

A Simple Planning Problem for COVID-19 Lock-down, Testing, and Tracing[†]

By FERNANDO ALVAREZ, DAVID ARGENTE, AND FRANCESCO LIPPI*

We study the optimal lock-down for a planner who controls the fatalities of COVID-19 while minimizing the output costs of the lock-down. The policy prescribes a severe lock-down beginning a few weeks after the outbreak, covering almost 50 percent of the population after a month, with a total duration shy of 4 months. The intensity of the optimal lock-down depends on the gradient of the fatality rate with respect to the infected and the availability of antibody testing, which yields a welfare gain of 2 percent of GDP. We also study test-tracing-quarantine, which we show to be complementary to lock-down. (JEL E23, I12, I15, I18)

We adopt a variation of the SIR epidemiology model of Kermack and McKendrick (1927) to characterize the optimal policy response to the COVID-19 outbreak under several scenarios. The typical approach in the epidemiology literature is to study the dynamics of the pandemic for infected, deaths, and recovered as functions of some exogenously chosen diffusion parameters, which are in turn related to various policies, such as the partial lock-down of schools, businesses, and other measures of diffusion mitigation, where the diffusion parameters are stratified by individual covariates—see, for example, Ferguson et al. (2020). We use a simple version of these models to analyze how to optimally balance the fatalities induced by the epidemic with the output costs of the lock-down policy. Our model thus features a key trade-off between lives saved versus forgone

*Alvarez: University of Chicago (email: f-alvarez1@uchicago.edu); Argente: Pennsylvania State University (email: dargente@psu.edu); Lippi: LUISS University (email: francescolippi@gmail.com). Pete Klenow was coeditor for this article. First draft, March 23, 2020. We benefited from the comments of Andrew Atkeson, Gadi Barlevy, Carlo Favero, 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, University of Chicago Economics Virtual Seminar, OVP’s COVID-19 Forum—Social Sciences Focus, the Riksbank, the Banque de France Virtual Seminar, the Pennsylvania State University Virtual Seminar, and the Bank of Mexico Virtual Seminar. Olivia Bordeu-Gazmuri provided excellent research assistance. We are grateful to Andy Atkeson for suggesting us to work on the tracing-testing-quarantine extension of our simple planning problem. The authors declare to have no conflict of interest to disclose regarding the research on this paper.

[†]Go to <https://doi.org/10.1257/aeri.20200201> to visit the article page for additional materials and author disclosure statement(s).

production. The novel aspect of our analysis is to explicitly formulate and solve a control problem, where the diffusion parameter is affected by the lock-down, chosen to maximize a social objective while taking into account the dynamics of the system.¹ A reason to write a planning problem directly is that with social interactions, there is an externality to be corrected. As understood and analyzed in Eichenbaum, Rebelo, and Trabandt (2020); Farboodi, Jarosch, and Shimer (2020); and Toxvaerd (2020), the voluntary social distance in the absence of a government-imposed lock-down would not be socially optimal. We use global methods because in our setup, the interaction of the law of motion of the SIR model and the lock-down policy makes the problem nonconvex.²

By computing the optimal policy and the associated trajectories, we aim to gauge the key elements that determine the *intensity* and *duration* of the lock-down. We solve the problem under different scenarios, which include congestion effects in the health care system, the effectiveness of the lock-down in reducing the diffusion of the virus, the possibility of testing for antibodies, and the possibility to trace and quarantine infected agents.

We parametrize the model using a range of estimates about the COVID-19 epidemic. Since we recognize that several parameters are highly uncertain, we explore a range of values concerning the severity of the congestion effects on the fatality rate, a range of valuations for the cost of lost lives, and the possibility of testing and releasing the recovered agents from lock-down.

In our baseline parameterization, conditional on a 1 percent fraction of infected agents at the outbreak and no cure for the disease, the optimal policy prescribes a lock-down starting four weeks after the outbreak and covering 45 percent of the population after 8 weeks. The lock-down is kept tight for about a full month and is gradually withdrawn, covering 30 percent of the population 3 months after the initial outbreak. The output cost of the lock-down is high, equivalent to losing 6 percent of one year's GDP (or, equivalently, a permanent reduction of 0.3 percent of output). The total welfare costs is more than four times bigger due to the cost of deaths (see panel A in Figure 1 and Table 1).

The intensity of the optimal lock-down 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 lock-down are significantly reduced and, in some cases, completely eliminated, even though the welfare cost of the pandemic remains high. On the other hand, the value of the statistical life we use in our benchmark case (40 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 makes the abandonment of the lock-down more gradual.

Our benchmark scenario assumes, as seems realistic following the outbreak, that there is no antibody test that allows those who recover to be issued an immunity card

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

²Using first-order conditions to solve for the path of states, controls, and co-states, as well as transversality, provides a necessary but not sufficient condition. That is why we use a discrete-time discrete-state finite difference approach.

and go back to work. We also analyze the problem in a scenario with such a test, in which case the optimal lock-down is longer, but overall it involves similar total number of lost working hours (forgone GDP, see Figure 1 and Table 1). The most salient feature of the case where a test is not available is that the lock-down ends up sooner and more abruptly. The dynamics of the epidemiological model explain why this is optimal: as time goes by, the fraction of those recovered increases, and thus the lock-down becomes 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 an antibody test yields a large welfare gain, in the order of 2 percent of one year's GDP. A by-product of the calculations is the benefit of the lock-down policy, measured as a percentage of permanent GDP flow of following the optimal policy versus the case of no lock-down (see Table 1). Under our preferred values, the total welfare cost of the virus is equivalent to a loss of 28 percent to 32 percent of one year's GDP. From this loss, the part due to the forgone GDP is between 6 percent and 8 percent of one year's GDP.

We conclude with two extensions that are relevant for applications to an actual economy. The first considers a setting where the planner's controls include a trace-test-quarantine (TTQ) instrument. The goal is to understand whether TTQ is a complement or substitute of the lock-down policy. Our findings show that there is a large overlap between the region of the state space where the lock-down is used and the one where the TTQ is used. The second extension takes a step toward realism by considering that one consequence of the lock-down is to "buy time" and allow for better treatment technologies to be available. Such technologies include, among others, antibody tests, more ICU capacity, test-tracing apps, and eventually a vaccine. The expectation of improved instruments to deal with the epidemic creates a *dynamic complementarity*: the incentives for lock-down are strengthened because, by delaying the diffusion, the planner will face the problem with better instruments and thus incur smaller losses. Such complementarities can be powerful: the prospect of a smart-tracing technology may lead to an immediate lock-down in an economy where, absent such prospect, there would be no lock-down at all.

Our simple analysis has limitations: the underlying model has no heterogeneity in fatality rates nor in diffusion rates, and the lock-down policy cannot be differentiated across agent's type; see Favero, Ichino, and Rustichini (2020) for a seminal exploration of the idea of an age- and industry-specific lock-down and Acemoglu et al. (2020) for an extension of our framework that allows for group-specific lock-downs. We also ignore direct health interventions that might be put in place to mitigate the consequences of the disease (e.g., compulsory social distancing and new emergency hospitals).

Several recent papers study similar control problems. Among the most closely related are the following: Farboodi, Jarosch, and Shimer (2020) discuss the quadratic search for both an optimal allocation and an equilibrium; Eichenbaum, Rebelo, and Trabandt (2020) focus on a competitive equilibrium where a consumption tax is used to slow down economic activity and the epidemic diffusion; Garriga, Manuelli, and Siddhartha (2020) study in detail the behavior before and after the appearance of a vaccine; and Gonzalez-Eiras and Niepelt (2020) provide useful analytical characterizations.

I. A Model of Lock-down, Testing, and Tracing

We start with a modified version of the SIR model by Kermack and McKendrick (1927) 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 ; that is,

$$(1) \quad N_t = S_t + I_t + R_t \quad \text{for all } t \geq 0.$$

The “recovered” include those that have been infected, survived the disease, and are assumed to be (forever) immune. Since we only include those that are alive, N_t is changing through time, and normalize the initial population to $N_0 = 1$. The planner can lock-down a fraction $L_t \in [0, \bar{L}]$ of the population, where $\bar{L} \leq 1$ allows us to consider that even in a crisis scenario some economic activity, such as energy and basic food production, will continue. We assume that the lock-down is only partially effective in eliminating the transmission of the virus. When L agents are in lock-down, then $(1 - \theta L)$ agents can transmit the virus, where $\theta \in (0, 1]$ is a measure of the lock-down 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 lock-down, we allow $\theta < 1$.

In addition to the lock-down policy, we assume the planner can “test and trace” infected agents and place them into quarantine. We refer to this policy as a TTQ. Let $Q_t \leq I_t$ denote the stock of quarantined agents at time t . The quarantined agents are removed from the pool of the active infected and thus do not contribute to the propagation of the new infections.

The law of motion of the susceptible agents then is

$$(2) \quad \dot{S}_t = -\beta[S_t(1 - \theta L_t)][(I_t - Q_t)(1 - \theta L_t)].$$

In the case where no control is exercised, $L = 0$ and $Q = 0$, the uncontrolled evolution of the system obeys the well-known $\dot{S} = -\beta SI$ equation, where $\beta > 0$ is the number of susceptible agents per unit of time to whom an infected agent can transmit the virus.

It is evident that locking down a part of the population can be 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, the new infections are reduced by the *square* of the lock-down rate.³ Likewise, if a fraction of the infected is quarantined, that is, if $Q > 0$, the reproduction rate of new infections is controlled by reducing the number of infected agents who have contacts with others, namely $(I - Q)$.

A fraction γ (per unit of time) of the infected recovers, thus:

$$(3) \quad \dot{I}_t = -\dot{S}_t - \gamma I_t.$$

³In search theory, Diamond and Maskin (1979, 1981) aptly named this feature “quadratic search.”

The stock of quarantined agents follows the law of motion

$$(4) \quad \dot{Q}_t = T_t - \gamma Q_t,$$

where $T_t \leq \bar{T}$ denotes the flow per unit of time of agents that are traced, tested (positive), and placed into quarantine, and \bar{T} is a capacity constraint on the number of agents that can be traced per unit of time.

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

$$(5) \quad -\dot{N}_t = \phi(I_t)I_t.$$

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

$$(6) \quad \phi(I) = [\varphi + \kappa I]\gamma.$$

The term $[\varphi + \kappa I] \in (0, 1)$ is the proportion of infected persons that die (“infected fatality rate,” or IFR). It appears that the IFR is increasing with I , an assumption that reflects congestion effects in the health care system. The multiplication by γ gives the fatality rate per unit of time.

Planner’s Objectives: We assume that each agent alive produces w units of output when not in lock-down. Agents are assumed to live forever, unless they die from the infection. The time discount rate is $r > 0$, and we assume that with probability ν per unit of time both a vaccine and a cure appear, so that the planner discount rate is $r + \nu$. The problem consists of minimizing the following present value:

$$(7) \quad \mathcal{V}(S_0, I_0, Q_0) \\ = \min_{\{L_t, T_t\}} \int_0^\infty e^{-(r+\nu)t} \{ wQ_t + wL_t[\tau(S_t + I_t - Q_t) + (1 - \tau)(1 - Q_t)] \\ + \nu s l \phi(I_t)I_t + c(T_t; S_t, I_t, Q_t) \} dt$$

subject to the laws of motion equation (2), equation (3), and equation (4) and an initial condition (S_0, I_0, Q_0) with $I_0 > 0$ and $S_0 + I_0 \geq N_0$. Note that, as the vaccine and cure arrive, there is no more cost, and the continuation value is zero.

The flow cost for the planner of having state (S, I, Q) at t and selecting control L, T has three components. The first one is the output lost due to the lock-down.

Since those in quarantine do not work, there is a cost wQ_t in the period return function. Note also that the lock-down L_t applies to the remaining agents, whose precise value depends on whether there is an antibody test or not. Without an antibody test, or $\tau = 0$, the lock-down implies that all agents not in quarantine are forced out of work. With a test, or $\tau = 1$, the lock-down does not apply to the recovered agents so that only the susceptible and the infected (not in quarantine) are out of production. Note that infected not in lock-down are assumed to produce

as much as those susceptible or recovered not in lock-down. Conversely, agents in lock-down produce zero.⁴

The second component 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), discussed later. In particular, if there are I infected, the deaths per unit of time are given by $\phi(I)I$. The third component is the cost of tracing-testing and quarantining T infected agents per unit of time $c(T; \cdot)$, which we allow to depend on the state (S, I, Q) . The reason is that the cost of tracing one infected agent generally depends on the difficulty to detect such agents, or their prevalence in the population; see Pollinger (2020); Chari, Kirpalani, and Phelan (2020); and Collard et al. (2020).

In what follows we will solve three distinct versions of this problem. The first problem has no testing ($\tau = 0$) and no TTQ ($\bar{T} = 0$). The second problem has testing ($\tau = 1$) and no TTQ. The third problem has both testing ($\tau = 1$) and TTQ ($\bar{T} > 0$). We solve the problem by discretizing the model to daily intervals, using value function iteration over a dense grid for (S, I) .⁵ In the first two cases the state space is two dimensional (since $Q = 0$). Also, in both cases the value function $V(S, I) = \mathcal{V}(S, I, 0)$ has analytic expressions on the boundary of its domain, where the lock-down 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 $V(0, I) = vsl \cdot I\gamma((\varphi/(r + \nu + \gamma)) + (\kappa I/(r + \nu + 2\gamma)))$ for all $I \in (0, 1)$. In the third case, discussed in Section IV, the state is three dimensional.

II. Parameterization of the Baseline Model (without TTQ)

We calibrate β (the propagation rate) to 0.13×365 (that is, 0.13 per day). 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). These parameter values imply an $R_0 = 2.34$, which is close but slightly higher than the one estimated for the United States by Fernández-Villaverde and Jones (2020) (Table 1) but lower than the typical values estimated by Atkeson, Kopecky, and Zha (2020).⁶ We set the fatality rate $\varphi = 0.0068$, or 0.68 percent, which is consistent with the age-adjusted fatality rate estimated from data of 37 countries reported in Verity et al. (2020). We set $\kappa = 0.034$ so the fatality rate is 2 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 uncertain. We set the planner's discount factor r to be consistent with a 5 percent annual interest rate and the rate ν assuming that a vaccine and a cure will appear on average in 1.5 years.

We normalize output $w = 1$ and adopt a baseline value of a statistical life of 40 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, and the shadow cost of each life lost used by the planner is 40 times annual GDP per capita, or about \$2.6 million. Our

⁴Both assumptions can be easily changed by rewriting the flow values of the objective function.

⁵We use the finite difference method described in detail in the online Appendix.

⁶Although Atkeson, Kopecky, and Zha (2020) do not report an R_0 for the entire country, in the states/cities common with Fernández-Villaverde and Jones (2020) (e.g., New York, Pennsylvania, and California), they report higher numbers (at least 25 percent higher).

choice of the benchmark value for $vsI = 40$ is in line with Hall, Jones, and Klenow (2020). These authors use a utilitarian criterion to value the extra years of life lost among those likely to die due to the infection, obtaining a cost of about 60 times per capita annual consumption, which is very close to our benchmark.⁷ The value of 40 annual per capita GDP is much lower than the typical figures for statistical value of life (approximately \$10 million or about 150 GDP per capita), which use a higher life expectancy; see Kniesner and Viscusi (forthcoming). Below we will also report results considering alternative vsI values of 50 and 70 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 operate. These sectors combined account for 25–30 percent of GDP (2018), and we thus 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.

III. Results for the Baseline Model (without TTQ)

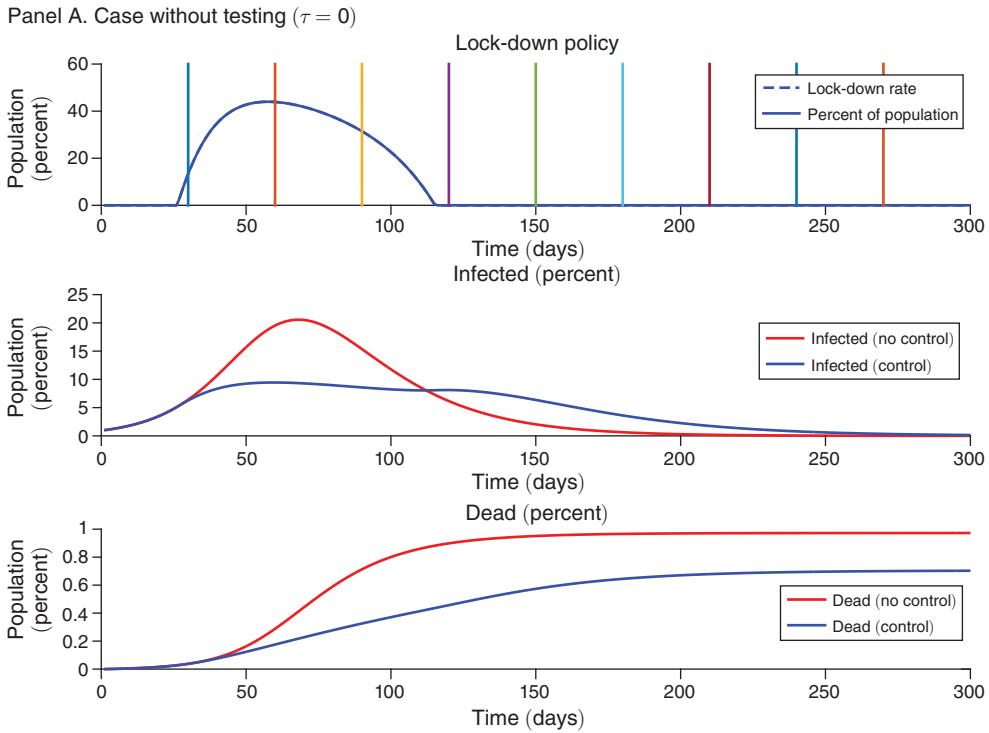
We display the time path of the optimal policy starting at $I_0 = 0.01$, that is, 1 percent of population infected at $t = 0$ for our benchmark parameter values under the assumption that the policymaker does not have access to a TTQ policy.⁸ In particular, we display the time path of the optimal lock-down policy L_t as function of time, the fraction of the population for which lock-down 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 interpreted as fraction of the initial population. In these graphs, conditional on the cure-vaccine not occurring, the horizontal axis denotes time after the outbreak. For comparison, we also plot the path if there is no lock-down policy, that is, for $L_t = 0$ for all $t \geq 0$.

Benchmark Case without Testing: Panel A of Figure 1 shows the results of the case with no test, $\tau = 0$, where the lock-down applies to anybody in the population including those that have recovered from the virus. In this case it is less efficient to lock-down agents because the recovered are also in lock-down, which reduces output without the benefit of reducing the transmission of the virus. As a result, the lock-down peaks at 40 percent about 2 months after the outbreak and ends abruptly before month 4 is reached. 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 lock-down) versus the blue line (optimal policy).

Benchmark Case with Testing: Panel B of Figure 1 presents the result for the case with testing. The lock-down starts four weeks after the outbreak. The fraction of the population in lock-down peaks at 40 percent, about 2 months after the

⁷Following Hall, Jones, and Klenow (2020), one can use that a year of life lost is valued as six 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, $6 \times 10 \text{ years} \times \text{annual consumption per capita} = 6 \times 10 \text{ years} \times 2/3 \times \text{annual GDP per capita} = 40 \times \text{annual GDP per capita}$.

⁸Mathematically this amounts to setting $\bar{T} = 0$. We assume that the initial fraction of the population susceptible is 97 percent, or $S_0 = 0.97$.



(Continued)

FIGURE 1. TIME PATHS UNDER BASELINE PARAMETERS

outbreak, and ends about 6 months after. Interestingly, the lock-down involves similar costs with or without testing in terms of cumulated forgone output, since in the absence of testing the lock-down duration is shorter, but it applies to a larger fraction of people. However, we show in what follows that welfare with testing is higher, in the order of a permanent 0.1 percent GDP flow, which is equivalent to a one-time payment of 2 percent of GDP.

One feature of these two benchmark cases is that the lock-down policy started sooner than the implementation observed in the United States.

Quantifying the Welfare Cost of COVID-19: In Table 1 we summarize the value of following the optimal policy versus the value where there is no lock-down, for different scenarios. Our preferred summary measure is to report $rV(S_0, I_0)/w$, which is the total expected discounted sum of future losses, both due to the lost GDP as well as by the values of the lost lives, where every life is evaluated using *vs**l*. The multiplication by r in $rV(S_0, I_0)/w$, converts the expected present value 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.⁹ The last column displays the present

⁹Since $Q_t = 0$, the output cost of the lock-down is $rw \int_0^\infty e^{-(r+\nu)t} [1 - \tau + \tau L_t(S_t + I_t)] dt$.

Panel B. Case with testing ($\tau = 1$)

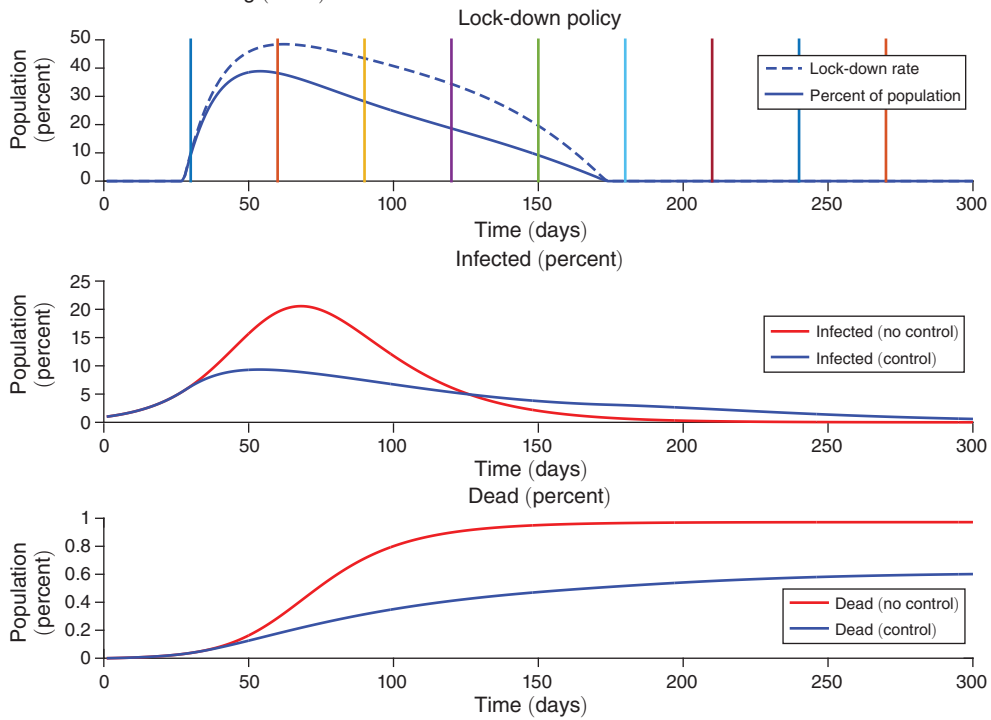


FIGURE 1. TIME PATHS UNDER BASELINE PARAMETERS (CONTINUED)

Notes: Panel A considers the case where the test is not available ($\tau = 0$), panel B the case where a test is available ($\tau = 1$). The red lines correspond to the scenario where no lock-down is exercised, the blue lines to the optimal control case. The parameters are (see Section II) $r = 0.05, \nu = 1/1.5, w = 1, vsI = 40, \gamma = 1/18 \cdot 365, \beta = 0.13 \cdot 365, \varphi = 0.0068, \kappa = 0.034, \theta = 0.5, \bar{L} = 0.70, \bar{T} = 0$. The initial condition is $I_0 = 0.01$ and $S_0 = 0.97$.

discounted values of the cost if $L_t = 0$ for all t , which we label as “No Policy” in Table 1. In all cases we express losses in percentage points.

The first two rows of Table 1 explore the different values of effectiveness of the lock-down. For the benchmark case, in the second row of the top panel, the optimal policy implies a permanent loss of approximately 1.4 percent of output. In other words, as a consequence of the virus, welfare under the optimal policy is equivalent to being 1.4 percent permanently poorer; of this welfare loss about 0.3 percent is due to direct output losses, the rest is due to lost lives.

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 (vsI) of 40 annual GDP per capita. For a higher vsI equal to 50 annual GDP per capita, the part due to output loss is a permanent flow of 0.4 percent (as opposed to 0.3 in the benchmark case). The output costs increases further, and so does the duration of the lock-down, if we consider a vsI equal to 70.

The third panel corresponds to the case where the case fatality rate $\phi(I)$ is constant at $\varphi = 0.0068$, or equivalently $\kappa = 0$. In this case, the optimal policy has no lock-down (so that it coincides with “no policy”), and the welfare losses are much

TABLE 1—WELFARE LOSSES $\left(\frac{rV(S,I)}{w}\right)$ WITH OPTIMAL POLICY VERSUS WITHOUT INTERVENTION

Case	Parameters	Optimal policy (percent)		No policy (percent)
		Welfare loss	Output loss	Welfare loss
<i>Case without antibody test ($\tau = 0$)</i>				
Low effectiveness	$\theta = 0.3$	1.5	0.2	1.6
Medium effectiveness (benchmark)	$\theta = 0.5$	1.4	0.3	1.6
<i>Alternative values of statistical life</i>				
$vsl = 50 \times$ GDP per capita		1.6	0.4	2.0
$vsl = 70 \times$ GDP per capita		2.0	0.6	2.8
<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
<i>Case with antibody test ($\tau = 1$)</i>				
$vsl = 40 \times$ GDP per capita (benchmark)		1.3	0.4	1.6
$vsl = 50 \times$ GDP per capita		1.5	0.4	2.0
$vsl = 70 \times$ GDP per capita		2.0	0.6	2.7
$vsl = 40 \times$ GDP per capita, TTQ	$\zeta = 1$	0.02	0.02	0.9
$vsl = 50 \times$ GDP per capita, TTQ	$\zeta = 1$	0.02	0.02	1.2
$vsl = 70 \times$ GDP per capita, TTQ	$\zeta = 1$	0.03	0.03	1.6
<i>Less pessimistic parameter values</i>				
Lower speed of spread of the virus	$\beta = 0.1$	1.0	0.15	1.0
Lower fatality rate	$\varphi = 0.005$	1.2	0.3	1.3
<i>Quadratic lock-down losses</i>				
Benchmark case		1.3	0.2	1.6
Benchmark case with testing	$\tau = 1$	1.1	0.2	1.6

Notes: 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 are the welfare cost component due to the reduced level of economic activity (that is, excluding fatalities). 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 $S_0 = 0.97$.

smaller than in the benchmark case 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 effects in the health care system.

The fourth panel corresponds to the case with antibody test ($\tau = 1$). 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 first row of the fourth 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 percent (10 basis points) lower with test than without test, i.e., 1.3 percent versus 1.4 percent. For our preferred parameter values (i.e., vsl between 40 and 50 times annual GDP per capita), the value of the test is equivalent to between 2 percent and 4 percent of one year's GDP.

The bottom part of the table contains more scenarios: the fifth panel considers less pessimistic values for the virus and the bottom panel considers a case where the per period output costs are proportional to a quadratic function of the total hours in

lock-down.¹⁰ In the quadratic cost case, the path of the lock-down policy starts at strictly positive values and is smoother through time.

IV. Allowing for a TTQ Policy

In the case where the planner has access to a TTQ protocol, the planner chooses two controls as a function of the state: the rate of tracing-testing-quarantining as well as the lock-down rate. The goal is to understand whether this policy is complementary or substitutable with the lock-down policy and to explore the parts of the state space in which each policy used.

The set-up allowing for TTQ was presented in Section I, where the state space is three dimensional. We focus here on a special case, where we can analyze the optimal policy in a simple yet interesting two-state problem. This case has computational advantages and, more importantly, it facilitates the interpretation and comparison with the previous case when the lock-down policy is the only one available.

Let us define X as the stock of those infected, not in quarantine:

$$(8) \quad X = I - Q$$

so that $\dot{X}_t = \dot{I}_t - \dot{Q}_t$, thus we can write:

$$(9) \quad \dot{S}_t = -\beta S_t X_t (1 - \theta L)^2,$$

$$(10) \quad \dot{X}_t = \beta S_t X_t (1 - \theta L)^2 - T_t - \gamma X_t.$$

The initial conditions of interest are $X_0 = I_0$ and $S_0 = 1 - X_0$, since there is quarantine, and note that $S + X \leq 1$.

The state space reduction is based on two assumptions. To eliminate $\{Q_t\}$ from the state, we rewrite the expected discounted cost of output forgone for those in quarantine, denoted by $\mathcal{C}(\{Q\})$ using integration by parts and the law of motion of Q , as follows:

$$\mathcal{C}(\{Q\}) \equiv \int_0^\infty e^{-(r+\nu)t} Q_t dt = \frac{Q_0}{r + \nu + \gamma} + \int_0^\infty \frac{e^{-(r+\nu)t}}{r + \nu + \gamma} T_t dt$$

which is equivalent to “loading” the expected discounted output cost every time someone is traced and placed in quarantine. The advantage of this formulation is that we can keep track of the output cost due to the quarantine using the contemporaneous control T_t .

Note that this is not yet enough to write the problem as a two-state (S, X) problem, since the return function still requires to have Q . The state-space reduction is based on three assumptions. First, we consider the case with $\tau = 1$, i.e., the one with an antibody test, which is convenient since the flow cost of forgone output cost of lock-down is simply $S + X$.

¹⁰The constant of proportionality is chosen so that if the same path for lock-down is followed in the benchmark and quadratic cases, the expected discounted costs are equal. See online Appendix for more details.

Second, we assume the following function for the cost of tracing T agents (the cost does not include the forgone output):

$$(11) \quad c(T, S, X) = \eta\left(T\left(\frac{S+X}{X}\right)^{1-\zeta}\right),$$

where η is a weakly increasing, positive, and convex function and where $\zeta \in [0, 1]$ indexes how smart the tracing is. If $\zeta = 0$, then there is no tracing and it is just random sampling. If $\zeta = 1$, then tracing is very powerful, the fraction in the population is immaterial, and the cost depends only on the number to be traced. Note that we allow the cost to depend on the flow of those traced-tested-quarantined, T , as well as on the composition of the state (S, X) . We assume that the function $c(T; S, X)$ is increasing and convex in T , for fixed (S, X) , and that $c(0; S, X) = 0$.

Equation (11) implies that the cost of finding T people that are infected (not currently in quarantine) depends on the size of X in the population that is being traced. In one extreme, if testing is random ($\zeta = 0$), the number of people that have to be tested to identify T is $T(S+X)/X$. Simply put, it is harder to find someone infected if there are very few infected in the population and we search at random. If, instead, there is a smart tracing technology ($\zeta = 1$), the cost is independent of the (S, X) state (composition of the pool).

Third, we note that the last term where Q shows up is the number of deaths per unit of time $\phi(X+Q)(X+Q) = \phi(I)I$, which depends on the total number of infected $I = Q+X$, regardless of whether they are in quarantine or not. This can be dispensed with if we consider the case in which the fatality rate function $\phi(\cdot)$ is constant, that is, $\kappa = 0$, so that $\phi(X+Q) = \varphi\gamma(X+Q)$, where φ and γ are constant parameters.¹¹

Numerical Examples: We display the heat map of policies L and T and the value function for the baseline parameters, setting $\tau = 1$ and $\kappa = 0$ as assumed previously. We set the upper bound on tracing-testing flow to $\bar{T} = 1$, that is, testing at a speed such that the entire population will be tested in a year, and we set the testing-tracing cost to be quadratic, that is, $\eta(z) = z^2\alpha/2$, with $\alpha = 0.02$. With this specification, if $z = T((S+X)/X)^{1-\zeta} = 1$, the cost will be $\alpha/2 = 0.01$. We consider the two extreme values of the effectiveness of tracing $\zeta \in \{0, 1\}$. For instance, if $\zeta = 1$, then $z = 1$ means that with “perfect” tracing it takes one task to test-track one infected agent. Instead, with $\zeta = 0$, tracing one infected agent requires running $((S+X)/X)$ tasks.

Figure 2 shows the optimal policy and value function obtained from the baseline parameterization under two alternative assumptions about the efficiency of tracing. The upper panel assumes random tracing ($\zeta = 0$)—that is, to find infected individuals, the population has to be tested at random. This makes testing very expensive, especially when there is a small number of infected, as shown by equation (11), which diverges as $X \rightarrow 0$. As a result, the policy makes no use of testing, nor of the lock-down. This is illustrated by the path highlighted in the phase diagram of

¹¹ See proposition in the online Appendix.

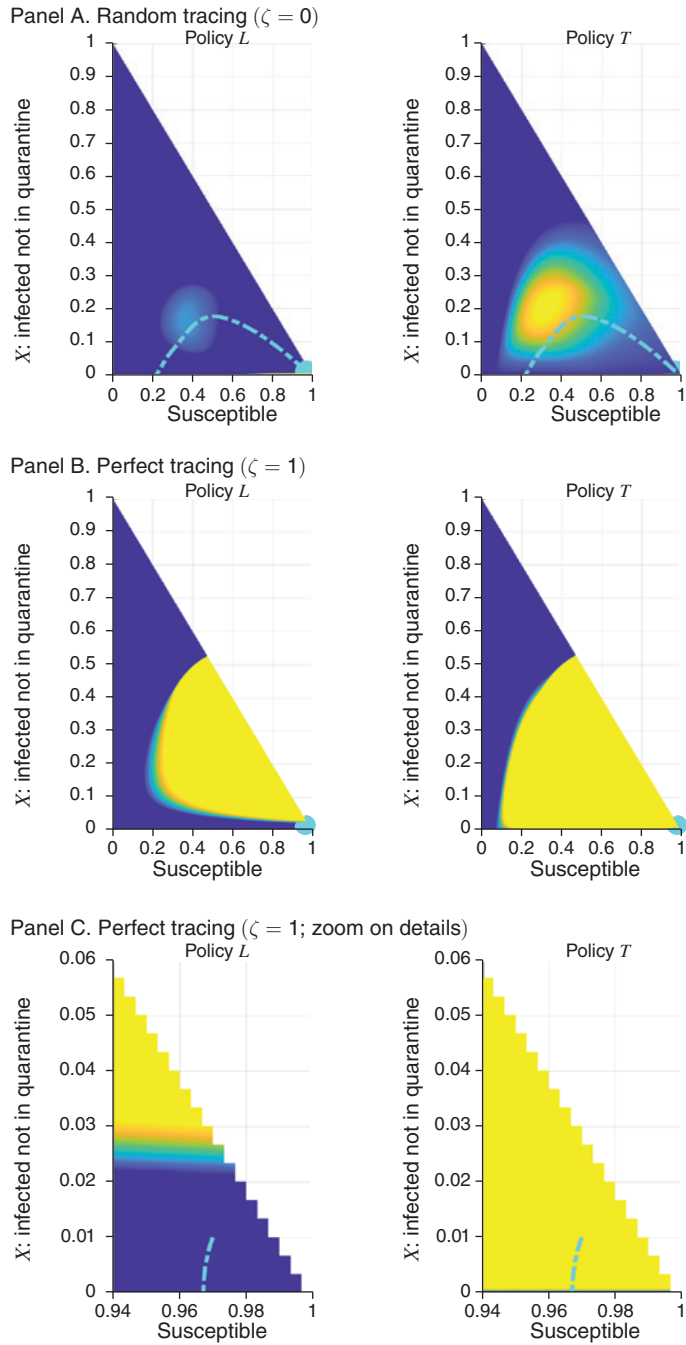


FIGURE 2. OPTIMAL POLICY WITH TTQ

Notes: Panel A assumes random tracing ($\zeta = 0$). Panels B and C assume perfect tracing ($\zeta = 1$). The baseline parameters are as in Figure 1 with $\tau = 1$ and $\kappa = 0$ (see Section IV); the TTQ parameters are $\bar{T} = 1$ and $\alpha = 0.02$.

panel A. Recall that, as explained previously, the model assumes a constant fatality rate ($\kappa = 0$) so that lock-down was not chosen even in the absence of the TTQ policy. Therefore, the overall welfare cost of this economy is the same one that was recorded in the absence of the TTQ policy following an outbreak with a 1 percent of infected agents. This cost is about 0.9 percent of annual GDP (see the third panel of Table 1).

Instead, when tracing is perfectly efficient ($\zeta = 1$), the policy changes substantially, as can be seen from panel B. Tracing, testing, and quarantining even a tiny fraction of the population becomes substantially cheaper now, since they can be immediately identified as opposed to being searched at random. Since those agents can be easily discovered, it is optimal for the policymaker to trace them and quarantine them. The resulting plan, following an outbreak with a 1 percent of infected agents, yields a cost that is about 0.2 of annual GDP, a value that is about 4 times smaller than the cost of the benchmark case without the TTQ policy (panel C of Figure 2 shows that the optimal policy squashes the fraction of infected to zero and eradicates the virus).

V. Dynamic Complementarities of Lock-down

This section extends the basic model by considering that one consequence of the lock-down is to allow the planner to “freeze the state” while setting up better technologies to deal with the epidemic in the near future.¹² Such technologies, unavailable at the time of the outbreak, include among others antibody tests, more ICU capacity, test-tracing apps, and eventually a vaccine. The expectation of improved instruments to deal with the epidemic creates a *dynamic complementarity*: the incentives for lock-down are strengthened because, by delaying the diffusion, the planner will face the problem with better means and thus incur smaller losses. Indeed, we find that the delayed optimal lock-down that was found in the analysis of Section III, as illustrated, for example, in Figure 1 where the lock-down is prescribed about a month after the outbreak, may transform into a policy of immediate and full lock-down by such dynamic complementarities.

Next we solve the optimal control problem described in Section I and assume a deterministic timeline for the set of policy instruments available to the planner. To keep this simple, we assume that at the time of the outbreak (conventionally $t = 0$) the only tool available to the planner is a generalized lock-down, that is, one with $\tau = 0$. We then assume that at a future date $t = \mathcal{T} > 0$, the planner’s toolkit is enriched by a new technology. Thus, when computing the present value of a policy at any $t \in (0, \mathcal{T})$, the planner takes into account that the future costs of the infection will be easier to manage because of the new technology.

Different scenarios can be analyzed that differ in terms of *what technologies* become available at date \mathcal{T} . For concreteness we focus here on a scenario where both the antibody test ($\tau = 1$) and the TTQ protocol analyzed in Section IV become

¹²We thank Andy Atkeson who suggested for us to work on this question.

available for $t \geq \mathcal{T}$.¹³ Since we already computed the value functions for each of these technologies, the problem is easily solved by backward induction.

Our interest is to compare the solution of the problem with $\mathcal{T} < \infty$ with the one where the new technology is never available for the same initial conditions, that is, the same (Q_0, X_0, S_0) . The difference in the path of the lock-down policy for $t \in [0, \mathcal{T})$ is informative about the dynamic complementarity of lock-down with other policies-technologies.

We used the parameters of the benchmark calibration from Section IV ($vsl = 40, \kappa = 0$) and the same initial condition of an outbreak starting with a 1 percent of infected. Moreover, we assume that starting in 60 days from the outbreak ($\mathcal{T} = 60$ days), the policymaker will have access to antibody testing and a smart tracing technology ($\zeta = 1$). Recall that in the baseline case where there is no anticipation of future policy improvements, the optimal policy for this case involves no lock-down at all due to the lack of congestion effects ($\kappa = 0$; see Table 1 and panel A of Figure 2). Instead, in this new setting the anticipation of the expected policy improvements leads the policymaker to immediately jump into a full lock-down as soon as the outbreak occurs, and to keep it until $t = \mathcal{T}$. Thus, lock-downs are dynamically complementary to the arrival of TTQ. This is intuitive, since smart tracing is best used when there are few infected. It is also easy to show that the strength of this dynamic complementarity is weaker if the time until the innovation arrival is longer (higher \mathcal{T}) or the vsl is smaller.

VI. Future Work

There are several extensions of interest of our setup. We overlooked the fact that a long lock-down could have “scarring” effects on the economy that could delay its restart (for example, 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 speeds of transmission seem worth exploring. It would also be interesting to explore the optimal lock-down policy in a setup where social distancing is endogenous, since behavioral changes take place independently of the lock-down. Third, 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.

REFERENCES

- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston.** 2020. “A Multi-risk SIR Model with Optimally Targeted Lockdown.” <https://mitsloan.mit.edu/shared/ods/documents/?PublicationDocumentID=7393>.
- Alvarez, Fernando, David Argente, and Francesco Lippi.** 2021. “Replication Data for: A Simple Planning Problem for COVID-19 Lockdown, Testing, and Tracing.” American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E120689V1>.

¹³This can be easily extended to other cases. For instance, we have also analyzed the case with congestion effects (i.e., $\kappa > 0$), where for $t \geq \mathcal{T}$ the antibody testing becomes available (but not the TTQ), so that $\tau_t = 1$ from $t \geq \mathcal{T}$. In this case, using the parameters of the benchmark calibration, the optimal lock-down policy starts a few days earlier.

- Atkeson, Andrew.** 2020. "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios." Federal Reserve Bank of Minneapolis Staff Report 595.
- Atkeson, Andrew, Karen Kopecky, and Tao Zha.** 2020. "Estimating and Forecasting Disease Scenarios for COVID-19 with an SIR Model." National Bureau of Economic Research Working Paper 27335.
- Chari, Varadarajan V., Rishabh Kirpalani, and Christopher Phelan.** 2020. "The Hammer and the Scalpel: On the Economics of Indiscriminate versus Targeted Isolation Policies during Pandemics." National Bureau of Economic Research Working Paper 27232.
- Collard, Fabrice, Christian Hellwig, Tiziana Assenza, Sumudu Kankanamge, Martial Dupaigne, Nicolas Werquin, and Patrick Fève.** 2020. "The Hammer and the Dance: Equilibrium and Optimal Policy During a Pandemic Crisis." Center for Economic and Policy Research Working Paper 14731.
- Diamond, Peter A., 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): 165–95.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt.** 2020. "The Macroeconomics of Epidemics." National Bureau of Economic Research Working Paper 26882.
- Farhoodi, Maryam, Gregor Jarosch, and Robert Shimer.** 2020. "Internal and External Effects of Social Distancing in a Pandemic." National Bureau of Economic Research Working Paper 27059.
- Favero, Carlo A., Andrea Ichino, and Aldo Rustichini.** 2020. "Restarting the Economy While Saving Lives under COVID-19." Center for Economic and Policy Research Discussion Paper 14664.
- Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia et al.** 2020. *Impact of Non-pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand.* London: Imperial College COVID-19 Response Team.
- Fernández-Villaverde, Jesús, and Charles I. Jones.** 2020. "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities." National Bureau of Economic Research Working Paper 27128.
- Garriga, Carlos, Rody Manuelli, and Sangh Siddhartha.** 2020. "Optimal Management of an Epidemic When a Vaccine is on the Horizon: An Application to COVID-19." Federal Reserve Bank of St. Louis.
- Gonzalez-Eiras, Martín, and Dirk Niepelt.** 2020. "On the Optimal 'Lockdown' during an Epidemic." CESifo Working Paper 8240.
- Hall, Robert E., Charles I. Jones, and Peter J. Klenow.** 2020. "Trading Off Consumption and COVID-19 Deaths." National Bureau of Economic Research Working Paper 27340.
- Hansen, Elsa, and Troy Day.** 2011. "Optimal Control of Epidemics with Limited Resources." *Journal of Mathematical Biology* 62: 423–51.
- Kermack, William Ogilvy, and Anderson G. McKendrick.** 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 115 (772): 700–721.
- Kniesner, Thomas J., and W. Kip Viscusi.** Forthcoming. "The Value of a Statistical Life." *Oxford Research Encyclopedia of Economics and Finance.* <https://doi.org/10.1093/acrefore/9780190625979.013.138>.
- Pollinger, Stefan.** 2020. "Optimal Case Detection and Social Distancing Policies to Suppress COVID-19." Toulouse School of Economics Working Paper 1109.
- Toxvaerd, Flavio.** 2020. "Equilibrium Social Distancing." Cambridge Institute for New Economic Thinking Working Paper 2020/08.
- Verity, Robert, Lucy C. Okell, Ilaria Dorigatti, Peter Winskill, Charles Whittaker, Natsuko Imai, Gina Cuomo-Dannenburg et al.** 2020. "Estimates of the Severity of Coronavirus Disease 2019: A Model-Based Analysis." *Lancet Infectious Diseases* 20 (6): P669–77.