

Gonzalez-Eiras, Martín; Niepelt, Dirk

Working Paper

On the Optimal "Lockdown" during an Epidemic

CESifo Working Paper, No. 8240

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Gonzalez-Eiras, Martín; Niepelt, Dirk (2020) : On the Optimal "Lockdown" during an Epidemic, CESifo Working Paper, No. 8240, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/216636>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

On the Optimal “Lockdown” during an Epidemic

Martín Gonzalez-Eiras, Dirk Niepelt

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

On the Optimal “Lockdown” during an Epidemic

Abstract

We embed a lockdown choice in a simplified epidemiological model and derive formulas for the optimal lockdown intensity and duration. The optimal policy reflects the rate of time preference, epidemiological factors, the hazard rate of vaccine discovery, learning effects in the health care sector, and the severity of output losses due to a lockdown. In our baseline specification a Covid-19 shock as currently experienced by the US optimally triggers a reduction in economic activity by two thirds, for about 50 days, or approximately 9.5 percent of annual GDP.

JEL-Codes: I180.

Keywords: epidemic, pandemic, lockdown, social distancing, production shortfall, health care system, Covid-19, SIR model, logistic model.

Martín Gonzalez-Eiras
University of Copenhagen
Oster Farimagsgade 5
Denmark – 1353 Copenhagen K
mge@alum.mit.edu

Dirk Niepelt
Study Center Gerzensee
P.O. Box 21
Switzerland – 3115 Gerzensee
dirk.niepelt@sgerzensee.ch
www.niepelt.ch

April 13, 2020

1 Introduction

To “flatten the curve” of severe respiratory tract infections caused by Covid-19, policy makers around the world have imposed strict social distancing measures and partial lockdowns. In that context, a first-order policy question is how strict such measures should be and for how long they should be imposed. In this note, we propose two simple models based on the classical epidemiological framework with an embedded policy choice to address this question.

Our analysis starts from a framework with “susceptible,” “infected,” and “removed” (deceased or fully recovered) persons in the tradition of the classical article by Kermack and McKendrick (1927). In that framework transitions between the subgroups with different health status are governed by epidemiological parameters. We augment this framework by allowing for a policy choice—reflecting the level of economic activity—to affect infection rates. Higher activity increases production but also raises the rate of infections, causing future production shortfalls due to death as well as an overburdening of the health care system. Since current and future production shortfalls and health care costs enter society’s loss function the government’s program is a dynamic one: to select the optimal path for activity (or social distancing or lockdowns).

This dynamic program cannot be solved in closed form. In parallel, ongoing work Alvarez et al. (2020) therefore numerically solve for the optimal policy path. We pursue a complementary approach: We simplify the epidemiological framework slightly and build two nested, much more tractable models. One of them can be solved in closed form and the other can “nearly” be solved in closed form. Together, the two models offer transparent and easy-to-compute answers to the policy question at hand. We believe that this is valuable, in particular when information about an infectious disease—like Covid-19 now—is sparse and the task is to gain a first, basic understanding of the tradeoffs at work.

We allow the government’s program to reflect several factors that prominently feature in policy discussions. For example, one of our models features convex costs of flows from the susceptible to the infected population, introducing a role for policies that flatten the curve. Similarly, our other model includes learning effects in the health care sector which introduce a role for delaying such flows. The learning effects reflect the fact that except for a few countries mostly in the Far East (that had experienced similar outbreaks in the past), governments and public health agencies across the globe were left scrambling after the surge in Covid-19 infections; over time, we should expect experience and more adequate supplies of equipment to relax some of the current bottlenecks.

When we calibrate the models using information about the projected death toll, health care stress, and output losses in the US due to the current Covid-19 shock we find that the optimal lockdown is quite severe and prolonged: Activity is optimally reduced by two thirds, for roughly 50 days. We conduct a series of robustness checks and find that all resulting model predictions are in the same ballpark.

As mentioned before our work is closely related to ongoing work by Alvarez et al. (2020). Other recent contributions that merge basic epidemiology and economics include Atkeson (2020), Eichenbaum et al. (2020), and Stock (2020). For discussions of the broader policy options, see for example Baldwin and Weder di Mauro (2020a; 2020b).

2 The Model

Our analysis is based on the canonical epidemiological model (the SIR model) due to Kermack and McKendrick (1927). We simplify that framework to improve tractability and imbed policy decisions that capture the severity and duration of a “lockdown.” In this section we review the SIR model and introduce policy objective and instrument.

2.1 SIR Model

The SIR model specifies laws of motion in continuous time for the population shares of three groups that differ with respect to health status. The three groups are the “susceptible,” the “infected,” and the “removed,” and their respective population shares at time $t \geq 0$ are denoted by $x(t)$, $y(t)$, and $z(t)$, where $x(t) + y(t) + z(t) = 1$. We normalize the population size to unity. Accordingly, the population shares $x(t)$, $y(t)$, and $z(t)$ correspond to the “number” of susceptible, infected, and removed persons.

At time $t = 0$ the population consists of $x(0)$ susceptible persons and a few infected persons, $y(0)$. There are no removed persons at this time, $z(0) = 0$. In each instant after time $t = 0$, the infected transmit their infection to the susceptible and a fraction of the infected either dies or develops resistance. Formally, following Bohner et al. (2019), the change of the number of susceptible, infected, and removed persons, respectively, satisfies

$$\dot{x}(t) = -b(t)x(t)\frac{y(t)}{x(t) + y(t)}, \quad (1)$$

$$\dot{y}(t) = -\dot{x}(t) - (c^d + c^r)y(t), \quad (2)$$

$$\dot{z}(t) = (c^d + c^r)y(t). \quad (3)$$

Here, $b(t)$ denotes a possibly time-varying infection rate. As in Bohner et al. (2019) it reflects epidemiological factors which we take as exogenously given. Unlike Bohner et al. (2019) we allow $b(t)$ to also reflect government policy (see below). The extent to which susceptible persons are infected depends on their number, $x(t)$; the infection rate, $b(t)$; and the share of the infected in the susceptible or infected population.

The number of infected persons increases one-to-one with each susceptible that gets infected. At the same time, a share $c \equiv c^d + c^r$ of the infected population dies or recovers; the coefficients c^d and c^r parameterize the flow into death and recovery, respectively.

The system (1)–(3) can be solved (see appendix A) for

$$x(t) = x(0)e^{\int_0^t \frac{-\kappa b(u)}{\kappa + e^{\int_0^u (c-b(s))ds}} du}, \quad (4)$$

$$y(t) = y(0)e^{\int_0^t \frac{b(u)}{1 + \kappa e^{\int_0^u (b(s)-c)ds}} - c du}, \quad (5)$$

$$z(t) = 1 - x(t) \left(1 + \kappa e^{\int_0^t (b(s)-c)ds} \right) \text{ s.t. (4),} \quad (6)$$

where $\kappa \equiv y(0)/x(0)$.




Figure 1: Dynamics in the SIR model: $x(t)$ (solid), $y(t)$ (dashed), and $z(t)$ (recovered and deceased, dotted).

Figure 1 illustrates the dynamics when we let $b(t) = \beta$, the fundamental infection rate absent any policy intervention.¹ We measure time in days and let $\beta = 0.2$, $c^r = (0.05)(0.99)$, and $c^d = (0.05)(0.01)$.²

2.2 Policy Objective and Instrument

The basic tradeoff we are interested in is the conflict between fostering economic activity and slowing down the spread of infections. Almost all countries that have responded to the spread of Covid-19 by imposing severe restrictions on mobility and economic activity have motivated these restrictions with the aim to delay infections or to “flatten the curve,” i.e., to reduce the speed at which infections occur. The main argument for delay is to gain time in which health care providers can prepare for the higher case load. The main argument for flattening the curve is to limit the stress that Covid-19 infections impose on the health care system—specifically on intensive care units—because this stress increases fatality rates. In the SIR model outlined above, both measures to delay and to flattening the curve correspond to policy interventions that push $b(t)$ below β .

Let $a(t) \in \mathcal{A}$ denote a measure of economic activity inversely related to lockdown policies such as social distancing, forced shutdowns of businesses, etc. The maximum element of \mathcal{A} is unity, representing the regular level of activity. The minimum element of \mathcal{A} (which is nonnegative) represents the lower bound on activity or upper bound on lockdown policies. This minimum could be strictly positive, reflecting the fact that even

¹Equation (2) implies that at the beginning of an epidemic when $x(t) \approx 1$ and $z(t) \approx 0$, rate β equals the growth rate of the number of persons who are or were infected:

$$\frac{\dot{y}(t) + \dot{z}(t)}{y(t) + z(t)} = \beta y(t) \frac{x(t)}{x(t) + y(t)} \frac{1}{y(t) + z(t)} \approx \beta.$$

²We take the value for β from Alvarez et al. (2020) and assume that 5 percent of the infected are removed, of which 1 percent dies.

during an extreme lockdown elementary goods and services need to be produced (e.g., in the food, energy, or health care sector) or that political constraints prevent extreme containment policies.

Activity $a(t)$ increases the spread of infections, which imposes a burden on the health care system, and it raises output. We assume that the infection rate satisfies

$$b(t) = \beta f(a(t))$$

for some smooth increasing function f . Moreover, we assume that output is a smooth increasing function g of activity which may also depend on the population shares,

$$\text{output}(t) = g(a(t), x(t), y(t), z(t)),$$

and satisfies $g(1, 1, 0, 0) = 1$ (i.e., we normalize output in “normal” times to unity). Finally, we assume that the burden that new infections impose on the health care system is a smooth increasing function h of $\dot{y}(t)$,

$$\text{burden}(t) = h(\dot{y}(t), t).$$

Let ρ denote the rate of time preference and ν the hazard rate with which a new vaccine is discovered. The policy problem then reads

$$\begin{aligned} \max_{(a(t))_{t=0}^{\infty}} \quad & \int_0^{\infty} e^{-(\rho+\nu)t} \{g(a(t), x(t), y(t), z(t)) - h(\dot{y}(t), t) + \nu V(1, x(t), y(t), z(t))\} dt \\ \text{s.t.} \quad & (1), (2), (3), \quad b(t) = \beta f(a(t)), \quad a(t) \in \mathcal{A}, \quad x(0) \text{ given.} \end{aligned}$$

The last term in the integral reflects the probability weighted value, V , of exiting the lockdown due to the discovery of a vaccine.

We are interested in analytical characterizations of optimal paths for $a(t)$ and the implied paths for $x(t)$, $y(t)$, and $z(t)$. The system (1)–(3) or (4)–(6) is not suitable for such characterizations. When we leave $a(t)$ unrestricted (subject to $a(t) \in \mathcal{A}$) and form the Hamiltonian that reflects (1)–(3) and the policy objective then the Hamiltonian does not yield closed-form solutions even if we assume tractable functional forms for f , g , and h .³ Similarly, restricting $a(t)$ to belong to a class of functions that is parameterized by a few parameters and maximizing the intertemporal objective subject to the constraints (4)–(6) very quickly becomes analytically intractable as well.

Against this background, we simplify the epidemiological framework in order to express health dynamics in terms of a single rather than two state variables.⁴ In the next section, we lay out these simplifications and solve the policy problem.

3 Analysis

We analyze two specialized models which are nested by the general model introduced above. We assume, realistically for most countries, that available tests for infection and

³See Alvarez et al. (2020) for a numerical approach to solving a related problem.

⁴The SIR model features two state variables, $x(t)$ and $y(t)$. The third variable, $z(t)$, is implied by the former two.

immunity are scarce, limiting the government's options to indiscriminate lockdowns of varying intensity and duration.

3.1 Model 1

To obtain the first model we simplify along two dimensions. First, we neglect deaths and let $c^d = 0$.⁵ Importantly, this does not mean that we disregard the burden that infections impose on the health care system, to the contrary. This burden depends on the outflow from susceptibles, not on the number of deceased.

Second, we blur the distinction between infected and recovered. While we maintain the feature of the SIR model that infection rates reflect the interaction between susceptible and infected persons we assume that infected persons are as productive as healthy ones. Formally, we let $c^r = 0$ such that infection is an absorbing state and $z(t) = 0$, and we assume that production does not depend on the population shares $x(t)$ and $y(t)$. Stated differently, we view $x(t)$ as the population share of the “not yet infected” and $y(t) = 1 - x(t)$ as the share of the “infected but still productive.” Since members of the two groups are equally productive the function g does not depend on population shares and V satisfies

$$V(1, x(t), y(t), z(t)) = \int_{j=0}^{\infty} e^{-\rho j} g(1, 1, 0, 0) dj = g(1, 1, 0, 0)/\rho = \rho^{-1}.$$

Regarding functional forms, we let $f(a(t)) = a(t)$, $\text{output}(t) = a(t)$, and $h(\dot{y}(t), t) = h_1 e^{-\lambda t} \dot{y}(t)$.⁶ That is, we let activity have a proportional effect on the infection rate and on output and we assume that stress in the health care system is proportional to $\dot{y}(t)$ and a factor $h_1 e^{-\lambda t}$. The parameter λ represents the speed of learning or efficiency enhancing measures in the health care sector. A strictly positive λ generates a motive to delay infections until society is better equipped to confront the stress imposed on the health care sector.⁷ Finally, we let $\mathcal{A} = [\bar{a}, 1]$ with $\bar{a} > 0$. Accordingly, the government's program reads

$$\begin{aligned} \max_{(a(t))_{t=0}^{\infty}} \quad & \int_0^{\infty} e^{-(\rho+\nu)t} \{a(t) - h_1 e^{-\lambda t} \dot{y}(t) + \nu/\rho\} dt \\ \text{s.t.} \quad & \dot{y}(t) = a(t)\beta y(t)(1 - y(t)), \quad a(t) \in [\bar{a}, 1], \quad y(0) \text{ given.} \end{aligned}$$

When we abstract from policy model 1 is identical to the SIR model except that $c = 0$; that is, the time paths of $x(t)$ and $y(t)$ follow logistic curves. Figure 2 illustrates the dynamics when $a(t) = 1$ and when we reduce β by a factor of 0.8 relative to the value underlying figure 1 in order to better match the dynamics of $x(t)$ in the SIR model.

⁵Recall from figure 1 that the population share of deceased is small even in the absence of any policy intervention.

⁶Unlike Alvarez et al. (2020) we assume that the effect of activity on the infection rate is linear rather than quadratic. Recall from equation (1) that the infection rate $\dot{x}(t)$ depends on $x(t)$ as well as the number of infected relative to the number of infected or susceptible. The latter ratio does not change with a lockdown.

⁷In model 2, we explicitly model the motivation to flatten the curve in order to smooth convex stress in the health care system over time.




Figure 2: Dynamics in model 1 absent policy intervention: $x(t)$ (solid), $\dot{y}(t)$ (scaled, dashed), and $y(t) = (1 - x(t))$ (dotted).

Note that the time paths of the shares of “infected” and “not yet infected” are very similar in the two models.

Optimal Policy To characterize the optimal policy we form the current value Hamiltonian

$$H^c(t) = a(t) - h_1 e^{-\lambda t} a(t) \beta y(t) (1 - y(t)) + \frac{\nu}{\rho} + \mu(t) a(t) \beta y(t) (1 - y(t)),$$

where $\mu(t)$ denotes the co-state variable associated with the state variable $y(t)$. The derivative of $H^c(t)$ with respect to the control variable $a(t)$ yields

$$\frac{dH^c(t)}{da(t)} = 1 - \beta y(t) (1 - y(t)) (h_1 e^{-\lambda t} - \mu(t)). \quad (7)$$

Since this derivative does not depend on $a(t)$ the control variable typically is in a corner: either $a(t) = \bar{a}$ or $a(t) = 1$.

The law of motion for the co-state variable is given by

$$\dot{\mu}(t) = -\frac{dH^c(t)}{dy(t)} + (\rho + \nu)\mu(t) = a(t)\beta(1 - 2y(t))(h_1 e^{-\lambda t} - \mu(t)) + (\rho + \nu)\mu(t).$$

Finally, the time derivative of the effect of the control variable on the Hamiltonian equals

$$\begin{aligned} \left(\frac{dH^c(t)}{da(t)} \right) &= -\beta(1 - 2y(t))(h_1 e^{-\lambda t} - \mu(t))a(t)\beta y(t)(1 - y(t)) \\ &\quad + \beta y(t)(1 - y(t)) [a(t)\beta(1 - 2y(t))(h_1 e^{-\lambda t} - \mu(t)) + (\rho + \nu)\mu(t) + \lambda e^{-\lambda t}], \\ &= \beta y(t)(1 - y(t)) [(\rho + \nu)\mu(t) + \lambda e^{-\lambda t}]. \end{aligned} \quad (8)$$

Note that $\mu(t)$ is the shadow value of the population share of the infected, $y(t)$. An increase in this share has no direct effect on output but *reduces* the future burden on the health care system since $\dot{y}(t)$ is strictly positive until everybody is infected.⁸ Accordingly,

⁸Recall that $\bar{a} > 0$.

$\mu(t) > 0$. Combining this result with equation (8) implies that the effect of $a(t)$ on the Hamiltonian is monotonically increasing over time.

We conclude that there are two cases to distinguish: Either $dH^c(0)/(da(0)) > 0$ and the optimal policy does not involve a lockdown. For given $y(0)$ this condition is satisfied when h_1 , which parameterizes the burden that infections impose on the health care system, is low. Or, if $dH^c(0)/(da(0)) < 0$ (which is the case for sufficiently high h_1) the optimal policy immediately imposes a lockdown. Such a lockdown cannot be permanent however and in fact, it ends before all persons have been infected. For as long as $\mu(t)$ is bounded the effect of the control on the Hamiltonian sooner or later becomes positive since $\lim_{t \rightarrow \infty} y(t) = 1$ and therefore $\lim_{t \rightarrow \infty} dH^c(t)/da(t) = 1 > 0$.⁹

3.2 Model 2

To obtain the second model we simplify along different dimensions. First, we neglect recovery ($c^r = 0$) and assume a strictly positive death rate ($c^d > 0$) such that everyone who transits from susceptible to infected eventually dies. Second, we blur the distinction between susceptible and infected assuming that the two groups are equally productive and that only their total share, $x(t) + y(t) = 1 - z(t)$, is relevant for the economy's dynamics. That is, we assume that $x(t)$ and $y(t)$ can be characterized by a single state variable.

Given the laws of motions (1) and (2) this requires that the relative share $x(t)/(x(t) + y(t))$ remains constant over time. Since $y(0)/x(0) = \kappa$ this implies $y(t)/(x(t) + y(t)) = \kappa(1 + \kappa)^{-1}$ and consistency with the laws of motion then entails $c^d = \beta$, which we assume to hold.¹⁰ Absent policy, the system (1)–(3) therefore simplifies to

$$\begin{aligned} (1 - z(t)) &= -\tilde{\beta}(1 - z(t)) \quad (= -\beta y(t)), \\ \dot{z}(t) &= \tilde{\beta}(1 - z(t)) \quad (= \beta y(t)), \end{aligned}$$

where $\tilde{\beta} \equiv \beta\kappa/(1 + \kappa)$ denotes the fatality rate. Constancy of the fatality rate is an unreasonable feature over longer periods, due to the eventual slowdown of infections; we therefore view model 2 as a useful approximation only for the short run.

Since susceptible and infected persons are equally productive and the deceased do not contribute to production we have

$$V(1, x(t), y(t), z(t)) = \int_{j=0}^{\infty} e^{-\rho j} g(1, 1 - z(t), 0, z(t)) dj = (1 - z(t))\rho^{-1},$$

where we assume that productivity returns to normal levels once the vaccine is discovered. Regarding the (other) functional forms, we let $f(a(t)) = a(t)$, $\text{output}(t) = a(t)(1 - z(t))$, and

$$h(\dot{y}(t)) = \frac{h_2}{2} (\dot{y}(t))^2 \left(\frac{1 + \kappa}{\kappa} \right)^2 = \frac{h_2}{2} \left((1 - z(t)) \frac{\kappa}{1 + \kappa} \right)^2 \left(\frac{1 + \kappa}{\kappa} \right)^2 = \frac{h_2}{2} (\dot{z}(t))^2.$$

⁹This requires that $\bar{a} > 0$ as we assumed. If $a(t)$ fell to zero the economy would shut down and infections would no longer spread.

¹⁰When we introduce the policy choice $a(t)$ we disregard the fact that this would in principle also affect the condition $c^d = \beta$ and thereby modify the laws of motion.

That is, we let activity have a proportional effect on the infection rate and on per-capita output and we assume that the stress in the health care system is quadratic, with coefficient $h_2/2$.¹¹ Finally, we let $\mathcal{A} = [\bar{a}, 1]$. The government's program thus reads

$$\begin{aligned} \max_{(a(t))_{t=0}^{\infty}} \quad & \int_0^{\infty} e^{-(\rho+\nu)t} \left\{ a(t)(1-z(t)) - \frac{h_2}{2} (\dot{z}(t))^2 + \frac{\nu}{\rho}(1-z(t)) \right\} dt \\ \text{s.t.} \quad & \dot{z}(t) = a(t)\tilde{\beta}(1-z(t)), \quad a(t) \in [\bar{a}, 1], \quad z(0) \text{ given.} \end{aligned}$$

Optimal Policy The current value Hamiltonian now reads

$$H^c(t) = a(t)(1-z(t)) - \frac{h_2}{2} a(t)^2 \tilde{\beta}^2 (1-z(t))^2 + \frac{\nu}{\rho}(1-z(t)) + \mu(t)a(t)\tilde{\beta}(1-z(t))$$

and its derivative with respect to the control variable $a(t)$ is given by

$$\frac{dH^c(t)}{da(t)} = (1-z(t))(1+\tilde{\beta}\mu(t)) - h_2 a(t) \tilde{\beta}^2 (1-z(t))^2.$$

We conjecture that $a(t)$ is interior and thus satisfies

$$a(t) = \frac{1 + \tilde{\beta}\mu(t)}{h_2 \tilde{\beta}^2 (1-z(t))}. \quad (9)$$

The product $a(t)(1-z(t))$ then does not directly depend on $z(t)$ and the same holds true for all terms in the Hamiltonian except the third one.

The law of motion for the co-state variable is given by

$$\dot{\mu}(t) = -\frac{dH^c(t)}{dz(t)} + (\rho + \nu)\mu(t) = a(t)(1 + \tilde{\beta}\mu(t)) - h_2 a(t)^2 \tilde{\beta}^2 (1-z(t)) + \frac{\nu}{\rho} + (\rho + \nu)\mu(t).$$

With an interior choice of $a(t)$ the first two terms in this law of motion cancel. The resulting differential equation integrates to

$$\mu(t) = \left(\mu(0) + \frac{\nu}{\rho(\rho + \nu)} \right) e^{(\rho+\nu)t} - \frac{\nu}{\rho(\rho + \nu)}. \quad (10)$$

Recall that $\dot{z}(t) = a(t)\tilde{\beta}(1-z(t)) = (1 + \tilde{\beta}\mu(t))/(h_2 \tilde{\beta})$ where we use equation (9). From condition (10) we therefore have

$$z(t) = z(0) + \left(\frac{1}{h_2 \tilde{\beta}} - \frac{\nu}{h_2 \rho(\rho + \nu)} \right) t + \left(\mu(0) + \frac{\nu}{\rho(\rho + \nu)} \right) \frac{1}{h_2(\rho + \nu)} (e^{(\rho+\nu)t} - 1). \quad (11)$$

To find $\mu(0)$ we compute the value of the objective function under the optimal policy, W say, and differentiate it with respect to $z(0)$. Under our conjecture that the $a(t)$ path is interior, all terms in the objective that are proportional to $a(t)(1-z(t))$ do not directly

¹¹This implies that lockdown is motivated by the aim to flatten the curve.

depend on $z(0)$, so we can neglect them. Moreover, from condition (11) the integral over $e^{-(\rho+\nu)t}\nu(1-z(t))/\rho$ only depends on $z(0)$ through the term

$$\int_0^\infty e^{-(\rho+\nu)t} \frac{\nu}{\rho} (1-z(0)) dt = (1-z(0)) \frac{\nu}{\rho} \frac{1}{\rho+\nu}.$$

We conclude that

$$\mu(0) = \frac{dW}{dz(0)} = -\frac{\nu}{\rho} \frac{1}{\rho+\nu}$$

and thus, from condition (10), $\mu(t) = \mu(0)$.

This implies that the optimal path of $a(t)$ satisfies

$$a(t) = \frac{1 - \frac{\tilde{\beta}\nu}{\rho(\rho+\nu)}}{h_2 \tilde{\beta}^2 (1-z(t))}$$

provided that this solution lies in $[\bar{a}, 1]$. That is, during the short term (when the number of dead increases from $z(t) \approx 0$ to a small population share) the optimal size of the lockdown approximately equals

$$\left(1 - \frac{\tilde{\beta}\nu}{\rho(\rho+\nu)}\right) / (h_2 \tilde{\beta}^2).$$

Higher values for $\tilde{\beta}$ or h_2 , that is, a higher fatality rate or higher costs in the health care sector thus increase the optimal severity of the lockdown. A more likely discovery of a vaccine (higher ν) increases the stringency of the optimal containment measures because it renders $\mu(0) = \mu(t)$ more negative; this lowers $\dot{z}(t)$ and shortens the expected duration of the lockdown.

The implied solution for $z(t)$ is given by

$$z(t) = z(0) + \frac{1}{h_2} \frac{\rho(\rho+\nu) - \nu\tilde{\beta}}{\tilde{\beta}\rho(\rho+\nu)} t,$$

which is a valid solution only if $\rho(\rho+\nu) > \nu\tilde{\beta}$. Under this restriction $a(t)$ is indeed interior.¹²

Recall that in the absence of policy $z(t) = 1 - (1-z(0))e^{-\tilde{\beta}t}$ and thus $\dot{z}(t) \approx \tilde{\beta}e^{-\tilde{\beta}t}$. Comparing this expression with the time derivate of the preceding equality we conclude that the optimal policy reduces the number of new deaths at time t by

$$\tilde{\beta}e^{-\tilde{\beta}t} - \frac{1}{h_2} \frac{\rho(\rho+\nu) - \nu\tilde{\beta}}{\tilde{\beta}\rho(\rho+\nu)}.$$

¹²Plugging the expression for $z(t)$ into the expression for the optimal value of $a(t)$ derived above yields the duration until $a(t)$ reaches the activity level 1—the duration of the lockdown. Since we view the model as a model of the short run we do not emphasize this implied duration.

3.3 Taking Stock

Model 1 takes the exogenous lower bound \bar{a} as given and predicts the optimal duration of a lockdown. When the burden that infections impose on the health care system is sufficiently high then the optimal policy immediately imposes a lockdown and abandons it before everybody is infected. If the burden is low, in contrast, then the optimal policy never imposes a lockdown.

Model 2 predicts an interior path for the control variable when a parametric condition is satisfied. When the fatality rate or the cost of stress in the health care sector are higher, or discovery of a vaccine is more likely then the optimal lockdown is tighter. Over time the lockdown is slowly relaxed as the number of deaths decreases.

We view the predictions of the two models as complementary. In the next section we calibrate the two frameworks and generate quantitative predictions.

4 Calibration and Quantitative Results

We calibrate the model such that one period in the model corresponds to one day. Following Alvarez et al. (2020) we assume an annual discount rate of five percent, which translates into a daily rate of $\rho = -\ln(0.95)/365$.¹³ Also following Alvarez et al. (2020) we let $\nu = 0.0018$; this corresponds to a probability of roughly 28 percent that a vaccine is discovered during half a year,¹⁴ or to an expected time until discovery of about one and a half years.

For model 1, we assume that the fundamental infection rate, β , equals $(0.2)(0.8)$. Alvarez et al. (2020) assume that this value equals 0.2; we reduce it to correct for the simplified law of motion (see the discussion relating figures 1 and 2). Moreover, we assume that $\lambda = -\ln(0.5)/182$ such that the cost of stress in the health care sector (conditional on $\dot{y}(t)$) falls by one half after half a year.

To calibrate the parameter h_1 we rely on estimates according to which an unchecked Covid-19 infection wave would have caused costs in the U.S. of 13 trillion dollars, corresponding to roughly 61 percent of annual U.S. GDP (Scherbina, 2020).¹⁵ Consistent with Alvarez et al. (2020) and Scherbina (2020) we assume that this damage would have occurred within half a year. In light of the model this implies

$$h_1 \int_0^{182} e^{-(\rho+\nu)t} e^{-\lambda t} \dot{y}(t) dt = 0.61 \cdot 365,$$

where the right-hand side accounts for the fact that daily output equals one in normal times. Letting $y(0) = 0.01$ and evaluating the integral numerically we find that $h_1 \approx 265$.

For model 2, we calibrate the fatality rate, $\tilde{\beta}$, based on estimates according to which an unchecked Covid-19 infection wave would have caused 1.9 million deaths in the U.S.,

¹³ $0.95 = e^{-\rho 365}$.

¹⁴ $1 - 0.28 \approx e^{-\nu 182}$.

¹⁵ $13/21.4 \approx 0.6075$ (BEA data).

corresponding to roughly 0.58 percent of the U.S. population (Scherbina, 2020).¹⁶ Consistent with Alvarez et al. (2020) and Scherbina (2020) we assume that most of this death toll would have occurred within half a year. This yields an estimate of $\tilde{\beta} = -\ln(1 - 0.0058)/182$. To calibrate the parameter h_2 we again use the cost estimate of 61 percent of annual U.S. GDP (Scherbina, 2020). From this cost, we subtract the present value of the permanent output losses due to lost lives after the end of the transition. The remainder of the cost estimate is what we attribute to the cost due to health care stress. Formally, we solve

$$\int_0^{182} e^{-(\rho+\nu)t} \frac{h_2}{2} \tilde{\beta}^2 (1 - z(t))^2 dt = 0.61 \cdot 365 - e^{-\rho 182} z(182) \int_0^\infty e^{-\rho t} dt$$

or

$$\int_0^{182} e^{-(\rho+\nu)t} \frac{h_2}{2} \tilde{\beta}^2 (1 - z(0))^2 \left(e^{-\tilde{\beta} t} \right)^2 dt = 0.61 \cdot 365 - \frac{e^{-\rho 182}}{\rho} \left(1 - (1 - z(0)) e^{-\tilde{\beta} 182} \right).$$

Letting $z(0) \approx 0$ and solving for h_2 yields $h_2 \approx 2.34 \times 10^9$. With these values the parametric condition discussed in subsection 3.2 is satisfied.

Quantitative Predictions We use the formula from model 2 to compute the optimal activity level during lockdown, $a^*(0)$ say. We find that the optimal lockdown is severe: activity is reduced to roughly 33 percent of normal. Not surprisingly this depends on the health care cost parameter, h_2 . As figure 3 illustrates $a^*(0)$ increases substantially (the lockdown is less extreme) when the cost is lower (and we keep the other parameter values unchanged).




Figure 3: Predicted optimal activity level, $a^*(0)$, for different costs due to stress in the health care system.

Next, we set \bar{a} in model 1 equal to $a^*(0)$ and numerically solve for the optimal duration of the lockdown, T^* say. We find that this duration equals nearly 52 days. After the lockdown, roughly 13 percent of the population are infected. Figure 4 illustrates how the objective of the government varies with the duration of the lockdown. The cost of getting

¹⁶ $1.9/330 \approx 0.0058$ (Census data).

T wrong is asymmetric: Setting T a bit smaller than T^* is less costly than setting it a bit higher. If the lockdown is kept in place over a very long period (longer than roughly 90 days) then this policy generates a lower value than no lockdown at all.




Figure 4: Value of the program for different durations of the lockdown, T . The dashed line indicates the value when there is no lockdown, $T = 0$.

Figure 5 illustrates how the optimal policy changes the dynamics of infections. The solid line in the figure depicts the optimal path: it is relatively up to the optimal exit time and increases quickly thereafter. The dashed line depicts the path of $y(t)$ in the absence of policy.




Figure 5: Share of infected under the optimal policy (solid) and in the absence of policy intervention (dashed).

The first line in table 1 summarizes these baseline results. The other lines in the table report how the predictions change when we alter the calibration. When we assume a lower discount rate then $a^*(0)$ falls and T^* rises slightly: a more patient planner reduces activity by more, for longer. The same holds true when we assume a higher discovery rate for a vaccine. When we double ν to 0.0036 then $a^*(0)$ falls to roughly 28 percent and T^* rises to roughly 83 days.

Changes in β do not affect the optimal severity of a lockdown. However, they do affect T^* . A reduction in β by 20 percent increases T^* to nearly 61 days, and an increase in β by 20 percent lowers it to 45 days. Lowering $\tilde{\beta}$ by 20 percent increases $a^*(0)$ to 33.32 and reduces T^* to 51.32, while increasing it by 20 percent implies $a^*(0) \approx 32.62$ and $T^* \approx 51.98$. The model predictions thus are fairly robust to changes in the fatality rate.

Finally, when we increase λ by 10 percent such that after half a year, efficiency in the health care sector has increased by roughly 53 percent (for instance because preparations for the wave of infections have been particularly bad), then the optimal duration rises to more than 56 days. When we strongly reduce λ , however, the optimal policy eventually involves no lockdown at all. That is, when the health care system is adequately prepared to deal with a pandemic such that there is no role for learning or efficiency improvements over time then it is optimal not to impose a lockdown.

Table 1: Optimal Lockdown

Calibration	$a^*(0)$ (in percent)	T^* (in days)
Baseline	32.98	51.64
Annual discount rate 3%	31.50	51.83
ν twice as high	28.12	82.83
β 20% lower	32.98	60.87
β 20% higher	32.98	44.70
$\tilde{\beta}$ 20% lower	33.32	51.32
$\tilde{\beta}$ 20% higher	32.62	51.98
λ 10% higher	32.98	56.31

Table 2: Quantitative predictions under different calibration assumptions. a^* denotes the optimal activity level relative to normal and T^* the optimal duration of the lockdown.

5 Conclusion

We embed a lockdown choice in a simplified epidemiological model and derive formulas for the optimal lockdown intensity and duration. The optimal policy reflects the rate of time preference, epidemiological factors, the hazard rate of vaccine discovery, learning effects in the health care sector, and the severity of output losses due to a lockdown.

In our baseline specification a Covid-19 shock as currently experienced by the US optimally triggers a reduction in economic activity by two thirds, for about 50 days. On an annual basis, this corresponds to a drop in GDP by 9.5 percent.

We hope that future research can build on our simplified frameworks.

A Solving the SIR Model

The system (1)–(3) can be solved as follows (Bohner et al., 2019): Let $\xi(t) \equiv x(t)/y(t)$ for $y(t) \neq 0$. We have

$$\dot{\xi}(t) = \frac{\dot{x}(t)y(t) - x(t)\dot{y}(t)}{y^2(t)} = (c - b(t))\xi(t),$$

such that

$$\xi(t) = \xi(0)e^{\int_0^t (c-b(s))ds} \Leftrightarrow y(t) = x(t)\kappa e^{\int_0^t (b(s)-c)ds}$$

where $\kappa \equiv y(0)/x(0)$. Substituting into equation (1) yields

$$\dot{x}(t) = -b(t)x(t) \frac{\kappa e^{\int_0^t (b(s)-c)ds}}{1 + \kappa e^{\int_0^t (b(s)-c)ds}},$$

which has the solution

$$x(t) = x(0)e^{\int_0^t \frac{-\kappa b(u)}{\kappa + e^{\int_0^u (c-b(s))ds}} du}. \quad (4)$$

Accordingly, we can solve equation (2) for

$$y(t) = y(0)e^{\int_0^t \frac{b(u)}{1 + \kappa e^{\int_0^u (b(s)-c)ds}} - c du} \quad (5)$$

and equation (3) for

$$z(t) = 1 - x(t) \left(1 + \kappa e^{\int_0^t (b(s)-c)ds} \right) \text{ s.t. (4),} \quad (6)$$

where we use the fact that the population size equals unity.

References

- Alvarez, F., Argente, D. and Lippi, F. (2020). A simple planning problem for COVID-19 lockdown. Unpublished, University of Chicago, Pennsylvania State University, LUISS and Einaudi Institute.
- Atkeson, A. (2020). What will be the economic impact of COVID-19 in the US? rough estimates of disease scenarios, *Working Paper 26867*, NBER, Cambridge, Massachusetts.
- Baldwin, R. and Weder di Mauro, B. (eds) (2020a). *Economics in the Time of COVID-19*, CEPR Press, London.
- Baldwin, R. and Weder di Mauro, B. (eds) (2020b). *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes*, CEPR Press, London.
- Bohner, M., Streipert, S. and Torres, D. F. M. (2019). Exact solution to a dynamic SIR model, *Nonlinear Analysis: Hybrid Systems* **32**: 228–238.
- Eichenbaum, M. S., Rebelo, S. and Trabandt, M. (2020). The macroeconomics of epidemics, *Working Paper 26882*, NBER, Cambridge, Massachusetts.
- Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics, *Proceedings of the Royal Society, Series A* **115**(772): 700–721.
- Scherbina, A. (2020). Determining the optimal duration of the COVID-19 suppression policy: A cost-benefit analysis, *Economics Working Paper 2020-03*, American Enterprise Institute, Washington.
- Stock, J. H. (2020). Coronavirus data gaps and the policy response to the novel coronavirus. Unpublished, Harvard University.