



TECHNISCHE
UNIVERSITÄT
WIEN

Operations
Research and
Control Systems

SWM
ORCOS

How Long Should the COVID-19 Lockdown Continue?

*J. Caulkins, D. Grass, G. Feichtinger, R. Hartl, Peter M. Kort,
A. Prskawetz, A. Seidl, S. Wrzaczek*

Research Report 2020-10

June 2020

ISSN 2521-313X

Operations Research and Control Systems

Institute of Statistics and Mathematical Methods in Economics
Vienna University of Technology

Research Unit ORCOS
Wiedner Hauptstraße 8 / E105-04
1040 Vienna, Austria
E-mail: orcosp@tuwien.ac.at

How Long Should the COVID-19 Lockdown Continue?

Jonathan Caulkins^{1,*}, Dieter Grass², Gustav Feichtinger^{3,4}, Richard Hartl⁵, Peter M. Kort^{6,7}, Alexia Prskawetz^{4,8}, Andrea Seidl⁵, Stefan Wrzaczek⁴

1 Heinz College, Carnegie Mellon University, Pittsburgh, USA

2 International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria

3 Research Unit Operations Research and Control Systems (ORCOS), Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria

4 Wittgenstein Centre for Demography and Global Human Capital (Univ. Vienna, IIASA, VID/OAW)

5 Department of Business Decisions and Analytics, University of Vienna, Vienna, Austria

6 Tilburg School of Economics and Management, Tilburg University, Tilburg, Netherlands

7 Department of Economics, University of Antwerp, Antwerp, Belgium

8 Research Unit Economics (ECON), Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria

* caulkins@andrew.cmu.edu

Abstract

Nations struggled to decide when and how to end COVID-19 inspired lockdowns, with sharply divergent views between those arguing for a resumption of economic activity and those arguing for continuing the lockdown in some form. We examine the choice between continuing or ending a full lockdown within a simple optimal control model that encompasses both health and economic outcomes, and pays particular attention to when need for care exceeds hospital capacity. The model shows that very different strategies can perform similarly well and even both be optimal for the same relative valuation on work and life because of the presence of a so-called Skiba threshold. Qualitatively the alternate strategies correspond to trying essentially to eradicate the virus or merely to flatten the curve so fewer people urgently need healthcare when hospitals are already filled to capacity.

Introduction

The novel SARS-CoV-2 virus has literally swept around the globe. A prominent counter-measure has been to “lock down” non-essential parts of the economy to reduce contagious spread through social interaction. Lockdowns have succeeded in “flattening the curve” but at a high price; job losses in many places are the highest ever recorded, and there have been effects on stock markets not seen with prior epidemics (see [5]). That raises the question of when such measures should be relaxed. Too soon and the epidemic will bounce back; too late and there is needless economic hardship.

This paper analyzes a simple optimal control model of that difficult balancing. It supplements an SLIR epidemic model with controls that temporarily remove a proportion of the population from both social interaction and the workforce. The objective function

encompasses both health and economic considerations. Specifically, it recognizes benefits from people being free to work and costs when infected people need hospitalization, with an extra penalty when the number needing hospitalization exceeds hospital capacity.

The modeling of policy here is very simple, involving just picking the beginning and ending times of the lockdown, denoted by τ_1 and τ_2 , respectively. In reality there are various ways of partially easing a lockdown without ending it completely, so the policy choice modeled here is stylized.

Since most jurisdictions had already begun their lockdown, we initially planned to optimize only over τ_2 , for varying values of τ_1 . The results were surprising.

As might be expected, the optimal duration of the lockdown depends sharply on when it began, but not always in the expected way. Whereas one might have thought that starting too late would require keeping the lockdown in place longer, we find that starting later can sometimes make it optimal to end the lockdown sooner.

To understand the origins of this behavior, we formulate the model more generally to optimize over both τ_1 and τ_2 . Doing so reveals alternative optimal strategies that differ only modestly concerning timing of lockdown initiation, but which differ markedly in terms of health and economic outcomes. Later starts lead to health outcomes that are just enough worse to offset for the earlier re-opening. Remarkable is that both optimal solutions prevail under the same prioritization regarding health versus economic growth. Alternative optimal solution trajectories in dynamic models have variously been called Skiba, Sethi-Skiba, DNS, and DNSS points recognizing the contributions of various pioneers in the field [13]. The one here appears to be of a new type. Rather than emerging from interactions between the state and control variables, here path dependence is generated by an interaction between timing and duration.

The literature on COVID-19 models is exploding almost as fast as the virus itself, and some papers do model the balancing of health and economic interests (see [22], [7] and [8] for a careful evaluation). [14] depart from the SIR dynamics and have economic activity as a control variable. Increasing this variable on the one hand raises output, but on the other hand it feeds the number of infections, raising the burden on the health care system. Calibrating their model based on the U.S. situation, they suggest a lockdown of 50 days in which economic activity is reduced by two-thirds.

[3] also employ the SIR dynamics, but their control variable is the fraction of the population going into lockdown. They particularly include the effect of testing. Conditional on the possibility of testing and after parametrizing their model using data from the World Health Organization (WHO), they find that it is optimal to start a lockdown one week after the outbreak, and after one month it will be gradually withdrawn. Absence of testing increases the economic costs of the lockdown and shortens its optimal duration.

[2] investigate a heterogeneous SIR model distinguishing between the “young”, the “middle-aged” and the “old”. They find that it is especially beneficial to have stricter lockdown policies on the oldest group. [16] consider a stochastic version of the SIR model. They use a continuous-time Markov chain model to study the value and optimal timing of two (sequential) options: the option to intervene and, after intervention has started, the option to end it.

We, like they, recognize that locking down the population in order to reduce transmission also reduces employment. Our primary additional innovation on the modeling side is distinguishing between the health consequences of infection when proper care is received from those when the healthcare system’s capacity to render appropriate care has been exceeded. The number of deaths is not simply proportional to the number of infections; there is an extra penalty for infections that happen when hospitals are overwhelmed. Our primary methodological contribution is exploration of threshold effects when the optimal solution “tips” from one strategy to another.

Like us, [24] extend the SIR framework by including a state for individuals exposed to the virus but not symptomatic, and they also consider limited capacity of the health system. Their main result is that testing generates considerable welfare gains.

Where the just mentioned papers, including ours, formulate planner problems, [11] consider a competitive equilibrium in which a consumption tax is used to slowdown economic activity and epidemic diffusion. [20] extend [11] by distinguishing goods by the degree to which they can be consumed at home rather than in a social (and thus possibly contagious) context.

The model analyzed here is quite simple. It does not, for example, distinguish between young and old or consider movement between jurisdictions. So although we choose parameters that produce epidemic curves broadly consistent with those of other models, we stress the qualitative conclusions only.

The paper is organized as follows. The next section presents the model and elaborates the choice of the parameter values. The numerical results and their implications are discussed in the section 'Results'. The section 'Conclusions' concludes.

The Model

SLIR Model

The backbone of the model is an open-population SLIR model [17] with a birth rate ν and extra mortality for individuals who are infected (μ_I) above and beyond that for those who are susceptible or recovered (μ).¹ Given the state variables:

$S(t)$: Number of susceptible individuals at time t

$L(t)$: Number of latent (asymptomatic and pre-symptomatic) individuals at time t

$I(t)$: Number of infected individuals who are showing symptoms at time t

$R(t)$: Number of recovered individuals at time t ,

the SLIR state equations are:

$$\dot{S}(t) = \nu N(t) - \beta \frac{S(t)(I(t) + fL(t))}{N(t)} - \mu S(t) \quad (1a)$$

$$\dot{L}(t) = \beta \frac{S(t)(I(t) + fL(t))}{N(t)} - (\mu + \varphi)L(t) \quad (1b)$$

$$\dot{I}(t) = \omega \varphi L(t) - (\alpha + \mu + \mu_I)I(t) \quad (1c)$$

$$\dot{R}(t) = (1 - \omega)\varphi L(t) + \alpha I(t) - \mu R(t) \quad (1d)$$

$$\beta := R_{\text{eff}}(t, \tau_1, \tau_2)\alpha \quad (1e)$$

$$N(t) := S(t) + L(t) + I(t) + R(t).$$

Lockdown's effect on epidemic spread

The lockdown directly affects the rate of social interaction, β , but the effective reproduction number $R_{\text{eff}}(t, \tau_1, \tau_2)$ is more readily interpretable, so we describe the lockdown phases in terms of effects on $R_{\text{eff}}(t, \tau_1, \tau_2)$ and adjust β accordingly (for a formal derivation of the basic reproduction number R_0 we refer to appendix S2).

¹We introduce the birth and background death rate parameters for completeness. For the qualitative results these are of no importance since the time horizon for lockdowns is relatively short.

We distinguish three periods: before, during and after the lockdown, denoted by the subscripts 1, 2, and 3, respectively. The decision maker gets to choose the times, τ_1 and τ_2 , at which the lockdown is initiated and ended.

The epidemic dynamics are identical during those three periods except that the lockdown alters $R_{\text{eff}}(t, \tau_1, \tau_2)$ as follows.

Before the lockdown, $R_0^1 = 2.5$. During the lockdown, the reproductive rate if everyone were susceptible is $R_0^2 = 0.8$. After a sustained lockdown, that reproductive number would only bounce back to $R_0^3 = 2.0$, because some aspects of social distancing will be maintained indefinitely, or at least until a reliable vaccine is available, even though others will only be maintained during the lockdown. However, if the lockdown is short, then the post lock-down reproductive number would return to its original value of R_0^1 . In particular, the gap between realized and potential value for R_0^3 decays exponentially in the length of the lockdown.

COVID-19 deaths

For some people, COVID-19 is relatively mild, perhaps akin to a bad seasonal flu. We ignore costs associated with those cases, apart from recognizing that they cannot work while sick. Instead we focus on those who require hospitalization, particularly, those who require critical care.

If p denotes the proportion of infected people who need critical care, and ξ_1 denotes the probability of death for those needing and receiving critical care, then $\mu_I = p\xi_1\alpha$, where α is the reciprocal of the average duration of symptoms, which we take to be nine days.

That level of modeling precision is adequate for the state dynamics, because deaths from COVID-19 are not common enough to appreciably alter population size, but within the objective function greater precision is needed.

A central challenge of COVID-19 is the surges in demand for care, so we distinguish between those who need critical care and receive it from those who need it when hospitals are full and so cannot be treated properly. That means health harms are driven not only by the number of people who get infected, but also by how peaked the epidemic is; flattening the curve lets more people receive appropriate care.

In particular, we distinguish two components to the flow of deaths from COVID-19. In addition to deaths that are proportional to $(pI(t))$ there is an extra penalty term proportional to $\max(\{0, pI - H_{\text{max}}\})$ where H_{max} is the number of critical care hospital beds available. The first term captures deaths that would occur even if there were no constraints on hospital capacity; the second captures the incremental risk of death if one needs critical care but does not receive it.

The max function is not differentiable. Furthermore, some costs arise when pI is smaller than but close to H_{max} . For example, medical personnel (doctors, nurses) cannot give the usual care to individual patients if the intensive care unit is almost fully occupied. So we choose the following smooth function $\max(\cdot, \zeta)$ that is increasing and approximately linear in $pI(t) - H_{\text{max}}$:

$$\max(\{0, pI - H_{\text{max}}\}, \zeta) := \frac{1}{\zeta} \log \left(1 + e^{\zeta(pI(t) - H_{\text{max}})} \right), \quad \zeta \gg 1.$$

Fig 1 shows this is an extremely close approximation when ζ is large.

Hence, deaths from COVID-19 in the objective function can be written as

$$\xi_1 pI + \xi_2 \max(\{0, pI - H_{\text{max}}\}, \zeta),$$

where ξ_1 is the death rate from COVID-19 of infected people who need and receive critical care, and ξ_2 is the additional, incremental death rate when such individuals do not receive that care.

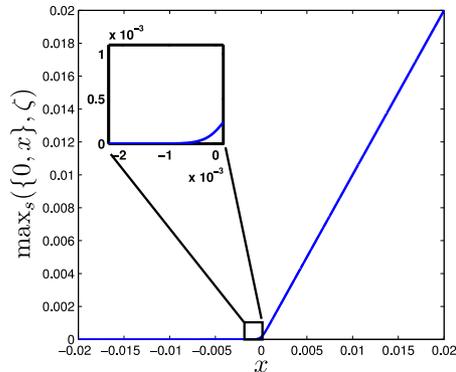


Fig 1. Approximation of the max function for $\zeta = 5000$.

Objective function

Deaths from COVID-19 are quick compared to those from chronic diseases or even cancer. The average time spent in hospital is about 12 days for those admitted to critical care and one-third that for others [9]. Since fatality rates are about 15% for all who are hospitalized and perhaps 45% for those receiving critical care, for any reasonable valuation of the social cost of a premature death, the cost of deaths is much greater than the cost of the hospitalization per se. Hence, health costs are represented by M , the cost per COVID-19 death, times the number of deaths.

There are instances of long-lasting morbidity associated with severe cases of COVID-19, with persistent damage to lungs or kidneys, for example. Should those turn out to be large compared to the costs of death, they could be accommodated through this same expression, just by using a larger value of M , if they are also driven by infections and unmet need for care.

Note this makes the health part of the objective function effectively piecewise linear in I . It is common to make cost be convex in an outcome, e.g., to use a quadratic function. However, that would imply the marginal cost of the 100th person who cannot receive needed care is greater than that of the 50th person. Making costs be (effectively) linear in the amount of unmet need values all people equally.

Economic activity is modeled as being proportional to the number of employed people raised to a power, as in a classic Cobb-Douglas model, with that exponent set to $\sigma = 2/3$ [1]. Since the model's time horizon is so short, capital is presumed to be fixed and so is subsumed into the objective function coefficient K for economic activity.

Susceptible, latent, and recovered individuals are eligible to work (symptomatic individuals are assumed to be too sick to work, or are in quarantine), but because of the lockdown the actual number working is only a proportion $\gamma(t)$ of those eligible to work $W(t)$. Before the lockdown $\gamma(t) = 1.0$, during it is reduced to 0.25, and afterwards it bounces back but only partially. The longer the lockdown, the more jobs are lost semi-permanently because firms go out of business. That recovery is modeled as decaying exponentially in the length of the lockdown with a time constant of 0.001 per day, so that if a lockdown ended after six months, 17% of jobs suspended during the lockdown would not reappear, at least until a vaccine became available.

It is presumed that after a vaccine has been widely deployed, there will not again be many deaths from COVID-19, but the economy will not necessarily snap back to full employment instantly, so the objective function includes a salvage function to capture that. In particular, the salvage value function reflects the reduction in economic activity at time T , when a vaccine is deployed, relative to what it was at time 0 and so would have been but for COVID-19 and the lockdown. That is a rate or flow and needs to be

multiplied by a factor Γ reflecting how long it takes for the economy to recover from that underemployment. No one knows the duration or shape of that recovery; which has variously been discussed as being V-shaped, U-shaped or W-shaped. For simplicity we set $\Gamma = 365$, as it would be for example if the recovery were linear but took two years, so the area is triangle-shaped.

Note. Omitting deaths after time T is a slight simplification, because some people who have an active infection at time T might still die after time T . However, because of the speed of the infection relative to the time horizon, it will be shown that the number of people infected at time T is not so large, so this is not a major concern.

Full model

The decision variables are τ_1 and τ_2 , the times when the lockdown begins and ends, and the full model can be written as:

$$V(X_0, \tau_1, \tau_2) := \tag{2a}$$

$$\int_0^T \left(M (\xi_1 p I(t) + \xi_2 \max(\{0, p I(t) - H_{\max}\}, \zeta)) - K \gamma(t, \tau_1, \tau_2)^\sigma W(t)^\sigma \right) dt + \tag{2b}$$

$$(T + \Gamma) K W(0)^\sigma \gamma(0, \tau_1, \tau_2)^\sigma - \Gamma K W(T)^\sigma \gamma(T, \tau_1, \tau_2)^\sigma \tag{2c}$$

$$V^*(X_0) := \min_{\tau_1, \tau_2} V(X_0, \tau_1, \tau_2), \quad X := (S, L, I, R), \quad W := S + L + R. \tag{2d}$$

$$\text{s.t. } \dot{X}(t) = \begin{cases} \text{SLIR}_1(X(t), \tau_1, \tau_2) & 0 \leq t < \tau_1 \\ \text{SLIR}_2(X(t), \tau_1, \tau_2), & \tau_1 \leq t \leq \tau_2 \\ \text{SLIR}_3(X(t), \tau_1, \tau_2) & \tau_2 < t \leq T \end{cases} \tag{2e}$$

$$X(0) = X_0 \geq 0 \tag{2f}$$

$$\gamma(t, \tau_1, \tau_2) := \begin{cases} \gamma_1 & 0 \leq t < \tau_1 \\ \gamma_2 & \tau_1 \leq t \leq \tau_2 \\ \gamma_3(\tau_1, \tau_2) := \gamma_2 + (\gamma_1 - \gamma_2) e^{\kappa_2(\tau_1 - \tau_2)} & \tau_2 < t \leq T \end{cases} \tag{2g}$$

$$R_{\text{eff}}(t, \tau_1, \tau_2) := \begin{cases} R_0^1 & 0 \leq t < \tau_1 \\ R_0^2 & \tau_1 \leq t \leq \tau_2 \\ R_0^3(\tau_1, \tau_2) := \bar{R}_0^3 + (R_0^1 - \bar{R}_0^3) e^{\kappa_1(\tau_1 - \tau_2)} & \tau_2 < t \leq T \end{cases} \tag{2h}$$

$$\text{with } R_0^2 \leq \bar{R}_0^3 \leq R_0^1. \tag{2i}$$

To refer to the health care term, the economic (labor) term, and the salvage functions the objective value (2c) we shortly write

$$V_h(X_0, \tau_1, \tau_2, M) := M \int_0^T (\xi_1 p I + \xi_2 \max(\{0, p I - H_{\max}\}, \zeta)) dt \tag{2j}$$

$$V_l(X_0, \tau_1, \tau_2, K) := T K W(0)^\sigma \gamma(0, \tau_1, \tau_2)^\sigma - K \int_0^T \gamma(t, \tau_1, \tau_2)^\sigma W(t)^\sigma dt \tag{2k}$$

$$V_s(T, \tau_1, \tau_2, K) := \Gamma K W(0)^\sigma \gamma(0, \tau_1, \tau_2)^\sigma - \Gamma K W(T)^\sigma \gamma(T, \tau_1, \tau_2). \tag{2l}$$

The derivation of the necessary optimality conditions can be found in the Appendix S1. The Matlab toolbox OCMAT is used for the numerical calculations. ²

²See http://orcos.tuwien.ac.at/research/ocmat_software.

Parameterization

The initial population is normalized to 1.0. Optimization begins with $L(0) = 0.1\%$ of the population being latent and 99.9% being susceptible and continues over a finite time horizon of $T = 365$ days, representing the time until a vaccine is hoped to be widely available.

The qualitative results are similar with $T = 730$ days.

Table 1 summarizes the base case parameters, several of which have already been discussed. The others are addressed here.

In July 2020, the U.S. Centers for Disease Control (2020) released new guidance for parameters in COVID-19 planning models. They suggest assuming that $\omega = 60\%$ of infections become symptomatic, that infectiousness in the L state is $f = 75\%$ of that in the I state. They suggest that presymptomatic individuals spend 6 days in the L state, so we make $\alpha = 1/9$, corresponding to a dwell time of 9 days in the I state, so symptomatic cases are infectious for a total of $6 + 9 = 15$ days.

The CDC advises that 50% of infections arise from people not manifesting symptoms (i.e., from the L state), and that implies that those who are never symptomatic remain infectious for 9 days. (Since that makes the average number of infections per person passing through the L state, or $(\omega \cdot 6 + (1 - \omega) \cdot 9) \times 75\%$ match the corresponding figure for infections coming from the I state.) That in turn implies that the average dwell time in the L state is $\omega \cdot 6 + (1 - \omega) \cdot 9 = 7.2$ days, so $\varphi = 1/7.2$.

The CDC estimates of the probabilities of death given hospitalization vary by age, but the weighted average, weighting by the age distribution of hospital cases, is roughly 15%. Dividing that by their estimate that 32% of COVID-19 hospital admissions need intensive care (ICU) suggests that the probability of death for those needing and receiving intensive care is about 45%. The parameter ξ_1 is the death rate per day for such people. Since the average dwell time in the I state is 9 days, the death rate per day is $\xi_1 = 5\%$.

ξ_2 is the additional, incremental death rate per day for infected people who need critical care but do not receive it. If the death rate for such individuals over an entire infection is 100% and the average dwell time in the I state is 9 days, then the incremental death rate per day is $\xi_2 = \alpha(1 - 45\%)$, or about 6.11%.

ICU beds are a smaller share of all hospital beds than is the proportion of COVID-19 hospital admissions needing ICU care, even after adjusting for different lengths of stay, so the care constraint will be on critical care beds.

[27] suggest that in the U.S., 58,166 of the existing 84,750 ICU beds could be made available for treating COVID-19 patients. Given the U.S. population is about 330 million, that is 0.176 per 1,000 people. The model acts as if patients who need critical care at some point need that care throughout their 9-day dwell time in the I state. CDC data suggest that the average time in hospital for those needing critical care at some point is about 12 days, but not all of that time is necessarily in the ICU, so the 9 days is probably about right. So we set $H_{\max} = 176$ per 1,000 or 0.000176. There are roughly ten times as many hospital beds as critical care beds, so sensitivity analysis with larger values of H_{\max} may be of interest if some regular beds could be converted over to critical care beds, although some readers have been skeptical about that possibility.

The probability of needing ICU care for those entering the I state, p , is the proportion of hospital admissions needing ICU care (32%) times the proportion of those showing symptoms who need hospitalization. The latter quantity can be computed indirectly because the infection fatality rate (which CDC gives as 0.65%) is the product of that quantity, the probability of death given hospitalization (15%) and the proportion of cases that are symptomatic ($\omega = 60\%$). That implies that the probability of hospitalization given symptoms is $0.65\% / (15\% \cdot 60\%) = 7.22\%$, and so $p = 7.22\% \cdot 32\% = 2.31\%$.

Dividing the infection fatality rate (0.65%) by $\omega = 60\%$ implies that the probability of death given symptoms is 1.0833%. Dividing that by the average dwell time in the I

state implies that $\mu_I = 0.0012037$.

There is not truly consensus in the literature about any of the key parameters, but the two for which the widest range of values seem plausible are the proportion of symptoms that become symptomatic, ω , and the social cost of a death, M , so we discuss them at length.

The CDC’s estimate that 60% of infections produce enough symptoms to be detected reflects the original understanding of the epidemic, but there have recently though been community-wide antibody tests in several European countries, in New York State, Santa Clara and Los Angeles Counties in California, and in a number of prisons, all suggesting that there may be far more undetected infections than was previously thought. For example, the Santa Clara study concluded that the actual number infected could be 50 to 85 times more than the number of confirmed cases [6].

That would suggest a substantially lower level of ω may be appropriate. [10], considering data from multiple European countries, suggest that infections could be ten times more common than previously supposed. That may be overstating, but we consider values down to $\omega = 12\%$ in other runs of the model, and the results below include sensitivity analysis with respect to this parameter.

Estimating the relative value of lost work vs. lost lives is tricky, to say the least, so without loss of generality we set $K = 1$ and consider a very wide range of values for M .

Still, it is helpful to determine at least roughly the size of M relative to K . [3] value a premature death at $20\times$ GDP per capita, while noting that [18] use a much greater value of $150\times$ GDP per capita. [15] offers multiple reasons why lower values may be preferred for analysis of COVID-19 in particular. For example, lower values would apply if one focused on years-of-life-lost, since most deaths are among the elderly, especially those with other pre-existing conditions. E.g., [25] report that the vast majority of those hospitalized for COVID-19 had prior serious comorbidities such as hypertension, obesity, and diabetes, to the extent that their estimated 10-year survival rate absent COVID-19 was only 53%. So we consider a range from $10\times$ to $150\times$ GDP per capita.

$K(\gamma W)^{2/3}$ measures GDP per day - K is the constant that we assume to capture everything except labor, so $365K(\gamma W)^{2/3}$ equals the nation’s GDP. Since the population size is normalized to 1.0, that implies values of a premature death, M , somewhere in the range from 3,650 to 54,750. We set $M = 21,716$ for analyses with a fixed M .

Table 1. Base case parameter values.

α	$R_0^{1,2,3}$	H_{\max}	p	M	K	$\gamma_{1,2,3}$	Γ	φ	f
$\frac{1}{9}$	2.5, 0.8, 2.0	1.76×10^{-4}	2.311×10^{-2}	21716	1	1, 0.25, 0.75	365	$\frac{1}{7.2}$	0.75
ω	κ_1	κ_2	σ	μ	ν	μ_I	ζ	ξ_1	ξ_2
0.6	2×10^{-3}	10^{-3}	$\frac{2}{3}$	$\frac{0.01}{365}$	$\frac{0.01}{365}$	$\frac{13}{10800}$	5000	0.05	$\frac{0.55}{9}$

Results

Optimizing only over the end time, τ_2

Countries have already started their lockdowns, so we start by optimizing only over τ_2 , but for various values of τ_1 because different places started their lockdowns at different times. Fig 2 plots the resulting optimized value of τ_2 (blue curve) vs. τ_1 . For convenience, the value of τ_1 is also shown, by the black line, so it is easy to visualize the duration of the lockdown as the height of the gap between the blue and black lines.

As τ_1 increases from 0 to a little more than 25, τ_2 increases by about the same amount, keeping the duration of the lockdown roughly constant. That makes intuitive sense. The later one begins the lockdown, the later it should be relaxed.

However, as τ_1 increases further, from about 25 to 44, the optimal τ_2 decreases. Hence, over that range, countries that started their lockdown later should end it sooner, something that is perhaps surprising. It may be that the timing of the ending of the lockdown is related to the accumulation of infections, which happens rapidly when the lockdown is not initiated early.

The solid red vertical line indicates the value of τ_1 that produces the lowest total cost, when τ_2 is optimized, and that is for $\tau_1 = 3.92$.

That this optimal value of τ_1 is greater than zero is perhaps a second surprising result. If lockdowns are costly, in this case because they reduce work, then it is possible to start a lockdown too soon, and for a later starting date to be preferred. Since the epidemic grows explosively, even exponentially, before the lockdown, some might have thought that earlier was always better.

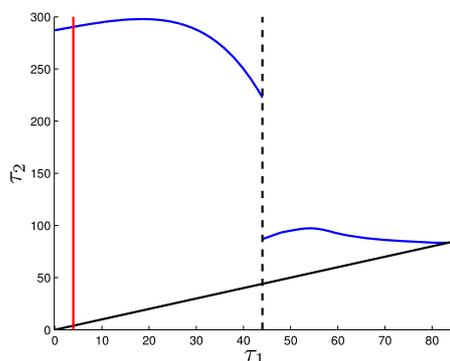


Fig 2. Fixed initial lockdown time τ_1 and optimally chosen time τ_2 . For $\tau_1 = 44$ there exists a Skiba solution, i.e. there are two different solution paths which deliver the same objective value.

Perhaps not surprisingly, if the initial reproductive rate R_0^1 is smaller then this optimal starting time comes later. This may explain some of the tension observed between residents of rural areas and urban areas over lockdown timing. The optimal date for starting the lockdown in a place like New York City, with high population density (so likely higher R_0^1) and early onset could be sooner than the optimal date for starting the lockdown elsewhere in the country.

When τ_1 equals the critical value of 44, indicated by the vertical dashed line, there are two alternative optimal solutions for τ_2 . Both $\tau_2 = 222.5$ (lockdown for over 7 months) and $\tau_2 = 86.9$ (a lockdown of less than three months) produce the same objective function value. And as τ_1 increases beyond 44, the best end time for the lockdown decreases further. Eventually, if the lockdown doesn't start until day 84, it is optimal not to start a lockdown at all.

The interpretation for this is that there are two broad strategies that can be pursued. One, might be called an 'eradication' strategy. It locks down long enough to push infections down to minimal levels, with just a modest rebound shortly before the end of the time horizon when a vaccine becomes available. The second, which might be called 'curve flattening', uses the lockdown to reduce the size of the initial spike in infections, so fewer come when infections exceed hospital capacity. The later the lockdown starts, the harder it is to pull off the eradication strategy, until at some critical point (in this case $\tau = 44$) it becomes optimal to switch to the 'curve flattening' strategy which requires a much shorter lockdown.

Optimizing over both τ_1 and τ_2 makes these results easier to understand. Fig 3 shows the optimal values of τ_1 and τ_2 (left panel) and the optimal solution value V^* (right panel) as a function of M , the cost of a COVID-19 death. It shows three regions. In region I, on the far left when M is very small, it is optimal to never lockdown and just let the epidemic run its course. Basically, if one does not care much about deaths, beyond their effect on economic productivity from reducing work, then the lockdown is not worth it.

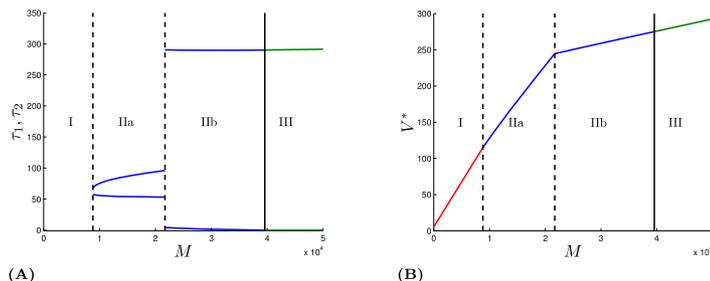


Fig 3. Figure showing the different regimes for varying M . The other parameter values are taken from Table 1. Mathematical description of the three regimes: Regime I: no lockdown, i.e. $\dot{X}(t) = \text{SLIR}_1(t), 0 \leq t \leq 365$, Regime II: lockdown in interval $(0, 365)$, i.e. $0 < \tau_1 < \tau_2 < 365$, Regime III, lockdown starts immediately, i.e. $0 = \tau_1 < \tau_2 < 365$.

In Region III, corresponding to large values of M , one should begin the lockdown immediately and keep it in place until shortly before the vaccine becomes available.

In Region II both broad strategies (flattening and eradication) can be considered, with flattening being preferred in Region IIa, eradication being preferred in Region IIb, and either being optimal when M is just exactly at the value separating Regions IIa and IIb. At that $M = 21,716$, there is a so-called Skiba point.

The existence of the Skiba point implies another surprising result. Both short and long lockdowns can be optimal for the exact same set of parameter values. Exact equality only occurs at that specific value of M , but for a range of values in that neighborhood, the two very different strategies perform nearly as well. So a single person can be indifferent, or nearly indifferent, between two very different approaches.

Furthermore, two people with modestly different relative valuations on work and health can favor very different policies, if their modestly different values of M lie on either side of the Skiba threshold.

As mentioned, plausible values of M range from 3,650 to 54,750. The lower end of that range lies within Region I, but increasing M within that range would carry the solution through Regions IIa and IIb, all the way into Region III. As mentioned earlier, we focus on qualitative results, not specific values. So we do not think the model shows that one strategy or another is necessarily the best. Rather, we would say that the model suggests that it is hard to be certain about that. Any one of the strategies could plausibly turn out to be optimal, depending on how uncertainty about the parameter values is resolved.

Fig 4 contrasts the two solutions when M is exactly at its Skiba threshold value of 21,716. The vertical lines indicate the start (τ_1) and end (τ_2) times of the lockdown for the curve flattening strategy (dashed lines; left panel) and the eradication strategy (dashed lines; right panel). The blue lines show the numbers who are infected under those two strategies, while the gray line shows the uncontrolled epidemic with no lockdown.

The number in the latent state tracks closely the number in the infected state. In particular, the number who are infected and symptomatic is a slightly lagged and smoothed parallel to the number in the latent state. The number of infected is lower than the number in the latent state in part because of the higher death rate for those

who are infected, but primarily because the latent state includes asymptomatic as well as presymptomatic individuals.

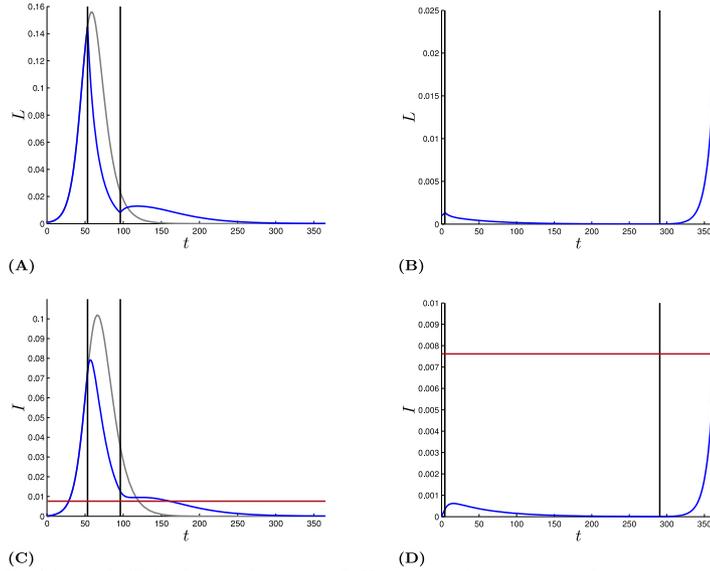


Fig 4. Panel (A) and (B) show the two different solutions paths emerging from the Skiba point for the number of latent $L(\cdot)$ (blue lines) and panel (C) and (D) show those paths for the number of infected $I(\cdot)$. The gray line in panels (A) and (C) depicts the uncontrolled path without lockdown for comparison. For the total costs and its shares see Table 2. The horizontal red line shows the health care capacity H_{\max}/p .

Note the very different vertical scales in the two panels. The eradication strategy keeps the number of infections well below the point where hospital capacity becomes binding, whereas the curve flattening strategy (greatly) exceeds that capacity, at times.

Here the optimal curve flattening strategy experiences a second, much smaller bump, but for other parameter values it can be a second spike. That is, it can be optimal to end the lockdown in a way that creates a resurgence of the epidemic. That the epidemic spikes again after ending a lockdown does not imply that ending the lockdown was a mistake.

Table 2 contrasts the valuations of the three components of the objective function under the two strategies, and the uncontrolled epidemic as a foil. The values can perhaps best be understood in percentage terms. In the absence of an epidemic, economic output would have been 365, and in the absence of controls the health cost would be 270. So under the shorter lockdown, economic output falls by about 9% for the year, and the costs of premature deaths are reduced by about one-quarter. With the longer lockdown, 86% of the health costs are averted, but economic output falls by nearly 50%.

Sensitivity analysis with respect to M

Fig 5 is a sensitivity analysis with respect to the value of M , i.e., the relative valuation placed on health as opposed to economic outcomes.

The upper left panel shows which type of solution is optimal for various values of the two most uncertain parameters, the proportion of infections that become symptomatic (ω) and the value of averting a death, expressed as a multiple of the nation's daily GDP per capita (M). Even holding ω at its base case value of 60 % and varying M

Table 2. The total costs, the share of costs of the health care system, profit from labor and the salvage value for the Skiba solution and the uncontrolled solution of Fig 4 ($M = 21, 716, K = 1$) and the lockdown times.

	Flattening	Eradication	Uncontrolled
$V^*(X_0)$	244.6	244.6	275.2
$V_h(X_0, \tau_1, \tau_2, M)$	206.8	38.0	269.9
$V_l(X_0, \tau_1, \tau_2, 1)$	31.9	177.9	4.1
$V_s(T, \tau_1, \tau_2, K)$	5.9	28.7	1.2
$[\tau_1, \tau_2]$	[53.3, 96.2]	[4.4, 290.5]	

between 0.365 and 5.475×10^4 spans all four regions. Adding in the variation in ω down to its lower plausible value of 12 % just accentuates the uncertainty as to which region pertains. That indicates that people can disagree about the best broad strategy depending on their judgments about how to balance health and economic outcomes and on the degree to which the new antibody testing results alter judgments about how many infections remain undetected. This shows that two reasonable, intelligent people who agree completely on how to think about the epidemiology, health, and economics of the COVID-19 crisis can still reach very different conclusions about what policy is best, just because of differences concerning these two key parameters, about which there is no scientific agreement, at least at present.

That upper left panel is drawn for $H_{\max} = 0.000176$. The upper right panel of Fig 5 is a bifurcation diagram showing which strategies are optimal as a function of both M and H_{\max} , which represents the number of critical care beds, while holding p at the value of 0.0231. That diagram shows that the qualitative results would not change until H_{\max} increased to a very large multiple of its base case value of 0.000176. For the most part, the curves in that panel slope up and to the right, meaning that if treatment capacity and the valuation of a life both increase, the structure of the solution remains the same. Eventually those curves turn vertical when hospital capacity is no longer binding.

The exception is that for M small we have, for increasing H_{\max} , first Region I then IIa and then I again. This curious pattern can be explained as follows. Because M is small, lockdowns will be short, if they exist at all, and so most people will become infected. If H_{\max} is very small then almost all of those infections will occur when hospital capacity has been exceeded. Conversely, if H_{\max} is very large, then almost all infections will occur when there is adequate hospital capacity, even if there is no lockdown. It is only for intermediate values of H_{\max} that the lockdown has the ability to tip a meaningful share of the infections from times when hospital capacity is exceeded to times when adequate care can be provided.

The lower left panel shows how the optimal solutions depend on both R_0^2 and M when p is fixed at the value 0.0231. Not surprisingly, lockdowns are preferred when M is large and/or R_0^2 is low, which explains why the boundaries are generally upward sloping.

Within the range of R_0^2 of greatest interest, the boundaries between no lockdown and short lockdown, and between short and long lockdowns are almost vertical, indicating that the precise value of R_0^2 is less important than the value of M for driving what strategy is optimal.

An exception is the curve between regions IIb and III. That is because for very low values of R_0^2 , the lockdown is so effective that it is not necessary to start the lockdown immediately to fight the epidemic.

For $R_0^2 = 1.1724$ and $M = 33, 550$ there is a so-called triple Skiba point, where three qualitatively different strategies perform equally well. The lower right panel depicts the time trajectories for $I(t)$ under those three strategies which are: (1) Locking down immediately (dashed line), (2) Locking down around day 24 and holding that nearly

until the vaccine arrives (dot-dash lines), and (3) merely flattening the curve with a relatively short lockdown from $\tau_1 = 48$ to $\tau_2 = 118$. Infections increase initially after lockdown under the two eradication strategies because R_0^2 is greater than one at that triple Skiba point.

One additional sensitivity analysis (not shown) considered the scenario when treatment is much more effective, so that the probability of dying if one needs and receives critical care is only 4.5% not 45%. That makes the extra penalty for exceeding hospital capacity gets much worse, 95.5% not 55%, but this does not change much about the structure of the solution.

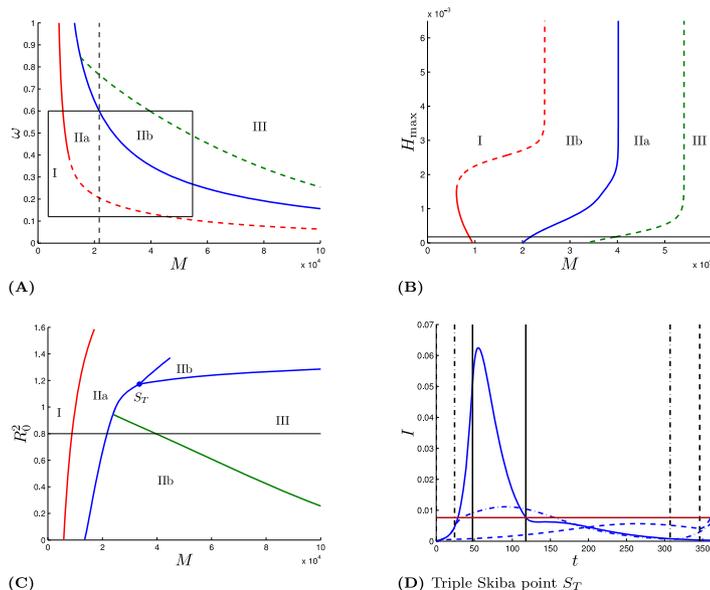


Fig 5. Bifurcation diagrams. Dashed curves correspond to a continuous regime transition and solid lines to a disruptive regime transition. Panel (A) shows the bifurcation diagram for parameter values of Table 1 while varying ω and M . Panel (B) shows the bifurcation diagram for parameter values of Table 1 and varying H_{\max} and M . Panel (C) shows the bifurcation diagram for M and R_0^2 . The black horizontal lines depict base case parameter values. The solid vertical lines in panel (B) delimit the region $M \in [0.365 \times 10^3, 5.475 \times 10^4]$. The dashed vertical line shows the M value of the base case (Skiba case). In panel (D) a triple Skiba solution is depicted (resulting for the (M, R_0^2) -parameter combination S_T of panel (C)). In that case two “eradication” solutions and one “curve flattening” solutions exist. The dashed solution starts with quarantine immediately.

Conclusions

Primary findings

This paper is about planning the timing of complete lockdowns in COVID-19 times. The aim of installing a lockdown is to limit social contact to reduce the number of people getting infected. The drawback is the concurrent reduction in economic activity. The latter also has an effect after the lockdown has ended, because it takes time for the economy to recover.

Our model is based on the epidemiological SLIR dynamics, which we combine with a basic Cobb-Douglas model of economic activity to develop a framework that can evaluate

the above-described tradeoff. In many countries the number of COVID-19 patients needing intensive care treatment came close to or even exceeded the available intensive care capacity, so that aspect is included in our model.

We find that essentially two different solution patterns can be optimal. One is what we call an *eradication strategy*, where a long lockdown significantly reduces not only the bad health effects of the epidemic but also economic activity. The other is a *curve flattening strategy* characterized by a relatively short lockdown period. The idea is to reduce the peak of the number of infected in order to limit the violation of the intensive care capacity constraint, where at the same time economic activity is not harmed as much.

Interesting is that the performance of both strategies can be quite similar. Indeed, there are specific parameter values such that two decision makers with the same preferences for health versus the economy can opt for completely different lockdown policies and both be choosing optimally.

We note that over the plausible ranges for two key parameters, namely the proportion of infections that become symptomatic and the valuation placed on preventing a COVID-19 related death, all types of solutions can be optimal. We view that as indicating that a degree of humility and open-mindedness may be appropriate; at least in our model, it is not necessarily clear whether a longer or a shorter lockdown is best.

Limitations

This model simplifies in many respects that could bear on the optimal timing and duration of a lockdown. That is why we stress qualitative results, not numerical results. These limitations include:

- Transmission mitigation is modeled crudely as a lockdown being in place or not. The intensity of the lockdown could instead be modeled as a continuous control variable because there are a wide range of non-pharmaceutical interventions that could be deployed in various combinations, not just all-or-nothing (see [12] and [23]). That would permit easing the lockdown gradually. In the present model the lockdown ends abruptly, although not completely; we model the reproductive rate as bouncing back to 2.0 but not its pre-lockdown level of 2.5 because we presume some protective measures are retained even after people go back to work. Also, the most important innovations might be efforts to re-engineer operations to make them safer, rather than either shutting them down or allowing them to operate normally.
- This model does not divide the population by gender, age or pre-existing health condition. Death rates are sharply higher for those who are older or who have pre-existing conditions. Strategies that have longer or more restrictive lockdowns for vulnerable populations may be prudent but cannot be modeled here. So the findings are directly relevant when policy is constrained not to impose different lockdowns on different socioeconomic groups.
- The technology for treating COVID-19 cases may improve over time. E.g., Remdesivir and convalescent plasma treatments are now being used. Improving technology could favor earlier and longer lockdowns if that defers cases until technology is better.
- The model does not consider seasonality. That may or may not be a serious limitation. The virus clearly rebounded sharply in many places during the Northern Hemisphere's fall, but our understanding is that there is still some debate as to whether that is seasonality per se, or the lagged effect of relaxation of policies

and ensuing vacationing. Still, there are variations over time in rates of social interaction, e.g., because of holidays and school calendars, and those are not considered at present.

- The model does not consider the possibility of long-term health effects of survivors of COVID-19, although there are reports of lasting harm not only to lungs, but also to the brain and kidney.
- The model ignores geography. In reality there is not one social planner governing one region, but multiple regions that affect each others' economies and epidemics. The optimal policy for one region may depend on what is happening in other regions to which it is connected by travel and migration, in terms of infection rates, policy choices, and economic dislocation.

Further work

A variety of extensions of this model are possible, including making the timing of the vaccine's arrival unknown, allowing for multiple lockdowns separated by periods of relaxation, and altering the epidemic's dynamics when the number of new infections is small enough that contact tracing can lead to quickly quarantining everyone that an incident case might have infected (see [19] and [21]).

The epidemic model could be enriched in various ways including making the rates of social interaction a distributed parameter that varies across a heterogeneous population, and modeling explicitly the population's patience with and commitment to lockdown restrictions.

Supporting information

S1 Appendix: Necessary Optimality Conditions Setting $\tau_0 := 0$ and $\tau_3 := T$ the Hamiltonians for the three stages $i = 1, 2, 3$ are

$$\mathcal{H}_i(X_i, \Lambda_i, \tau_1, \tau_2) := M(\xi_1 pI + \xi_2 \max(\{0, pI - H_{\max}\}, \zeta) - K\gamma(t, \tau_1, \tau_2)^\sigma W(t)^\sigma) + \Lambda_i' \text{SLIR}_i(X_i, \tau_1, \tau_2), \quad (3)$$

with

$$\begin{aligned} X_i(t) &:= X(t), \quad \tau_{i-1} \leq t \leq \tau_i \\ \Lambda_i(t) &:= (\lambda_1(t), \lambda_2(t), \lambda_3(t), \lambda_4(t)), \quad \tau_{i-1} \leq t \leq \tau_i. \end{aligned}$$

The costates satisfy the canonical system

$$\dot{\Lambda}_i(t) = -\frac{\partial}{\partial X_i} \mathcal{H}_i(X_i(t), \Lambda_i(t), \tau_1, \tau_2), \quad \tau_{i-1} < t < \tau_i. \quad (4)$$

Since the RHS of the ODEs (4) are continuously differentiable, the costates can continuously be extended to the left and right side of the interval $\tau_{i-1} < t < \tau_i$. Thus

$$\Lambda_i(\tau_{i-1}) := \lim_{t \rightarrow \tau_{i-1}^+} \Lambda_i(t) \quad \text{and} \quad \Lambda_i(\tau_i) := \lim_{t \rightarrow \tau_i^-} \Lambda_i(t)$$

uniquely exist. For the derivatives of the Hamiltonians with respect to the switching times we find

$$\frac{\partial}{\partial \tau_1} \mathcal{H}_i(X_i, \Lambda_i, \tau_1, \tau_2) = \begin{cases} 0 & i = 1, 2 \\ \gamma_3(\tau_1, \tau_2) K \sigma \gamma_3(\tau_1, \tau_2)^{\sigma-1} W^\sigma + \Lambda'_i \frac{\partial}{\partial \tau_1} \text{SLIR}_i(X_i, \tau_1, \tau_2) & i = 3 \end{cases}$$

$$\frac{\partial}{\partial \tau_2} \mathcal{H}_i(X_i, \Lambda_i, \tau_1, \tau_2) = \begin{cases} 0 & i = 1, 2 \\ -\gamma_3(\tau_1, \tau_2) K \sigma \gamma_3(\tau_1, \tau_2)^{\sigma-1} W^\sigma + \Lambda'_i \frac{\partial}{\partial \tau_2} \text{SLIR}_i(X_i, \tau_1, \tau_2) & i = 3 \end{cases}$$

The Hamiltonian in the third stage (after the lockdown) and the salvage value explicitly depend on the switching times τ_1 and τ_2 . Thus for $0 < \tau_1 < \tau_2 < T$ the necessary optimality conditions at the switching times τ_1 and τ_2 write as (cf. [26]).

$$\mathcal{H}_2(X_2(\tau_1), \Lambda_2(\tau_1), \tau_1, \tau_2) - \mathcal{H}_1(X_1(\tau_1), \Lambda_1(\tau_1), \tau_1, \tau_2) = \int_{\tau_2}^T \frac{d}{d\tau_1} \mathcal{H}_3(X_3(t), \Lambda_3(t), \tau_1, \tau_2) dt + K \frac{d}{d\tau_1} \gamma(T, \tau_1, \tau_2)$$

$$\mathcal{H}_3(X_3(\tau_2), \Lambda_3(\tau_2), \tau_1, \tau_2) - \mathcal{H}_2(X_2(\tau_2), \Lambda_2(\tau_2), \tau_1, \tau_2) = \int_{\tau_2}^T \frac{d}{d\tau_2} \mathcal{H}_3(X_3(t), \Lambda_3(t), \tau_1, \tau_2) dt + K \frac{d}{d\tau_2} \gamma(T, \tau_1, \tau_2)$$

and at τ_i , $i = 1, 2$ the costates satisfy

$$\Lambda_1(\tau_1) = \Lambda_2(\tau_1)$$

$$\Lambda_2(\tau_2) = \Lambda_3(\tau_2).$$

At the endtime T the costates satisfy the transversality condition

$$\Lambda(T) = \begin{pmatrix} \lambda_1(T) \\ \lambda_2(T) \\ \lambda_3(T) \\ \lambda_4(T) \end{pmatrix} = \begin{pmatrix} -\Gamma K \sigma W(T)^{\sigma-1} \gamma(T, \tau_1, \tau_2)^\sigma \\ -\Gamma K \sigma W(T)^{\sigma-1} \gamma(T, \tau_1, \tau_2)^\sigma \\ 0 \\ -\Gamma K \sigma W(T)^{\sigma-1} \gamma(T, \tau_1, \tau_2)^\sigma \end{pmatrix}.$$

S2 Appendix: Derivation of R_0 For the derivation of the reproductive rate R_0 of the SLIR model (1) we assume $\nu = \mu$. Then, for every S there exists a disease free equilibrium $(S, 0, 0, 0)$ of system (1). Using the expansion formula for the determinant along the first column it can easily be seen that the stability of this equilibrium only depends on the dynamics of the latent and infected population. Thus, we consider the linearization of the L, I system

$$\dot{L}(t) = \beta \frac{S(t)(I(t) + fL(t))}{N(t)} - (\mu + \varphi)L(t)$$

$$\dot{I}(t) = \omega\varphi L(t) - (\alpha + \mu + \mu_I)I(t)$$

yielding

$$J = \begin{pmatrix} f\beta - \varphi - \mu & \beta \\ \omega\varphi & -\alpha - \mu - \mu_I \end{pmatrix} \quad (5)$$

The characteristic equation of Eq. (5) yields

$$\xi^2 + \xi(\varphi + 2\mu + \alpha + \mu_I - f\beta) + (\alpha + \mu + \mu_I)(\varphi + \mu - f\beta) - \beta\omega\varphi.$$

A simple calculation shows that a necessary and sufficient condition for $\text{Re } \xi_{1,2} < 0$ is

$$(\alpha + \mu + \mu_I)(\varphi + \mu - f\beta) - \beta\omega\varphi > 0$$

which is equivalent to

$$\frac{\beta(\omega\varphi + f(\alpha + \mu + \mu_I))}{(\alpha + \mu + \mu_I)(\varphi + \mu)} < 1.$$

Setting

$$R_0 := \frac{\beta(\omega\varphi + f(\alpha + \mu + \mu_I))}{(\alpha + \mu + \mu_I)(\varphi + \mu)}$$

we find that the disease free equilibrium is locally asymptotically stable iff $R_0 < 1$, and R_0 being the reproductive rate.

References

1. Acemoglu, Daron Introduction to Modern Economic Growth Princeton University Press: 2009.
2. Acemoglu, Daron and Chernozhukov, Victor and Werning, Iván and Whinston, Michael D. A Multi-Risk SIR Model with Optimally Targeted Lockdown National Bureau of Economic Research, WP 27102; 2020.
3. Alvarez, Fernando E. and Argente, David and Lippi, Francesco A Simple Planning Problem for COVID-19 Lockdown National Bureau of Economic Research, WP 26981; 2020.
4. Atkeson, Andrew What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios National Bureau of Economic Research, Working Paper 26867; 2020.
5. Baker, Scott R. and Bloom, Nicholas and Davis, Steven J. and Kost, Kyle and Sammon, Marco and Viratyosin, Tasaneeya The unprecedented stock market reaction to COVID-19 The Review of Asset Pricing Studies; 2020.
6. Bendavid, Eran and Mulaney, Bianca and Sood, Neeraj and Shah, Soleil and Ling, Emilia and Bromley-Dulfano, Rebecca and Lai, Cara and Weissberg, Zoe and Saavedra-Walker, Rodrigo and Tedrow, James and Tversky, Dona and Bogan, Andrew and Kupiec, Thomas and Eichner, Daniel and Gupta, Ribhav and Ioannidis, John and Bhattacharya, Jay COVID-19 Antibody Seroprevalence in Santa Clara County, California medRxiv, Cold Spring Harbor Laboratory Press; 2020.
7. Bloom, David E and Kuhn, Michael and Prettner, Klaus Modern infectious disease: Macroeconomic impacts and policy responses Journal of Economic Literature; 2020.
8. Brodeur, Abel and Gray, David M. and Islam, Anik and Bhuiyan, Suraiya A Literature Review of the Economics of COVID-19. IZA Discussion paper No. 13411; 2020
9. CDC COVID-19 Pandemic Planning Scenarios National Center for Immunization and Respiratory Diseases (NCIRD), Division of Viral Diseases; 2020.
10. Delius, Gustav W and Powell, Benedict J and Bees, Martin A and Constable, George W A and MacKay, Niall J and Pitchford, Jonathan W More prevalent, less deadly? Bayesian inference of the COVID19 Infection Fatality Ratio from mortality data medRxiv, Cold Spring Harbor Laboratory Press; 2020.

11. Eichenbaum, Martin S. and Rebelo, Sergio and Trabandt, Mathias The Macroeconomics of Epidemics National Bureau of Economic Research, WP 26882; 2020.
12. Flaxman, Seth and Mishra, Swapnil and Gandy, Axel and Unwin, H. Juliette T. and Mellan, Thomas A. and Coupland, Helen and Whittaker, Charles and Zhu, Harrison and Berah, Tresnia and Eaton, Jeffrey W. and Monod, Mélodie Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe Nature 584, No. 7820, 257-261; 2020.
13. Grass, Dieter and Caulkins, Jonathan P. and Feichtinger, Gustav and Tragler, Gernot and Behrens, Doris A. Optimal Control of Nonlinear Processes: With Applications in Drugs, Corruption, and Terror. Berlin: Springer-Verlag.; 2008.
14. Gonzalez-Eiras, Martín and Niepelt, Dirk On the Optimal 'Lockdown' during an Epidemic Swiss National Bank, Study Center Gerzensee, CEPR Discussion Paper 14612; 2020.
15. Hammit, James K. Valuing mortality risk in the time of COVID-19. SSRN 3615314; 2020.
16. Huberts, Nick and Thijssen, Jacco Optimal Timing of Interventions during an Epidemic University of York; 2020.
17. Kermack, William Ogilvy and McKendrick, Anderson Gray Contributions to the mathematical theory of epidemics—I Proceedings of the Royal Society 115 (772), 700-721; 1927.
18. Kniesner, Thomas J. and Viscusi, W. Kip and Woock, Christopher and Ziliak, James P. The Value of a Statistical Life: Evidence from Panel Data The Review of Economics and Statistics 94(1), 74-87; 2012.
19. Kretzschmark, Mirjam E. and Rozhnova, Ganna and Bootsma, Martin C.J. and van Boven, Michiel and van de Wijgert, Janneke H.H.M. and Bonten, Mark J.M. Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modeling study. The Lancet Public Health 5, No. 8, e452-e459; 2020.
20. Krueger, Dirk and Uhlig, Harald and Xie, Taojun Macroeconomic Dynamics and Reallocation in an Epidemic National Bureau of Economic Research, WP 27047; 2020.
21. Lander, Willem and Abrams, Steven and Petrof, Oana and Coletti, Pietro and Kuylen, Elise and Libin, Pieter and Mogelmoose, Signe and Wambua, James and Herzog, Sereina A. and Faes, Christel and Beutels, Philippe and Hens, Niel The impact of contact tracing and household bubbles on deconfinement strategies for COVID-19: an individual-based modelling study. medRxiv; 2020.
22. Layard, Richard and Clark, Andrew E. and De Neve, Jan-Emmanuel and Krekel, Christian and Fancourt, Daisy and Hey, Nancy and O'Donnell, Gus When to release the lockdown? A wellbeing framework for analysing costs and benefits. IZA Institute of Labor Economics, IZA DP No. 13186; 2020.
23. Niel, Hens and Vranck, Pascal and Molenberghs, Geert The COVID-19 epidemics, its mortality, and the role of non-pharmaceutical interventions. European Heart Journal: Acute Cardiovascular Care 9, No. 3, 204-208; 2020.
24. Piguillem, Facundo and Shi, Liyan Optimal COVID-19 Quarantine and Testing Policies Einaudi Institute for Economics and Finance, EIEF Working Papers Series 2004; 2020.

25. Richardson, Safiya and Hirsch, Jamie S. and Narasimhan, Mangala and Crawford, James M. and McGinn, Thomas and Davidson, Karina W. and and the Northwell COVID-19 Research Consortium Presenting Characteristics, Comorbidities, and Outcomes Among 5700 Patients Hospitalized With COVID-19 in the New York City Area JAMA; 2020.
26. Tomiyama, Ken and Rossana, Robert J. Two-stage optimal control problems with an explicit switch point dependence: Optimality criteria and an example of delivery lags and investment Journal of Economic Dynamics and Control 13 (3), 319-337, 1989.
27. Tsai, Thomas C. and Jacobson, Benjamin H. and Jha, Ashish K. American Hospital Capacity And Projected Need for COVID-19 Patient Care Health Affairs Blog; 2020.