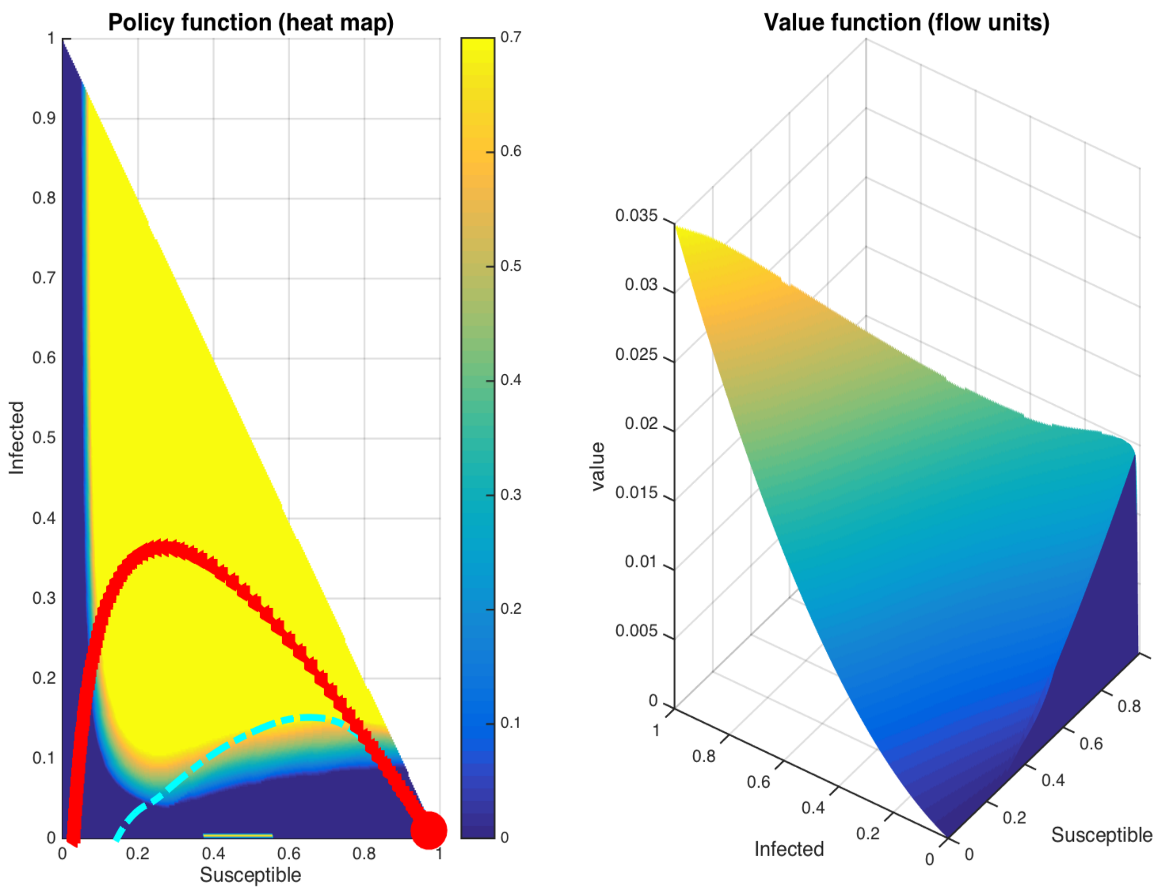

Note: The figure on the left shows the optimal policy for the benchmark parameter values. The blue area indicates lower values of lockdown and the yellow color higher values. The figure on the left depicts the value function. The units for the value function are permanent flow cost as a fraction of the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

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

Case	Parameters	Optimal Policy		No Policy
		Welfare Loss	Output Loss	Welfare Loss
Benchmark Case				
Low effectiveness	θ=0.3	1.7 %	0.3 %	1.9%
Medium effectiveness	θ=0.5	1.5 %	0.4 %	1.9%
High effectiveness	θ =0.7	1.4 %	0.4 %	1.9 %
Alternative Values of Statistical Life				
vsl = 10× GDP per capita		0.9 %	0.2 %	0.9 %
vsl = 30× GDP per capita		2.0 %	0.6 %	2.8 %
vsl = 80× GDP per capita		3.7 %	1.4 %	7.5 %
Constant fatality rate κ=0				
Low effectiveness	θ=0.3	0.9 %	0.0 %	0.9 %
Medium effectiveness	θ=0.5	0.9 %	0.0 %	0.9 %
High effectiveness	θ=0.7	0.9 %	0.0 %	0.9 %
No testing of the recovered $\tau = 0$				
vsl = 10× GDP per capita		0.9 %	0.1 %	0.9 %
vsl = 20× GDP per capita		1.6 %	0.4 %	1.9 %
vsl = 30× GDP per capita		2.2 %	0.6 %	2.8 %
vsl = 80× GDP per capita		4.5 %	2.5 %	7.5 %
Less pessimistic parameter values				
Lower speed of spread of the virus	$\beta = 0.1$	0.8 %	0.1 %	0.8 %
Lower fatality rate	$\varphi = 0.005$	1.1 %	0.4 %	1.5 %

Note: Welfare losses are measured by the permanent percent reduction in per capita GDP induced by the policy (or its absence) under various parameterizations. 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 $S(0) = 0.97$.

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.

Diamond, P. and Eric Maskin. 1979. "An Equilibrium Analysis of Search and Breach of Contracts, I: Steady States." *Bell Journal of Economics* 10:282–316.

Diamond, P.A and Eric Maskin. 1981. "An equilibrium analysis of search and breach of contract II. A non-steady state example." *Journal of Economic Theory* 25 (2):165 – 195.

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandtz. 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, Daniel Laydon, Gemma Nedjati Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, ZULMA Cucunuba Perez, Gina Cuomo-

Dannenburg 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 epidemilogia.” 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.