



## Decision Support

## The hammer and the jab: Are COVID-19 lockdowns and vaccinations complements or substitutes?



J. P. Caulkins<sup>a</sup>, D. Grass<sup>b</sup>, G. Feichtinger<sup>c</sup>, R. F. Hartl<sup>d,\*</sup>, P. M. Kort<sup>e</sup>, M. Kuhn<sup>b,f</sup>,  
A. Prskawetz<sup>b,f,g</sup>, M. Sanchez-Romero<sup>b,g</sup>, A. Seidl<sup>d</sup>, S. Wrzaczek<sup>b,f</sup>

<sup>a</sup> Heinz College, Carnegie Mellon University, Pittsburgh, USA

<sup>b</sup> International Institute for Applied Systems Analysis (IIASA), Schlossplatz 1, Laxenburg 2361, Austria

<sup>c</sup> Department for Operations Research and Control Systems, Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria

<sup>d</sup> Department of Business Decisions and Analytics, University of Vienna, Vienna, Austria

<sup>e</sup> Tilburg School of Economics and Management, Tilburg University, Tilburg, Netherlands

<sup>f</sup> Wittgenstein Centre for Demography and Global Human Capital (IIASA, VID/OeAW, University of Vienna), Austria

<sup>g</sup> Research Group Economics, Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria

## ARTICLE INFO

## Article history:

Received 14 April 2022

Accepted 19 April 2023

Available online 26 April 2023

## Keywords:

OR In government

COVID-19

Vaccinations

Dynamic optimization

SIR Models

## ABSTRACT

The COVID-19 pandemic has devastated lives and economies around the world. Initially a primary response was locking down parts of the economy to reduce social interactions and, hence, the virus' spread. After vaccines have been developed and produced in sufficient quantity, they can largely replace broad lock downs. This paper explores how lockdown policies should be varied during the year or so gap between when a vaccine is approved and when all who wish have been vaccinated. Are vaccines and lockdowns substitutes during that crucial time, in the sense that lockdowns should be reduced as vaccination rates rise? Or might they be complementary with the prospect of imminent vaccination increasing the value of stricter lockdowns, since hospitalization and death averted then may be permanently prevented, not just delayed? We investigate this question with a simple dynamic optimization model that captures both epidemiological and economic considerations. In this model, increasing the rate of vaccine deployment may increase or reduce the optimal total lockdown intensity and duration, depending on the values of other model parameters. That vaccines and lockdowns can act as either substitutes or complements even in a relatively simple model casts doubt on whether in more complicated models or the real world one should expect them to always be just one or the other. Within our model, for parameter values reflecting conditions in developed countries, the typical finding is to ease lockdown intensity gradually after substantial shares of the population have been vaccinated, but other strategies can be optimal for other parameter values. Reserving vaccines for those who have not yet been infected barely outperforms simpler strategies that ignore prior infection status. For certain parameter combinations, there are instances in which two quite different policies can perform equally well, and sometimes very small increases in vaccine capacity can tip the optimal solution to one that involves much longer and more intense lockdowns.

© 2023 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

## 1.1. Motivating question and countries' policies to date

The COVID-19 pandemic has been devastating for health and livelihoods. Millions died; more than one hundred million lost their jobs (see [International Labour Organization \(2021\)](#) and [World Health Organization \(2022\)](#)). The health consequences are familiar and need not be elaborated, except to note that they are heavily

\* Corresponding author.

E-mail addresses: [caulkins@andrew.cmu.edu](mailto:caulkins@andrew.cmu.edu) (J.P. Caulkins), [grass@iiasa.ac.at](mailto:grass@iiasa.ac.at) (D. Grass), [gustav.feichtinger@tuwien.ac.at](mailto:gustav.feichtinger@tuwien.ac.at) (G. Feichtinger), [richard.hartl@univie.ac.at](mailto:richard.hartl@univie.ac.at) (R.F. Hartl), [kort@tilburguniversity.edu](mailto:kort@tilburguniversity.edu) (P.M. Kort), [kuhn@iiasa.ac.at](mailto:kuhn@iiasa.ac.at) (M. Kuhn), [afp@econ.tuwien.ac.at](mailto:afp@econ.tuwien.ac.at) (A. Prskawetz), [Miguel.Sanchez@oeaw.ac.at](mailto:Miguel.Sanchez@oeaw.ac.at) (M. Sanchez-Romero), [andrea.seidl@univie.ac.at](mailto:andrea.seidl@univie.ac.at) (A. Seidl), [wrzaczek@iiasa.ac.at](mailto:wrzaczek@iiasa.ac.at) (S. Wrzaczek).

concentrated among older populations. Hence, most of the economic dislocation comes not directly from death and disease but rather from efforts to stem transmission.

Some measures, such as wearing masks and having knowledge workers telecommute, are relatively painless, but most governments also 'locked down' varying proportions of their economies. Lockdowns are expensive in terms of lost economic output and lost freedom to travel, and only a handful of countries succeeded in avoiding widespread infection altogether until the majority had been vaccinated.

Vaccines are relatively cheap and confer at least partial immunity for a period of time. Several were invented, tested, and approved in record-breaking time, but it took approximately 6–24 months (depending on the country) to produce and distribute enough to vaccinate everyone who wants it.

That raises the question of what should happen to lockdown intensity during a vaccine roll out. Put differently, the response to COVID-19 or any subsequent viral pandemic may have (at least) three phases. In phase I, the response is limited to non-pharmacological interventions (NPI), notably lockdowns. In phase III, widespread vaccination lets economic, social and travel activity return close to normal. In between, there is a crucial intermediate phase during which nations need to figure out how to balance lockdowns and vaccination.

Even before vaccines were developed, most countries adjusted lockdown intensity up and down in response to changing conditions, trying to strike some balance between economic hardship and infection. Intuitively, increasing rates of vaccination ought to affect that balancing, but it's not immediately clear how.

At least two arguments spring to mind, albeit pointing in opposite directions. One is that since vaccination removes people from the pool of susceptibles, thereby reducing the epidemic's effective "reproductive rate" if all else were held equal (Adam (2020)), one can relax the lockdown commensurately without triggering a resurgence. That approach views vaccines and lockdowns as substitutes; as vaccines arrive, they ease up on lockdowns.

Another view is that harsh lockdowns might not be feasible or desirable for decades on end, but if full vaccination is right around the corner, it would be foolish to risk getting sick now. Sticking with stiff lockdowns just a little longer might prevent death or lifelong complications of "long COVID". According to that view, vaccines and lockdowns can be complements. The vaccine increases the appeal of a strict lockdown that would otherwise be prohibitively long and expensive.

The stakes in this balancing are high. Total years of life lost in developed countries is approximately two to nine times that of the typical year of seasonal flu (Arolas et al. (2021)). A June 2021 Congressional Research Service report put that global total in just the first year or so at 3.7 million deaths, lost work equivalent to 255 million full-time jobs, 95 million more people living in extreme poverty and 80 million more people malnourished because of the combined effects of COVID-19 infections and lockdowns (Jackson, Weiss, Schwarzenberg, & Nelson (2021)).

Auld & Toxvaerd (2021) approach this question empirically, describing what countries actually did. They find substitution: countries that rolled out vaccines quickly relaxed their lockdown policies faster than did countries that rolled out vaccines slowly, although not so much as to override the health-protecting effects of vaccination.

We think the history is a bit more complicated. Fig. 1 plots the average lockdown intensity vs. proportion of the population that had been vaccinated to date using data from OurWorldInData.org. It is a vector field with an arrow for each country showing how those quantities changed between April and August, 2021 for a range of countries distinguished by their income level.

The horizontal axis measures the proportion of the population that is fully vaccinated plus 0.5 times those who are partially vaccinated (so two partially vaccinated people is equivalent to one fully vaccinated). The vertical axis is the average lockdown intensity over the preceding four months (averaging is preferred since the current month's lockdown intensity can be affected by particular outbreaks).

Whereas Auld & Toxvaerd (2021) find only substitution and we find below that countries should dynamically adjust lockdowns as vaccination expands, the flatness of most arrows in Fig. 1 suggests that most countries maintained a relatively stable lockdown intensity over this period. A few arrows point down, consistent with substitution, but there are almost as many up as down arrows on the left side of the graph (places that had barely started vaccination).

These descriptive historical data are interesting, but countries may or may not have behaved optimally for COVID-19, and the world may be challenged by other pandemics in the future. That makes it desirable to investigate this question within a prescriptive framework.

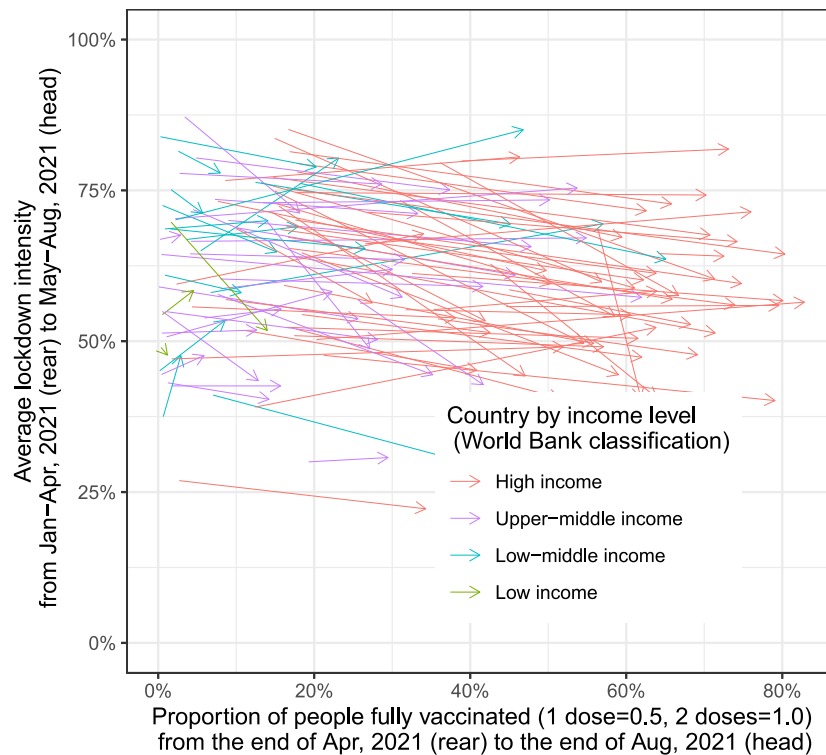
Hence, we investigate a simple dynamic optimization model of the epidemic's spread and resulting health and economic costs. The analysis starts with time 0 being the moment when the vaccine has been approved. The decision maker can then modulate lockdown intensity up or down as vaccinations are deployed.

We find that even within this relatively simple model, it can be sensible for lockdowns and vaccinations to co-exist, although there is not a monotonic relationship between lockdowns and vaccines either over time or in total. Sometimes it is optimal to ease lockdowns steadily as vaccination progresses, but sometimes it is optimal first to increase lockdown stringency for a time, and only begin reducing it later. Likewise, there are conditions under which expanding vaccine production capacity reduces the total lockdown effort, and conditions when the opposite is true. Section 4 investigates this analytically, and shows that when vaccination proceeds sufficiently slowly, there is a complementary relationship, and when vaccination proceeds sufficiently fast, they are substitutes. Numerical examples in Section 5 confirm that and find other surprises, including situations with multiple optimal solutions, times when it makes sense to "allow" some of the population to get infected before imposing strict lockdowns, and even instances in which expanding vaccination capacity slightly increases rather than reduces COVID-19 deaths under the optimal policy. Much depends on the parameter values and initial conditions (i.e., the state of the epidemic at the time the vaccine is approved), but in ways that support sensible economic interpretation.

## 2. Literature review

The literature on COVID-19 is vast. We confine ourselves to discussing some prescriptive societal-level dynamic optimization models that address at least two of the following three topics: COVID-19, lockdowns, and vaccination. For the papers on COVID-19 and lockdown we only concentrate on some of the first papers. An overview of the considered literature can be found in Table 1, where we characterize the papers by the type of model, the control instruments and whether vaccination and lockdowns are found to complement or substitute each other (if applicable).

Papers in which lockdowns and vaccination are simultaneously considered as controls are still scarce. Most papers have focused either on vaccination or lockdowns, but not on both. At the beginning of COVID-19 pandemic, the first set of papers modeled the optimal start and end time of a lockdown and assumed that the intensity of the lockdown was constant (Bosi, Camacho, & Desmarchelier, 2021; Caulkins et al., 2020). Bosi et al. (2021) show that the intensity of the lockdown depends on the degree of altruism



**Fig. 1.** Change of lockdown stringency (vertical axis) and proportion vaccinated (horizontal axis) over four months in 2021. Source: Authors' calculations using [Hannah et al. \(2020\)](#). Each arrow is associated with a single country. The color of the arrow represents the income level of the country according to the World Bank income level classification.

**Table 1**  
Summary of the literature on COVID-19.

Reference	Epidemiological model	Optimal policies/Controls		Complementarity vs. substitutability
		Lockdowns	Vaccination	
<a href="#">Acemoglu et al. (2020, 2021)</a>	Heterogeneous SIR	✓		Complements
<a href="#">Alfaro et al. (2020)</a>	Heterogeneous SIR	✓		
<a href="#">Alvarez et al. (2021)</a>	SIR	✓		Complements
<a href="#">Aspri et al. (2021)</a>	SEAIRD	✓		
<a href="#">Bosi et al. (2021)</a>	SIS	✓		
<a href="#">Buratto et al. (2022)</a>	SIR	✓	✓	Complements
<a href="#">Caulkins et al. (2021, 2022)</a>	SLIR	✓		
<a href="#">Caulkins et al. (2020)</a>	SIR	✓		
<a href="#">Eichenbaum et al. (2022)</a>	SIR	✓		Complements
<a href="#">Farboodi et al. (2021)</a>	Heterogeneous SIR	✓		
<a href="#">Federico et al. (2022)</a>	SIRS		✓	
<a href="#">Federico &amp; Ferrari (2021)</a>	SIR	✓		
<a href="#">Freiberger et al. (2022)</a>	Heterogeneous SIR	✓		
<a href="#">Fu et al. (2021)</a>	SIR	✓		
<a href="#">Fu et al. (2022)</a>	SIRV	✓	✓	Complements
<a href="#">Garriga et al. (2022)</a>	SIRSV	✓	✓	Complements
<a href="#">Giagheddu &amp; Papetti (2023)</a>	Heterogeneous SIR	✓		Complements
<a href="#">Gonzalez-Eiras &amp; Niepelt (2020)</a>	SIR	✓		Complements
<a href="#">Hsu et al. (2020)</a>	Heterogeneous SIR	✓		
<a href="#">Huberts &amp; Thijssen (2023)</a>	SIR	✓		
<a href="#">Jones et al. (2021)</a>	SIR	✓		
<a href="#">Kaplan et al. (2020)</a>	SEIR	✓		
<a href="#">Libotte et al. (2020)</a>	SIR		✓	
<a href="#">Mak et al. (2022)</a>	Heterogeneous SEIR		✓	
<a href="#">Makris (2021)</a>	Heterogeneous SIR	✓		
<a href="#">Piguillem &amp; Shi (2022)</a>	SEIR	✓		Both
<a href="#">Rao &amp; Brandeau (2021, 2022)</a>	Heterogeneous SIR		✓	

in a society. [Caulkins et al. \(2020\)](#) show, using the valuation of life by a decision marker, that the overall performance of an eradication strategy and a curve flattening strategy can be the same under certain conditions.

Various models started including the possibility of several full lockdowns ([Huberts & Thijssen, 2023](#)) or that the lockdown in-

tensity could vary over time. For instance, [Aspri, Beretta, Gandolfi, & Wasmer \(2021\)](#) determine the optimal end of lockdowns of an economy and distinguish different optimal policies related to how much life is valued. The authors find testing of particular importance for reducing mortality with only mild economic losses. Similar results as those presented by [Aspri et al. \(2021\)](#) are found

by Eichenbaum, Rebelo, & Trabandt (2022) in a SIR-macro model with different degrees of disutility for social distance. Alfaro, Faia, Lamersdorf, & Saidi (2020) also shows that the intensity of the lockdown and its evolution over time depend on social preferences (i.e., the degree of altruism, patience, etc.). Jones, Philippon, & Venkateswaran (2021) include the possibility of working from home and show that a social planner will reduce the number of fatalities more than if decisions about social distance are made by households, though at the expense of a sharper drop in consumption.

Caulkins et al. (2021, 2022) confirm that decision makers with similar preferences can prefer dramatically different lockdown strategies, extending Caulkins et al. (2020) by studying the optimal lockdown intensity and including lockdown fatigue effects. Federico & Ferrari (2021) identify the optimal lockdown intensity for different phases of a pandemic, where the transmission rate evolves according to a stochastic differential equation. Makris (2021) consider the social distance (i.e. the lockdown) as individual decision in a non-cooperative game.

Fu, Xiang, Jin, & Wang (2021) allow for the possibility of imposing a different lockdown to susceptible and infected people. Kaplan, Moll, & Violante (2020) study the costs associated with lockdowns across individuals working in different economic sectors and how to distribute the costs. Freiberger, Grass, Kuhn, Seidl, & Wrzaczek (2022) determine the optimal target transmission rate of the disease (which is influenced by social distancing measures such as isolation and lockdowns) and the optimal testing efforts in a model that accounts for different population groups or regions. Lockdowns in context of multiple regions are also studied in Hsu, Lin, & Yang (2020), which show that lockdown policies may differ across countries and may lead to different welfare effects depending on the epidemiological situation.

Various authors have started to introduce the possibility of a vaccine arrival exogeneously into models with optimal lockdown strategies. Some examples are Acemoglu, Chernozhukov, Werning, & Whinston (2020, 2021); Alvarez, Argente, & Lippi (2021); Farboodi, Jarosch, & Shimer (2021); Giagheddu & Papetti (2023); Gonzalez-Eiras & Niepelt (2020); Piguillem & Shi (2022) which investigate the consequences of different vaccine arrival times, which is frequently modeled through a Poisson process, on the strength of the lockdown. In many of these models, once that the vaccine is made available, it is assumed that all individuals are vaccinated and the pandemic is over. They find that an earlier arrival of the vaccine implies a stringent lockdown, which Acemoglu et al. (2020) denote as “wait-for-the-vaccine” effect and Alvarez et al. (2021) denote this effect as “dynamic complementarity”. Buratto, Muttoni, Wrzaczek, & Freiberger (2022) also address a similar question with a stochastic arrival time of a vaccine. They find that at the time when the vaccine becomes available, the lockdown intensity is intensified discontinuously, which is a result of the additional possibility for susceptibles to obtain protection against the disease. However, Piguillem & Shi (2022) also find that the impact of the vaccine arrival may have a non-monotonic effect on lockdowns. In particular, if vaccines are either not expected or they are expected to arrive immediately, the optimal strategy is not to use lockdowns, while if vaccines are expected in the near future they find that stringent lockdowns can also be optimal.

Another set of models considers that only the vaccination rate can be optimally decided. For instance, Libotte, Lobato, Platt, & Silva Neto (2020) consider an optimal control model where the control is the vaccination rate and the objective is to minimize the number of infected individuals as well as the amount of vaccines used. Rao & Brandeau (2021, 2022) include heterogeneous agents and additional objectives such as minimizing deaths or the quality-adjusted life years lost. Federico, Ferrari, & Torrente (2022) solve an optimal vaccination problem in which immunized individuals

can become susceptible. Mak, Tiglong, & Tang (2022) analyze the inventory dynamics associated with rolling out a COVID-19 vaccine that requires two doses to be spaced apart by a fixed time interval. However, neither Mak et al. (2022), Federico et al. (2022); Rao & Brandeau (2021, 2022), nor Libotte et al. (2020) deal with whether vaccination should replace or complement a lockdown.

Closer to our paper are Fu, Jin, Xiang, & Wang (2022) and Garriga, Manuelli, & Sanghi (2022), which analyze the impact of vaccines on lockdown policies and hence whether lockdowns and vaccines are complements or substitutes. Fu et al. (2022) consider an optimal control problem in which the policy maker optimizes lockdown intensity while recognizing that takes time to vaccinate the population. They find it necessary to establish a lockdown upon vaccine arrival, however, the lockdown should gradually be relaxed as vaccination expands. Garriga et al. (2022) design a dynamic macro model in which a lockdown decreases output but reduces infections within a standard SIR model with endogenously chosen vaccination rates. Our model differs from Garriga et al. (2022) in that we force the decision maker to grapple with limitations in intensive care capacity. While Garriga et al. (2022) find that lower vaccination capacities lead to a less restrictive lockdown policy, our results show that lockdown intensity and vaccination can be substitutes, not just complements.

### 3. The model

The backbone of the present model is a so-called *SIR* model (see Kermack & McKendrick (1927)) in which people with active infections (denoted by state  $I(t)$ ) spread the infection to those who are susceptible (state  $S(t)$ ) before moving on to a non-infectious or ‘recovered’ state. The familiar Bass model from marketing (compare Bass (1969)) is very similar, with early adopters of a new product ‘infecting’ others through word of mouth or ‘viral’ spread of new product adoption.

We adapt the COVID-19 *SIR* model presented in Caulkins et al. (2021) in two ways. First, there are now two distinct ways for someone to be ‘recovered’. One, denoted  $R_1(t)$ , is for people who have already been infected before time  $t$  but not vaccinated. The second,  $R_2(t)$ , is for people who have been vaccinated. It is possible for someone who has recovered in the first sense to then be vaccinated, in which case they move from state  $R_1(t)$  to state  $R_2(t)$ .

Second, for simplicity, we do not consider the possibility that the population rebels and undermines lockdown restrictions, referred to as lockdown fatigue in Caulkins et al. (2021). Thus, the instantaneous effectiveness of locking down is assumed to be determined by its current intensity and is independent of its past history.

#### 3.1. Control variables

The key time-varying policy variable is  $\gamma(t)$ , the proportion of the economy that is allowed to operate, i.e., that is not ‘locked down’. This proportion is allowed to move continuously between 0 and 1 (0 means full lockdown, 1 means no lockdown). In practice, lockdowns tend to be chunky. E.g., restaurants may be restricted to delivery and take-out or they may allowed also to operate in-person dining at 25%, 50%, 75% or full capacity. That is only five options, not a continuum. However, restaurant restrictions can be applied to just some jurisdictions within the country and not others, and they can be mixed and matched with similar restrictions on other industries. Considering all of those many combinations, it is not unreasonable to view the extent of locking down as a continuously varying control variable.

Having big parts of the economy shut down is costly, but adjusting rates of employment takes time and is costly too. That is,



**Table 2**  
Notation.

Control variables:	
Lockdown intensity adjustment	$u(t)$
<b>State variables:</b>	
Susceptibles	$S(t)$
Infected	$I(t)$
Recovered from infection	$R_1(t)$
Vaccinated	$R_2(t)$
Proportion of economy that is not 'locked down'	$\gamma(t)$
<b>Functions:</b>	
Infection rate	$\beta(\gamma(t))$
Lockdown adjustment costs	$V_u(u(t))$
Economic costs	$V_l(W(t), \gamma(t))$
Health costs	$V_h(I(t))$
Residual economic costs at time $T$	$S(X(T))$

bouncing back and forth between having 10 or 30% percent of the economy shut down is more costly than maintaining a steady 20% lockdown. So rather than setting  $\gamma(t)$  directly, it is adjusted by the control variable  $u(t)$ , which is costly to employ.

As in [Caulkins et al. \(2021\)](#), we make those costs quadratic (so rushed changes are more costly than slow and steady adjustments) and allow re-opening activity to be more difficult (costly) than closing down. Here, though, both costs are only one-tenth as large as in our previous papers.<sup>1</sup> This permits  $\gamma(t)$  to adjust quickly so the specific value chosen for  $\gamma(0)$  is not terribly important.

$$V_u(u(t)) := \begin{cases} c_l u^2(t) & u(t) \leq 0 \\ c_r u^2(t) & u(t) > 0 \end{cases} \quad (1)$$

In general,  $c_l$  and  $c_r$  can differ, meaning that increasing or relaxing the lockdown can involve different costs.

The lockdown affects economic output (discussed below) and decreases the SIR model's infection rate  $\beta$  through some combination of fewer contacts (e.g., people staying home) and lower risk of infection per contact (e.g., because of rules requiring minimum separation distances). The infection rate is assumed to be decreasing and convex with respect to the intensity of the lockdown, since an intelligent social planner would close first those businesses that generate the highest infection risks per unit of employment or economic output (e.g. night clubs or large public events might be closed first). A simple functional form that captures this idea is:

$$\beta(\gamma) = \beta_1 \gamma^\theta, \quad \beta_1 > 0, \theta > 1. \quad (2)$$

In theory this equation permits the social planner to eliminate all transmission by shutting down the entire economy (i.e., pushing  $\gamma$  down to 0), which may not be realistic since various essential services must continue to operate. That might suggest a need to place a rigid lower bound on how low  $\gamma$  can become. However, the diminishing returns in this relationship and the objective function make exceedingly low  $\gamma$  suboptimal, even if they are formally allowed. In practice, optimal trajectories tend to keep  $\gamma$  above 0.5 or 0.6, which correspond to reducing transmission by two-thirds to three-quarters, not driving it to 0.

Different countries were in different stages of the epidemic when vaccines first became available, but with some notable exceptions (e.g., Australia, New Zealand) in many developed countries between two (Germany) and seven (US) percent of the population was already known to have been infected. Given how many cases are asymptomatic or otherwise not recorded, we start the model with ten percent of the population in the  $R_1$  state. Also at that time, the number of new cases detected per day was roughly in the range of 20 (Canada and the EU) to 90 per 100,000 (in the

US and UK).<sup>2</sup> That suggests initial rates of infection in the range of 0.6% to 2.7%, assuming people remain in the  $I$  state for 15 days and half of all infections are detected. So we set  $I(0) = 1.67\%$ .

Having recovered from a past infection is at least partially protective even if it does not confer full immunity. Hence, the infection-minimizing strategy would focus vaccination on people who are susceptible, and administer it to previously infected individuals in the  $R_1$  state later. However, most countries did not do this. So in our basic model vaccinations are randomly spread across susceptible and those in the  $R_1$  state. In a sensitivity analysis we explore how much better a country could do if it were to prioritize vaccinating those who had not yet been infected.

We assume that the rate of vaccination is constrained by some exogenous rate,  $b$ , that reflects physical limits on production and/or distribution. We focus primarily on results for  $b$  large enough that it does not take very long until substantially everyone could be vaccinated. That time appears to have been on the order of six months for Israel, the U.K., U.S., Europe and some Gulf Arab states, and as long as two years in some places, so we vary this parameter in sensitivity analysis below.

The dynamic optimization continues through a fixed time horizon  $T$  that is long enough for everyone in such countries to get vaccinated; specifically, we set  $T = 1095$  days, or three years. Some sensitivity analyses will show results for smaller values of  $b$ , meaning that the country's own vaccine production capacity is insufficient to vaccinate everyone within 3 years. For those countries, terminating the dynamic modeling after three years is equivalent to assuming that at that point, some international body swoops in and vaccinates everyone who is still susceptible. It is somewhat artificial to model such a vaccine "end game" as happening instantaneously, but we are primarily interested in  $b$  values large enough that this termination condition would not come into play.

Vaccines are extremely cheap compared to both the COVID-19's health harms and the economic harms created by lockdowns. Indeed, vaccines are so cheap that their monetary costs are effectively negligible for rich countries and poorer countries often receive them free or at subsidized rates.<sup>3</sup> Consequently, vaccination costs are not considered and vaccination is never withheld; the rate of vaccination continues at a pace of  $b$  until demand for vaccination has been met. That happens at a time that is approximately the reciprocal of  $b$ , but the total number of people who get vaccinated depends slightly on how many get infected and die, so the exact end point of the dynamic optimization is not trivial to compute a priori.

This approach ignores "vaccine hesitancy". We of course understand that not everyone will get vaccinated, but modeling in this way does not assume that. It only assumes that enough people will get vaccinated – or infected – that economy-wide lockdowns are no longer optimal.

We do think there are interesting dynamic optimization questions that arise when a large proportion of the population is vaccine hesitant, including the issuing of "vaccine passports" and possibly opening some activities only to those who have been vaccinated, but those issues are not the focus here.

### 3.2. State equations

Infected individuals recover at rate  $\alpha$ , which is the inverse of the average dwell time in the infectious state, taken here to be 15 days. Since we assume an infection fatality rate (IFR) of 1% (dis-

<sup>2</sup> <https://91-divoc.com/pages/covid-visualization/>

<sup>3</sup> Vaccines cost about \$40 per person. So vaccinating everyone in the United States would cost about  $\$40 * 330$  million people = \$13B, which is about the GDP produced in 5–6 hours. Even partial lockdowns extended for months destroy vastly more economic value than that.

<sup>1</sup> The parameters used in the present paper are summarized in [Section 5, Table 3](#).

**Table 3**  
Parameter values.

Parameter	Value	Description
$\alpha$	0.066667	reciprocal of the 15-day average duration of the infection
$\beta_1$	0.13333	leading coefficient governing how lockdowns affect infection risk
$\theta$	2	exponent of lockdown efficiency in the infection risk term
$\nu$	$2.7397e - 005$	birth rate (1 percent per year as a daily rate)
$\mu$	$2.7397e - 005$	death rate not caused by COVID-19 (also 1 percent per year as a daily rate)
$\mu_I$	$0.01 * \alpha$	Daily rate of death for infected COVID patients, equal to an IFR of $0.01 * \alpha$
$p$	0.02311	probability that an infected person needs critical care
$\zeta$	5000	parameter in the continuous approximation of the max-function
$I_{max}$	0.000176	capacity of intensive care units to treat COVID-19 patients
$\sigma$	0.66667	labor elasticity in the Cobb-Douglas production function
$K$	1	coefficient of economic activity; it defines the units with which the objective function is measured to be daily GDP per capita
$M$	*	social cost of a premature death due to COVID-19
$\xi_1$	0.03	daily death rate of infected individuals who receive critical care, set equal to a 45% probability times alpha to convert to a daily rate
$\xi_2$	0.036667	incremental death rate if ICU capacity is exceeded, set equal to $55% * \alpha$
$c_l$	100	parameter in adjustment costs representing the costs of shutting businesses down
$c_r$	500	corresponding adjustment cost parameter for the costs of reopening businesses
$b$	*	daily vaccination rate
$\bar{T}$	365	constant in salvage value function

discussed below), that suggests including an outflow from the  $I$  state due to COVID-19 deaths at a rate  $\mu_I = 0.01/15$ . We also allow natural fertility and mortality at rates,  $\nu$  and  $\mu$ , of 1% per year.

COVID-19 deaths in the objective function are modeled using a more sophisticated two-part function that recognizes an incremental death risk when hospitals' critical care capacity is swamped. That is omitted from the state dynamics for simplicity; COVID-19 deaths are sufficiently rare that it matters little whether this incremental bump in COVID-19 death risk is modeled explicitly.

Because the planning horizon is relatively short, we do not consider the possibility that recovered individuals may lose their immunity, i.e. there is no backflow from  $R_1(t)$  or  $R_2(t)$  to  $S(t)$ .

As a result the model dynamics are (suppressing the time argument)

$$\begin{aligned} \dot{S} &= \nu N - \beta(\gamma) \frac{SI}{N} - \frac{bS}{S + R_1} - \mu S \\ \dot{I} &= \beta(\gamma) \frac{SI}{N} - (\mu + \mu_I + \alpha)I \\ \dot{R}_1 &= \alpha I - \frac{bR_1}{S + R_1} - \mu R_1 \\ \dot{R}_2 &= b - \mu R_2 \\ \dot{\gamma} &= u \\ \beta(\gamma) &:= \beta_1 + \beta_2 \gamma^\theta \\ 0 &\leq \gamma(t) \leq 1, \quad \forall t, \end{aligned} \tag{3}$$

where  $X_0 = (S(0), I(0), R_1(0), R_2(0), \gamma(0))$  is a set of initial conditions that correspond to the state of the epidemic when the vaccine is approved.

### 3.3. Objective function

As in [Caulkins et al. \(2021\)](#), the decision maker optimizes an objective function which consists of i) lockdown adjustment costs  $V_u(u(t))$  (as stated in (1)), ii) economic costs, and iii) health costs.

Economic costs concern lost economic output. We assume production obeys a Cobb-Douglas production function where capital  $K$  is constant (since the time horizon is short) and  $\sigma$  is the labor share (set to  $\frac{2}{3}$ ). It is assumed that infected people are sick or quarantined and so cannot work, i.e. in the absence of a lockdown, the number of workers would be  $W(t) = S(t) + R_1(t) + R_2(t)$ . Since  $(1 - \gamma(t))$  denotes the proportion of the economy that is shut down,  $W(t)\gamma(t)$  is the number of people who are able to work. That means that economic output at time  $t$  is  $KW^\sigma(t)\gamma^\sigma(t)$ .

Since we assume the birth rate matches the death rate,  $\nu = \mu$ , if there had never been an epidemic, the number of workers would have been a fixed constant which without loss of generality is scaled to 1.0. There also would have been no reason to lockdown, so the rate of production would have simply been  $K$ . Thus the economic loss at any given time is  $V_l(W(t), \gamma(t)) := K - KW^\sigma(t)\gamma^\sigma(t)$ .

Dynamic modeling of infections ends at time  $T$ , but there may be residual economic costs that accrue beyond that time if the economy is below full employment at time  $T$ . We model those as being equal to the output gap at time  $T$  times a proportionality constant  $\bar{T}$ :

$$\bar{T}(K - KW^\sigma(T)\gamma^\sigma(T)). \tag{4}$$

We set  $\bar{T}$  to 1 year. One way to interpret that value of  $\bar{T}$  is that residual unemployment decays linearly over the course of two years. This salvage value function is included for completeness but it turns out not to be large compared to costs incurred before time  $T$ .

Health costs related to COVID-19 are dominated by costs of lost lives and are primarily modeled based on the CDC's best guess scenario (Scenario 5) as described in their May 2021 guidance for COVID-19 planning scenarios. We motivate the key parameters here, but see [Caulkins et al. \(2020, 2021, 2022\)](#) for additional details. Note that our epidemic parameters are intended to reflect the virus as of the time that vaccination began in late 2020. In reality, the delta and omicron variants subsequently came to dominate, and their parameters would be different.

The infection fatality rate (IFR) grows exponentially with age, so the population-wide average IFR depends on the age distribution of infections. Weighting the CDC's age-specific IFR's (as of May 2021) by the overall age distribution of the U.S. population would suggest an average IFR of 1.3%, but younger people tend to have more social contacts and so may be more likely to get infected (even if the consequences of infection are much more severe for older people), so we round this down to 1%.

Rather than just imposing a single fixed IFR, we distinguish the death risk when adequate critical care is available from the situation in which so many people are infected that many cannot receive adequate care. In particular, we estimate the probability  $p$  that an infected person will need an ICU bed or other critical care, and then the death rate per day for such people when they do receive that care ( $\xi_1$ ) and the additional, incremental death risk when they do not ( $\xi_2$ ).

Weighting the CDC’s age-specific probabilities of death given hospitalizations by the age-distribution of COVID-19 hospitalizations reported by COVID-NET gives an overall probability of death given hospitalization of 11.1%, which in turn implies a probability of hospitalization given infection of 1% over 11.1% = 9%. The similarly age-weighted CDC probabilities of needing an ICU given hospitalization is 25%, suggesting that the probability of needing an ICU given infection is 2.25%. Our past papers set this parameter  $p$  to 2.311%, which is nearly identical, so we retain the previous value for continuity.

Apart from the initial surge in certain places, notably New York City, for the most part demand for critical care did not exceed supply in the U.S. Hence, dividing the overall IFR by this probability  $p$  suggests that those who reached the point of needing critical care had a 45% probability of dying even if such a person received critical care. If all such very sick individuals would have died if they did not receive care, that implies an incremental death risk of 55%, and so the corresponding COVID-19 specific death rates per day are  $\xi_1 = 3\%$  and  $\xi_2 = 3.667\%$ .

Tsai, Jacobson, & Jha (2020) suggest that 58,166 of the 84,750 ICU beds in the U.S. could be used for treating COVID-19 patients. That is 176 per million people. The age-weighted average time in hospital for those needing critical care is 12.4 days according to CDC scenario planning guidance. Since that is close to the average dwell time in the  $I$  state, we set  $H_{\max} = 0.000176$ , which is rather similar to Charpentier, Elie, Laurière, & Tran (2020)’s value of 0.0002 based on data from France.

Thus we would like to model COVID-19 deaths as

$$(\xi_1 pl(t) + \xi_2 \max\{0, pl(t) - H_{\max}\}), \tag{5}$$

but the maximum-function is not differentiable at  $pl = H_{\max}$ , which hinders a qualitative analysis of the numerical solution. Thus we use  $\frac{1}{\zeta} \log(1 + e^{\zeta(pl - H_{\max})})$  as an approximation. The larger  $\zeta$  is, the better the approximation, and for the value of  $\zeta = 5000$  used here, the approximation is excellent, as illustrated in Caulkins et al. (2020, 2021). Thus, introducing a parameter  $M$  for the cost per premature death due to COVID-19, the health costs can be modeled as:

$$V_h(I(t)) := M \left( \xi_1 pl(t) + \frac{\xi_2}{\zeta} \log(1 + e^{\zeta(pl(t) - H_{\max})}) \right). \tag{6}$$

It is very difficult to place a relative value on lost work vs. lost lives, but without loss of generality we set  $K = 1$  (meaning that the value function is denominated in days-of-GDP) and then consider a very wide range for  $M$ . Economic analyses often value preventing a premature death at between 20× and 150× annual GDP per capita (Alvarez et al. (2021), Kniesner, Kip Viscusi, Woock, & Ziliak (2012)), but Hammit (2020) argues that lower values may be preferred for analysis of COVID-19 deaths, so we consider a range from 10× to 150× annual GDP per capita.

One argument for lower values is that COVID-19 is much more deadly for people who are older, which might suggest that years-of-life lost (YLL) per death is low. But the age distribution of COVID-19 deaths resembles that for all other causes combined. E.g., in the U.S. about 70 per cent of COVID-19 deaths are among those 70 years old and older, which is only slightly greater than that group’s 64 per cent share of all deaths. Hence, the years of life lost per fatality is around 11.7 years (Goldstein & Lee (2020)). That is less than for deaths from homicide or traffic crashes, but more than for some diseases of old age, such as Alzheimer’s.

Of course, deaths are not the only health-related cost. There are also costs of treatment and illness, but for COVID-19 they tend to be smaller than the social costs of death. Since the model’s controls do not materially alter the ratio of severe vs. mild cases, the costs of non-fatal illness can be incorporated by adopting some-

what larger values of  $M$  than one would to account for just deaths alone.

Conversely, readers who believe that lockdowns harm outcomes not captured in current GDP (e.g., mental health anguish from isolation, loss of education, etc.) may prefer smaller values of  $M$ . There is no one right value of  $M$ , so it is treated as a bifurcation parameter to explore the implications of a full range of  $M$  values.

In sum, the policy maker aims to minimize the following objective function, where  $(X(t))$  denotes a vector of all state variables at  $t$ ):

$$\mathcal{V}(X_0) := \int_0^T V_l(W(t), \gamma(t)) + V_h(I(t)) + V_u(u(t)) dt + S(X(T)), \tag{7}$$

where

$$\begin{aligned} V_l(W(t), \gamma(t)) &:= K(1 - W^\sigma(t))\gamma^\sigma(t) \\ V_h(I(t)) &:= M \left( \xi_1 pl(t) + \frac{\xi_2}{\zeta} \log(1 + e^{\zeta(pl(t) - H_{\max})}) \right) \\ V_u(u(t)) &:= \begin{cases} c_l u^2(t) & u(t) \leq 0 \\ c_r u^2(t) & u(t) > 0 \end{cases} \\ S(X(T)) &:= \bar{T}(K - KW^\sigma(T))\gamma^\sigma(T) \\ W(t) &:= S(t) + R_1(t) + R_2(t). \end{aligned} \tag{8}$$

Table 2 summarizes the state and control variables, as well as the function of the model.

#### 4. Analytical insights

##### Overview of analytical results


Our central question is how the arrival of vaccines should affect the overall duration and intensity of the optimal lockdown policy. Here we address that question analytically, but numerical analyses are also needed (see Section 5) because the answer can depend on the values of two key parameters:  $M$ , which stands for the value of preventing a COVID-19 death, and  $b$  which denotes the rate of vaccine production.

To be clear, the outcome of interest here is the total amount of locking down done over the entirety of the planning horizon, not the intensity of the lockdown at any single point in time. Formally, it is the integral over time of one minus  $\gamma$ .

A natural jargon for this question is the economic concept of substitutes vs. complements. As vaccines become “cheaper” (more available) should society “buy” more or less locking down? To the extent that vaccines are an alternative way to “purchase” the saving of lives, vaccines may substitute for lockdowns, meaning that there should be less locking down if the rate of vaccination is higher. Alternately, lockdowns – which slow the costly (“bad”) way of exiting the susceptible state – might become more appealing if vaccines create a complementary (“good”) exit path, in which case one would “spend” more (in terms of lost economic output) buying more lockdown protection if vaccine availability increased.

One can also think about this by analogy to the Slutsky equation that decomposes a change in the price level of a good into a substitution and income effect. Increasing vaccine production  $b$  not only changes the relative price of saving lives via lockdowns or vaccination, it also makes society better off. So increasing vaccine production capacity  $b$  may have both an “income” and a “substitution” effect.

Slutsky applied that decomposition in a comparative statics framework, whereas here we have optimization over time. So we examine effects on the total (i.e., integral over time) amount of locking down, loss in economic production, and number of COVID-19 deaths. When we say that one policy uses lockdowns to a greater degree than does another policy, that doesn’t mean it is more stringent at every point in time. Indeed, there is no specific



**Fig. 2.** Optimal trajectories over time for employment (left panel) and infections (right panel) with a larger (solid lines) or smaller (dashed lines) vaccine production capacity, where both trajectories share the same total amount of locking down.

requirement as to the “shape” of the policy. Rather, it just means that in total over the planning horizon (or, equivalently, on average), there is more lockdown effort.

We derive two fundamental results. First, if the value of preventing a COVID-19 death,  $M$ , is small enough, then increasing the rate of vaccination will reduce the total amount of locking down (substitution) if  $b$  is sufficiently large, and increase it (complementarity) if  $b$  is sufficiently small. Second, when  $M$  is large enough, then increasing the rate of vaccination will always reduce the total amount of locking down (substitution).

Thinking about the  $b - M$  parameter space, these theorems explain what happens near two corners and along the opposite edge. They do not explain what happens in the middle or how large that middle is. The next section’s numerical solutions suggest that more complicated behavior can emerge in the middle, but that the middle is not so large. I.e., the qualitative insights summarized in the theorems for extreme cases are fair guides as to the character of the overall solution in most cases.

Before stating these theorems, we define three aggregate performance measures, with which different optimal solutions can be compared.

**Contour lines**

**Definition 1.** Consider the optimal control problem (7), (8) and (3) and let the superscript  $*$  denote variables that are optimal. Then the total amount of locking down  $\Gamma^*(b, M)$ , the total value of lost economic production  $\mathcal{V}_l^*(b, M)$ , and the total cost of COVID-19 deaths  $\mathcal{V}_h^*(b, M)$  are defined as

$$\begin{aligned} \Gamma^*(b, M) &:= \int_0^{T_1} (1 - \gamma^*(t))dt \\ \mathcal{V}_l^*(b, M) &:= \int_0^T V_l(W^*(t), \gamma^*(t))dt \\ \mathcal{V}_h^*(b, M) &:= \int_0^T V_h(I^*(t))dt, \end{aligned} \tag{9}$$

where  $T_1$  denotes the end time of the lockdown. In case of multiple lockdowns,  $T_1$  denotes the end time of the last one.

These performance measures give intuition about the qualitative change of the optimal solution when certain model parameters are changed. In particular, here and in Section 5 we focus on the parameters  $b$  and  $M$ .

The following lemma characterizes the slope of the three sets of contour lines defining  $b - M$  pairs that produce the same value of one of these performance measures. The proof of this lemma and all other results in this section can be found in Appendix C.

**Lemma 2.** Consider the optimal control problem (7), (8) and (3) and let the superscript  $*$  denote variables that are optimal. Then the derivative of  $b$  with respect to  $M$  along the contour lines of the performance measures defined in Definition 1 can be derived as follows:

$$\begin{aligned} \frac{db}{dM} \Big|_{\Gamma^*(b, M)=const.} &= - \frac{\int_0^{T_1} \frac{\lambda_M(t)}{\lambda_\gamma(t)} dt}{\int_0^{T_1} \frac{\lambda_b(t)}{\lambda_\gamma(t)} dt} \\ \frac{db}{dM} \Big|_{\mathcal{V}_l^*(b, M)=const.} &= - \frac{\int_0^T \frac{\lambda_M(t)}{\lambda_{V_l}(t)} dt}{\int_0^T \frac{\lambda_b(t)}{\lambda_{V_l}(t)} dt} \\ \frac{db}{dM} \Big|_{\mathcal{V}_h^*(b, M)=const.} &= - \frac{\int_0^T \frac{\lambda_M(t)}{\lambda_{V_h}(t)} dt}{\int_0^T \frac{\lambda_b(t)}{\lambda_{V_h}(t)} dt} \end{aligned} \tag{10}$$

where  $T_1$  denotes the time when the lockdown is finally terminated, i.e.  $\gamma(T_1) = 1$ , and  $\lambda_x(t)$  denotes the shadow price of  $x$ .

This analytic representation of the contour lines permits an intuitive interpretation. A ratio of shadow prices refers to the exchange rate (or, in other words, marginal rate of substitution) between two variables or parameters given that a certain performance measure remains constant. Consider, for instance, the contour line with respect to the total amount of locking down  $\Gamma^*(b, M)$ . The denominator equals the ratio of the shadow prices of  $b$  and  $\gamma$ . That denotes their exchange rate, i.e. how much  $\gamma$  would the decision maker trade for a marginal increase of  $b$ . The numerator denotes the exchange rate between  $\gamma$  and  $M$ . Both the numerator and denominator are integrated over time, since  $\Gamma^*(b, M)$  is defined as the total lockdown over the planning period. Thus the derivative of the contour line equals the aggregated exchange rate between  $M$  and  $\gamma$  over the aggregated exchange rate between  $b$  and  $\gamma$ . I.e. due to the intertemporal property of the total lockdown, both  $M$  and  $b$  are first exchanged to  $\gamma$  at every  $t$  and then related by the fraction of the summation of the instantaneous effects over time.

**Existence of both complementarity and substitution for sufficiently small  $M$**

It might be tempting to think that vaccination would always substitute for locking down, since they are both ways to reduce COVID-19 deaths, but the following theorem shows that is wrong. In particular, the following existence result formalizes conditions that determine whether the vaccination is a complement or a substitute to lockdowns for a fixed  $M$ .

The initial premise of the theorem is that when  $M = 0$  there is an optimal solution that does not lock down the economy. That is reasonable since if COVID-19 infections are costless, then there is little reason to take expensive steps to avert them. The premise is






Fig. 3. How the optimal total amount of locking down varies with vaccine production capacity for  $M = 7300$ .

not entirely innocuous because infections also keep people away from work for the (relatively brief) time people are in state  $I$ . But as a practical matter, few would advocate shutting down the economy if COVID-19's only downside was keeping people away from work while they are in state  $I$ . Thus, for parameter sets that reasonably model the real world, the premise of the theorem will be satisfied.

**Theorem 3** (Existence). *Let for  $M = b = 0$  and  $\gamma_0 = 1$  the optimal solution of problem (7), (8) and (3) be  $(\gamma^*(t), u^*(t)) = (1, 0)$  for all  $t$ . Then there exists  $\bar{M} > 0$  such that for every  $M$  with  $0 \leq M < \bar{M}$ ,  $\underline{b}$  and  $\bar{b}$  exist, where the lockdown intensity reaches a maximum with respect to  $b$ , at  $b^* \in [\underline{b}, \bar{b}]$ . For lower values of  $b$ , i.e.  $b < \underline{b}$ , vaccination and lockdown are complements. For higher values of  $b$ , i.e.  $b > \bar{b}$ , they are substitutes. This translates into*

$$\frac{d\Gamma^*(b, M)}{db} > 0 \quad \text{for } b < \underline{b} \tag{11}$$

$$\frac{d\Gamma^*(b, M)}{db} < 0 \quad \text{for } b > \bar{b} \tag{12}$$

If  $b^*$  is the unique maximum,  $\underline{b} = \bar{b} = b^*$  holds.

The theorem tells us that for values of  $M$  not too large, the total amount of locking down goes up with vaccination capacity  $b$  if  $b$  is small ( $b < \underline{b}$ ), whereas it goes down with  $b$  if  $b$  is large ( $b > \bar{b}$ ). In other words, the theorem proves the existence of both a complementary region located at  $b < \underline{b}$ , and a substitution region located at  $b > \bar{b}$ . In most cases (for example, see Figs. 3 and 4 in the next section) it holds that there is just one complementary region that transitions into the substitution region at  $b = \underline{b} = \bar{b}$ . However, uniqueness of the maximum of  $\Gamma^*$  is not guaranteed, implying that it does not always hold that  $\underline{b} = \bar{b}$ , so there may be not just one complementary and one substitution region.

**Corollary 4.** *Let the properties of Theorem 3 be satisfied, then  $\underline{\gamma}$  exists such that the result holds for all initial values of  $\gamma(\cdot)$  with  $\underline{\gamma} < \gamma(0) \leq 1$ .*

**Non-existence of only complementarity for sufficiently large  $M$**

When  $M$  is large, preventing deaths is of prime importance. So, when vaccination capacity is low, it is optimal to implement a strict lockdown. That lockdown can be relaxed when vaccination capacity increases (substitution), as is proved in the following theorem.

**Theorem 5** (Existence). *Consider the optimal solution of problem (7), (8) and (3). Then there exists a sufficiently high  $\bar{M}$  such that for all  $M \geq \bar{M}$ , all  $b \geq 0$  and all  $\gamma_0 \in (0, 1]$  vaccination and lockdown are substitutes.*

**First order condition and interpretation**

The optimal solution permits some additional, economic interpretation. The model above can be solved with the Maximum Principle of optimal control theory (see e.g. Grass, Caulkins, Feichtinger, Tragler, & Behrens (2008)), the derivations of which are included in Appendix A. From the Hamiltonian maximizing condition with respect to the lockdown adjustment control (see (15) in Appendix A) we get the following necessary first order condition that has to hold whenever there is a lockdown, i.e. when  $0 < \gamma(t) < 1$ ,

$$\frac{\partial \mathcal{V}_u(u)}{\partial u} = -\lambda_\gamma(T) - \int_t^T \left( \frac{\partial V_t(W, \gamma)}{\partial \gamma} + (\lambda_I - \lambda_S) \frac{\partial \beta(\gamma)}{\partial \gamma} \frac{SI}{N} \right) ds. \tag{13}$$

As usual in optimization theory this condition balances marginal costs and effects. The left hand side of (13) denotes the marginal adjustment costs at  $t$  and the right hand side the marginal net benefit of the relaxation of the lockdown. While the adjustment costs at  $t$  are static (accruing at  $t$  only), the effect of the marginal lockdown relaxation is intertemporal. It consists of the marginal effect of the lockdown on the economic loss during the economic recovery (term outside the integral), marginal production during the pandemic (first term in integral), and the marginal value of changes in transmissions times the contact probability times  $\lambda_I - \lambda_S$ . The latter two effects are intertemporal and therefore integrated over time.  $\lambda_I - \lambda_S$  is the difference in the marginal value of infected and susceptible people and can be interpreted as the value of one susceptible person getting infected.

As a final result, the following corollary formulates a certain property of the optimal solution under the condition that the optimal solution prescribes that the lockdown will not be intensified over the remaining planning period. In particular, we find that in that case vaccination and lockdown are substitutes. Such a situation typically occurs during a final time interval leading up to the point in time when all susceptibles are vaccinated.

**Corollary 6.** *Consider the optimal control problem (7), (8) and (3). For fixed  $b$  and  $M$  vaccination and lockdown are substitutes if the lockdown is never intensified during the planning period, i.e.  $\dot{\gamma}(t) \geq 0$  for  $t \in [0, T]$ .*

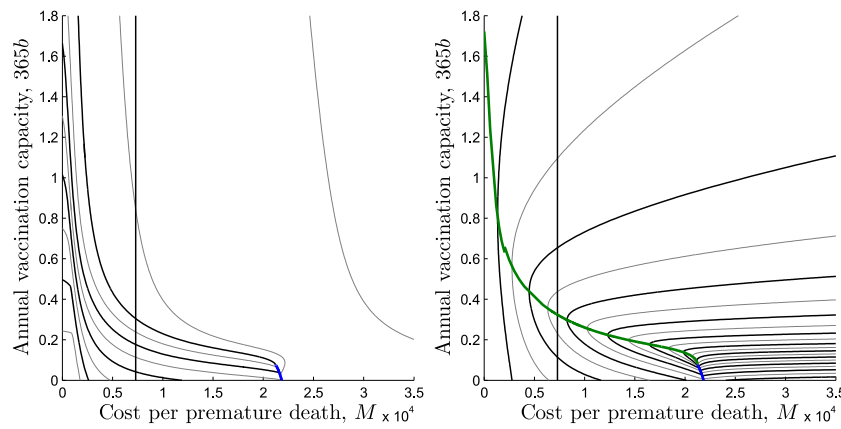
**5. Numerical results**

We next investigate numerically how should lockdown intensity change after a vaccine is approved? And does that change depend on whether the vaccine production capacity is higher or lower?

For the basecase parameter values given in Table 3 and a value of life equal to 20 times annual GDP per capita (so  $M = 7300$  since time is measured in days), Fig. 2 shows two very different optimal solutions depending on the vaccination capacity  $b$ . The left-hand panel plots the evolution over time of the state variable  $\gamma$ , which is one minus the optimal lockdown intensity, and the right-hand panel plots the resulting number of people who are infected. The fact that two sharply divergent trajectories can both be optimal shows that the overall answer to our central motivating question is "it depends". In this case it depends on the parameter values, specifically the vaccination rate.

The lines are color coded with blue indicating periods of increasing lockdown stringency, green indicating relaxation, and red indicating that the lockdown is over.

For the solid line (large  $b$ ), the change in lockdown intensity is monotonic over time, apart from a brief initial adjustment because



**Fig. 4.** Contour lines showing equal numbers of COVID-19 deaths (left panel, the value of the leftmost curve is  $\frac{110}{7300}$  and of the rightmost curve  $\frac{25}{7300}$ , the other values are 30, 40, 50, ...,  $110 \times \frac{1}{7300}$ ) and total amounts of locking down (right panel) for varying annual rates of vaccination (vertical axis) and values on preventing a COVID-19 death. Vertical line at  $M = 7300$  indicates the value of  $M$  used in the preceding figure.

the initial lockdown was not severe enough. If the parameter  $\gamma(0)$  had been chosen optimally, the solid time path would have been monotonic (i.e., green from the outset). As more people are vaccinated, less lockdown is needed, as described in Corollary 6. That is perhaps the expected pattern, but the dashed line shows that it is not the only type of solution that can be optimal.

The dashed curve has one additional segment. As with the solid curve, the brief initial (green in this case) phase is an adjustment phase of little consequence. The same initial lockdown intensity is now too heavy to be optimal when vaccination capacity  $b$  is low. Therefore, the decision maker starts by briefly reducing the lockdown, i.e., increasing  $\gamma$ .

Naturally infections increase (roughly doubling for these parameter values) until the lockdown is expanded enough to be fairly strong (midway through the blue part of the trajectory). At some point, i.e. when blue passes into green, enough people have been vaccinated (or infected) that it is possible to start relaxing the lockdown without the infection rate rebounding. This relaxation phase (green) ends when everyone is vaccinated (joins the red curve), which happens later than in the solid line case because the vaccination capacity,  $b$ , is smaller.

To summarize, the solid and dashed lines show almost opposite optimal responses to the onset of vaccinations, depending on the potential vaccination rate. In neither case is the lockdown trajectory monotonic over time, and the non-monotonicity in the dashed curve is fundamental. We include in Appendix B a more detailed analysis of the dynamics of the lockdown intensity.

In Fig. 2, the larger  $b$  induced a tougher lockdown at its peak, but the lockdown was shorter. The total amount of locking down is exactly the same for both trajectories. Of course, it is hard to assess from just those two runs what is the overall dependence of lockdowns on  $b$ . So Fig. 3 plots the total lockdown-induced suppression of employment on the vertical axis (i.e., the integral of one minus  $\gamma$  over time, or  $\Gamma^*$  from Definition 1) versus the vaccine production capacity parameter  $b$ .

Where this curve slopes upward, vaccines and lockdowns are complements, as a greater vaccination capacity produces a greater total lockdown effort. Where the curve slopes down, a greater vaccination capacity substitutes for locking down. That the curve slopes up to a peak at  $b = 8.9 \times 10^{-4}$  (i.e., requiring roughly three years to vaccinate everyone) before declining, shows that vaccines can be either complements or substitutes for lockdowns. I.e., it shows that the conditions of Theorem 3 pertain not just for some arbitrarily small  $M$ , but even for the basecase value of  $M$ .

These results all pertained to  $M = 7300$ , meaning that preventing a COVID-19 death is valued at 20 times annual GDP per capita.

Other things can happen for other values of  $M$ . For instance, when  $M$  is rather small, the analog to Fig. 3 has two humps, not just one (figure not shown). That is why in Theorem 3 in Section 4, there could be complementarity for all  $b < \underline{b}$  and substitution for all  $b > \underline{b}$ , but  $\underline{b}$  and  $\bar{b}$  could be different, with  $\underline{b} < \bar{b}$ .

Fig. 4 is a bifurcation diagram that simultaneously varies not only the vaccine production capacity ( $b$ , on the vertical axis), but also  $M$  (horizontal axis). For ease of interpretation the vertical axis measures the annual rate not the daily rate of vaccination (i.e., 365 times  $b$  not  $b$ ), so 0.5 corresponds to it taking 2 years to vaccinate everyone, 1 corresponds to being able to vaccinate everyone in one year, etc.

The contour lines in the left hand panel denote combinations of  $b$  and  $M$  that produce the same number of COVID-19 deaths, with large numbers of deaths in the lower left and few deaths in the upper right. The contour lines are widely spaced to the upper right of the black line, and densely packed below it. That is because in broad terms two different – indeed almost opposite – strategies can be optimal. When  $b$  and  $M$  are large enough, the optimal strategy is to lock down hard and long enough that relatively few people get infected. In that part of the graph, even big increments in  $M$  produce modest reductions in deaths, because there are already relatively few deaths. The other strategy is to lock down sparingly even though that means herd immunity is achieved by allowing people to get infected, not just through vaccination. Within that region, increasing either  $b$  or  $M$  alters the proportion of people who reach the recovered state via infection vs. vaccination, so it more greatly affects the number of deaths.

One surprising result is that in a small region with high  $M$  and small  $b$  the contour lines of death are backward bending. That means there are some  $M - b$  combinations such that increasing  $b$  moves to a higher not lower contour line. I.e., increasing vaccine production capacity can slightly increase not reduce the number of COVID-19 deaths when following the optimal strategy. The reason is that lockdowns are relaxed so much in response to greater vaccination rates, that more people get infected.

The contour lines in the right hand panel denote combinations of  $b$  and  $M$  that elicit the same total amount of locking down (i.e., the same integral of one minus  $\gamma(t)$ ). That amount is greatest in the lower right, when averting a COVID-19 death is highly valued (large  $M$ ) and the vaccinate production capacity is low (small  $b$ ), so it is optimal to maintain an intense lockdown for a long time, until enough vaccine is finally produced.

When  $M$  is large, increasing  $b$  sharply reduces the total amount of locking down. That is not because lockdowns become much less

intense, but rather because they can be shorter; the larger  $b$  means it does not take as long for everyone to be vaccinated.

On the far left side of the graph, when  $M$  is very low – or, equivalently, job losses are seen as very painful – it is optimal not to use lockdowns much regardless of the vaccine production capacity,  $b$ .

But for a broad range of  $M$  values that the literature views as reasonable – between about 10 and 75 times GDP per capita – increasing the vaccination rate  $b$  from 0 up to the black or blue line increases the total amount of locking down that is optimal. (Vaccination and lockdowns are complements.) Beyond that, further increases in vaccination capacity substitutes for locking down.

For the rest of this section, we will refer just to “the black line”. The blue line in the lower right is a Skiba curve where there are multiple optimal solutions. It reveals some mathematically fascinating behavior, but we defer its discussion to the next section because it only appears for combinations of  $M$  and  $b$  that would be rare (i.e., a wealthy country that highly values saving lives but which nonetheless has very limited access to the vaccine). The blue line is not exactly an extension of the black line; there is a tiny overlap, but for purposes of thinking about whether vaccines complement or substitute for lockdowns, it is not too great a simplification for now to think about the blue line as just extending the black line.

This model is highly stylized, but roughly speaking, one can position different countries in different places on this plane. The UK is in the upper right, with both intensive lockdowns and high rates of vaccinations. Brazil might be in the lower left, with low rates of vaccinations and less intense lockdowns. Arguably Portugal is towards the lower right (hard lockdown with lower – at least initially – vaccination rates) and the U.S. in the upper left, with high rates of vaccination following relatively low concern and lackadaisical lockdowns (in at least some states).

Unless new variants emerge that defeat current vaccines, most rich nations are probably above the black line, in the substitution region. They committed to protecting most of their population from infection by doing whatever amount of locking down was necessary. For them, increasing vaccination production capacity translates primarily into less (shorter) unemployment.

Other parts of the world had to wait longer for full vaccination (lower  $b$ ), and at least some may exhibit values of  $M$  that are below the black line. For them, in this model expanding vaccination capacity reduces COVID-19 deaths, but it might also reduce their economic output (if they feel obligated to use any vaccines received) because that can induce longer lockdowns.

Whether any given country is below the black line is hard to judge; the numerical value of  $M$  represents an ethical judgment, not a scientific fact, and people will differ. Larger values of  $M$  are not necessarily “better” or “more humane”. Nor should lower  $M$  values be seen as materialistic or shallow.  $M$  can be seen as trading off death by COVID-19 vs. poverty and death by poverty. It also embodies the tradeoff between COVID-19 deaths and other adverse health outcomes. Even in rich countries the “discretionary” healthcare interrupted by lockdowns included important services like cancer screening. In sub-Saharan Africa lockdowns also interrupted actions against major killers like malaria [Aborode et al. \(2020\)](#).

The contour lines for total economic loss (not shown) are very similar to those for the total amount of locking down because almost all of the reduction in economic output comes from locking down, not from people being sick or dying of COVID-19. That happens within this model even though it does not distinguish ages. In reality, since COVID-19 deaths are concentrated among retirees, even less of the reduction in economic production would be due to COVID-19 deaths than is shown in this model; almost all of it would be due to locking down.

We can integrate the two panels of [Fig. 4](#) by thinking about how a country should “spend” the potential wealth created by having greater vaccination capacity. That the contour lines in the upper right of the left hand panel are nearly vertical shows that rich countries should use greater vaccine production capacity almost exclusively to buy reductions in unemployment; for them, increasing  $b$  has almost no effect on deaths. The right hand panel shows that countries that start out below the black line are in the opposite situation. For them, there is always some increase in vaccine production capacity that leaves them on the same isolockdown contour line but on the other side of the black line, and so with far fewer people dying, but essentially the same amount of economic dislocation. Such countries would be using that greater vaccine capacity to buy reductions in COVID-19 deaths.

## 6. Sensitivity analysis

This section explores several sensitivity analyses.

### 6.1. Restricting vaccination to susceptibles

Only vaccinating those who are still susceptible would avert more infections than would distributing vaccines also to those who have already developed (at least partial) immunity by previously being infected. That is not the policy most nations are pursuing. However, [Fig. 5](#) shows how much better the outcomes could be in that case, while holding all other parameters at their base case values (including that 10 percent of the population begins in the  $R_1$  at time 0).

The left-hand panel shows that the bottom-line answer is “not much” because the optimal lockdown policy is similar (center panel) and there are only modestly fewer deaths (right panel), at least with these parameter values.

The benefit of wiser vaccine allocation is largest (biggest gap between dashed and solid lines in left-hand panel) when vaccine capacity is intermediate. When vaccination capacity is very small, there is less lost by not using it efficiently. When vaccine capacity is great, there is little penalty from “wasting” some vaccines on people who already have (some) immunity due to prior infection, although the benefit could be larger if more than 10 percent of the population began in the  $R_1$  state.

### 6.2. Skiba points

We return now to the blue curve in the lower right of [Fig. 4](#). It is hard to see, but the contour lines are discontinuous along the blue curve. That is because every point on that curve is a Skiba point from which two entirely different optimal solutions emanate, so two entirely different amounts of locking down are optimal.

One solution involves intense locking down and few deaths; it is an extension of the (unique) solutions that are optimal for parameter constellations to the upper right of the blue curve. The other, an extension of the (unique) solutions to the lower left, involves much less locking down and so more deaths. Exactly on that blue curve, either approach performs equally well in terms of total social cost.

For a thorough discussion of Skiba curves, also called DNS or DNSS curves, see [Grass et al. \(2008\)](#), but they can informally be thought of as the optimized version of “tipping points”. Here the thresholds are depicted in parameter space, but they are easier to think about in a space defined by initial conditions for state variables, although the principles are similar. Starting at any point on one side of a Skiba curve, there is one well-defined, unique optimal solution. By continuity these are all similar to each other

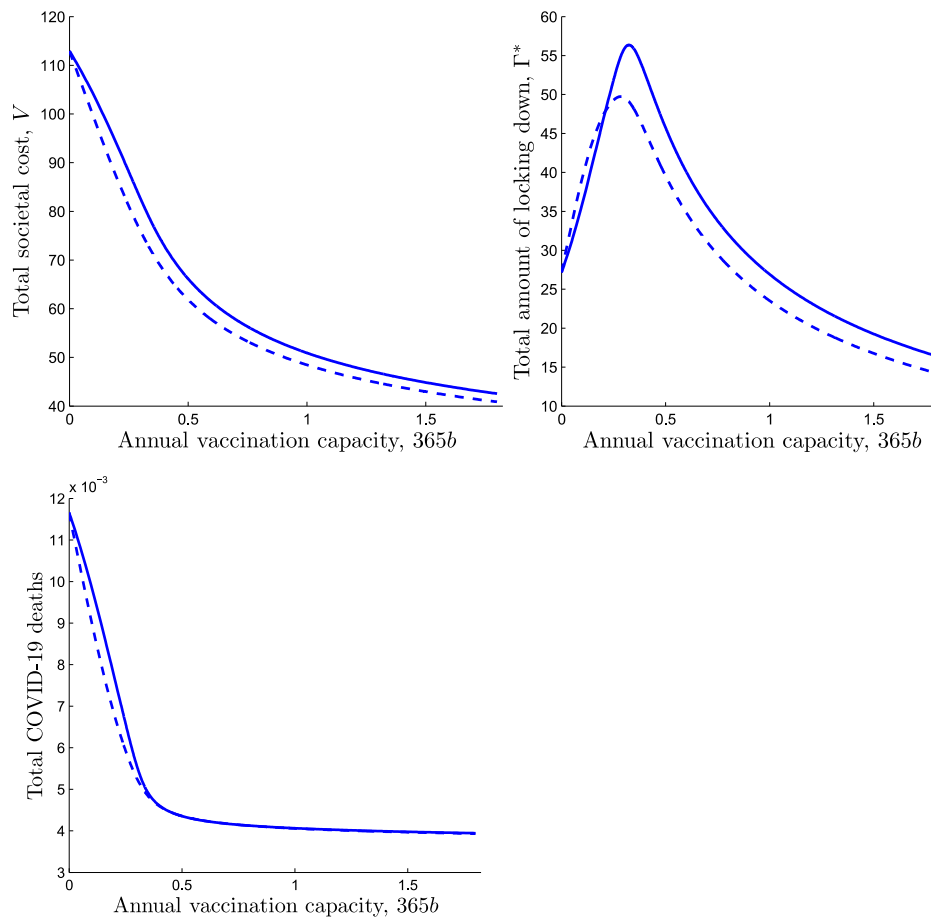


Fig. 5. Improvement when only vaccinating susceptibles (dashed lines) not also those previously infected (solid lines) for total social cost (left panel), the total amount of locking down  $\Gamma^*$  (center panel), and the number of deaths (right panel).

and so can thought of as a family reflecting one strategy or approach. Likewise, starting at any point on the other side of a Skiba curve, there is also one unique optimal solution, but those solutions are qualitatively different from the first family of solutions. If one starts exactly on the Skiba curve, then there is a member of the first family of solutions and a member of the second that perform equally well, so there are multiple optimal solution trajectories that can be entirely different.

In many models involving a societal "bad", such as drug abuse or terrorism, the Skiba thresholds separate strategies that "eradicate" the problem (if not literally, at least by driving it to a low-level) vs. an "accommodate" approach that still tries to reduce the problem, but more modestly. Loosely speaking that is also what happens here. To the upper right of the Skiba curve, it is optimal to lockdown so forcefully and for so long that relatively few people get infected or die. To the lower left of the Skiba curve, one should accept a higher rate of infection in order to be able to preserve more jobs.

The Skiba curve does not extend all the way across the  $b - M$  plane. For many  $b - M$  combinations, those two strategies are end points on a continuum with a smooth transition between them. But where the Skiba curve exists, there is a discrete point where one should make an abrupt pivot from one strategy to the other.

This is best visualized by replicating Fig. 3's plot of the optimal amount of locking down (vertical axis) vs. the vaccine production capacity ( $b$ , on the horizontal axis), but for a larger value of  $M$ . Fig. 6 does this when  $M = 21,900$ , or 60 times GDP per capita, and so is effectively a vertical slice through the  $b - M$  bifurcation diagram at that value of  $M$ .

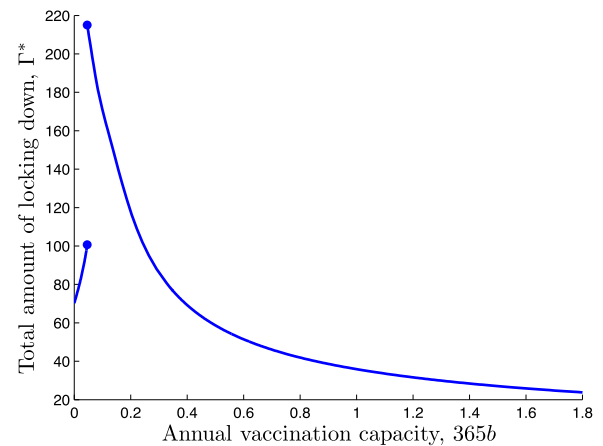


Fig. 6. For a larger value of  $M$ , the optimal total amount of locking down jumps up discretely as vaccination capacity crosses a threshold .

As in Fig. 3, the curve slopes upward for small  $b$ , indicating that vaccines and lockdowns are complements, and the opposite is true for large  $b$ . In Fig. 3 that hump-shaped curve was continuous, but in Fig. 6 it jumps discontinuously at  $365b = 0.0462$  since the point  $(21,900; 0.0462)$  is a Skiba point on the blue curve in the bifurcation diagram in Fig. 4.

Exactly at that point, it is optimal either to use a smaller amount of locking down (a little over 100 on the vertical axis) or a much greater amount of locking down (a little over 210). (The units






Fig. 7. Two different solution trajectories in terms of  $\gamma$  (left panel),  $I$  (center panel), and  $S$  (right panel) that are both optimal for the exact same parameter values and initial states.

are person-days of lockdown-induced unemployment measured on a scale such that 365 corresponds to everyone being unemployed for one year.)

Fig. 7 shows these two optimal solution trajectories for  $\gamma$  (left panel),  $I$  (center panel), and  $S$  (right panel). The solid line shows a strong lockdown, few people getting infected and a steady decline in  $S$  as people get vaccinated, but the slope is shallow because  $b$  is small. Few people get infected, but it takes a long time to reduce  $S$  via vaccination.

That solution is on the right side of the Skiba, in the substitution region, so if  $b$  were increased then it would be possible to achieve that relative low number of infections while using less locking down.

The alternate optimal solution (dashed lines) involves a sharp but modest spike in infections that is soon snuffed out by a lockdown that is aggressive but slightly delayed. Those infections reduce the proportion of the population that is susceptible by about 20 percent. Thereafter, the optimal strategy locks down enough to keep infections low, but the intensity of that requisite lockdown is less severe because that initial surge of infections created some partial herd immunity.

That solution is in the complementary region. That means that less (total) lockdown effort would be expended if  $b$  were a bit smaller, presumably because the lockdown would be delayed longer.

That is rather interesting on a substantive level. Past models tended to find that the optimal strategy was either to keep infections low throughout (as in the solid line) or to use lockdowns

only to flatten the curve a bit but still to have most people become infected. This dashed line solution shows something different, which is to let the epidemic run only for a bit before getting serious about pushing infections down to very low levels.

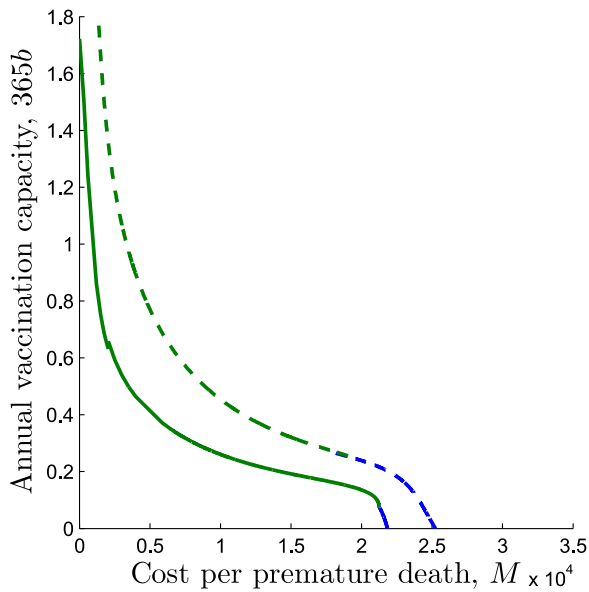
It also adds an interesting wrinkle to the question of whether vaccination and lockdowns are substitutes or complements. Right around the Skiba curve, infinitesimal increases in vaccination capacity can induce very large increases in the optimal total amount of locking down, which is a sort of extreme form of complementarity.

### 6.3. Virus variants

The analysis above was all done with an  $R_0 = 2$ . That value should be understood as the epidemic reproductive number after factoring in containment strategies that do not involve job losses, such as asking knowledge workers to telecommute and requiring masks to be worn.

That value would seem right for the original virus to someone who thinks such "soft" containment measures can only reduce infections by 20% – since 2.0 is 20% below the uncontrolled  $R_0$  estimate of 2.5 the CDC offers for the original virus. It may also seem correct for some more contagious variants if those soft measures are more effective. E.g., if the new variants'  $R_0$  were 4.0 and "soft" measures cut transmission by 50 percent, that would also be well-modeled by our base case  $R_0$ .

But it is also of interest to consider how the results might differ for variants that are considerably more contagious even after soft



**Fig. 8.**  $b - M$  bifurcation diagrams for  $R_0 = 2$  (solid lines) and  $R_0 = 4$  (dashed lines) show a bigger complementarity region when the epidemic's reproductive number is larger.

measures have been factored in. So Fig. 8 contrasts the  $b - M$  bifurcation diagrams with the original  $R_0 = 2$  (solid line) and a higher  $R_0 = 4$  (dashed line). Here we suppress the contour lines in favor of showing regions in  $b - M$  parameter space that correspond to substitution vs. complementarity.

Increasing  $R_0$  shifts the black line to the right, making the complementarity region bigger. The reason is that it is more expensive to contain a more contagious virus via lockdowns, so it takes a larger  $M$ -value to justify prolonged lockdowns when the vaccination rate is low.

The curves are not exactly hyperbolas, but away from the axes they are very roughly approximated by one, and the hyperbola's constant is about twice as great with  $R_0 = 4$  as with  $R_0 = 2$ . So as a very rough rule of thumb, one might say that if the product  $bMR_0$  is greater than some cut-off, then vaccines substitute for lockdowns, but below that threshold greater vaccination capacity should induce wider use of lockdowns.

## 7. Discussion and conclusions

### 7.1. Overview of results

There are at least three broad types of responses to a pandemic such as COVID-19. One is masking and other relatively painless forms of social distancing that do not seriously disrupt the economy. We assume such measures are always in place, and so consider here a baseline epidemic reproductive number  $R$  of 2.0, which is somewhat lower than the totally uncontrolled  $R_0$  for the original virus.

The second is locking down parts of the economy, despite the high cost of doing so in terms of lost jobs, income, and economic well-being. The third is vaccination – which is highly desirable but takes time to invent and deploy.

We start our optimization model at the moment when vaccines are first approved. We presume it is optimal to vaccinate people as rapidly as possible, since vaccines are so very cheap compared to death and unemployment, and investigate how lockdowns should be adjusted over time after vaccinations start, for lesser and greater rates of vaccine production (governed by the parameter  $b$ ).

We described circumstances in which increased vaccination capacity should substitute for lockdowns, effectively meaning greater vaccination capacity buys reduced unemployment. Within this simple model, that would be the case for countries that both highly value preventing COVID-19 deaths and can vaccinate their population quickly (within a year or two). However, when those two parameters are lower, increased vaccination capacity could stimulate greater not less locking down (complementarity), in which case greater vaccination capacity is used to “buy” better health outcomes. There are even so-called Skiba points where an infinitesimal increase in  $b$  can trigger a very large increase in the amount of locking down that is optimal.

We also identified parameters for which a rather interesting strategy is optimal, namely letting about 20 percent of the population get infected before implementing lockdowns severe enough to keep infections quite low for the remaining years until vaccinations confer herd immunity.

The model is sufficiently stylized that we do not claim it shows that any actual country should necessarily follow this prescription. That simplicity is actually a virtue, however, with respect to our central qualitative conclusion. If a highly complicated model produces the result “It depends on the parameters” one can be forgiven for wondering whether it is only arcane details of the model that produce that ambiguity. However, when a simple model shows complicated behavior, that increases the likelihood that the complexity emanates from something quite fundamental in the nature of the problem at hand.

For COVID-19, this intermediate stage of the COVID-19 pandemic is drawing to a close, so it is too late for policy makers to apply this model directly. However, given that further pandemics are possible looking at this question analytically may pay dividends if it prevents naive or uncritical application of “lessons learned” from COVID-19 to future epidemics. Indeed, the results even caution against unthinking transfer of what worked in one country to another country. E.g., rich countries might be wary of presuming that what was optimal for them will necessarily be optimal for countries facing very different economic constraints.

### 7.2. Strong modeling assumptions that may be limitations

COVID transmission dynamics and the effects of lockdowns on transmission, economic activity, and other social outcomes are all very complicated. The model in this paper necessarily made quite a few significant simplifications in order to remain tractable. We hope that the broad qualitative findings are robust to those simplifications, but we want to explicitly acknowledge here simplifications that could be seen as real limitations.

We do not distinguish between lockdown policy and actual behavior. There is a literature that worries about whether the public may relax its vigilance over time, creating gaps between policy and actual behavior (e.g., Gozzi, Bajardi, & Perra (2021) and Caulkins et al. (2021)). That complexity is omitted here on the grounds that vaccine rollout is a relatively short process.

We are modeling an intermediate stage that begins when vaccines are approved and ends when lockdowns are no longer needed. The terminal conditions at which point societies are willing to end lockdowns seem to vary from country to country, sometimes depending on which political party or politician is in power. To abstract away from such complexity, as well as the transition to endemic conditions with an ongoing sequence of virus variants, we assume that enough people are willing to be vaccinated that broad lockdowns will no longer be employed.

We look forward to a time when more empirical work has been done on the effectiveness of lockdowns at reducing transmission rates, so the parameters of Eq. 2 can have a firmer basis.

The model considers only one type of virus. There is no modeling of multiple competing virus variants, loss of vaccine effectiveness because of mutations, etc. Furthermore, it is assumed that there are no re-infections.

The model considers only one type of person. It does not distinguish by age, sex, occupation, or any other category besides disease state.

The model considers one society. There is no migration, travel across international borders, or macroeconomic consequences from borrowing, international trade, or other economic interactions with other countries or regions.

The model does not consider distributional issues. Lockdowns reduce employment of working age people to prevent deaths that would otherwise fall (mostly) on retirees. Hence, lockdowns create a massive reallocation of welfare from younger to older citizens, something that is not considered here.

Elementary SIR models assume random mixing, but there is heterogeneity across people in how many social contacts they have. Because of stochastic selectivity, the proportion of social interactions that are negated by past infections will be greater than the proportion of people who have been infected, and that is not reflected here.

Infection by asymptomatic individuals would be better captured by an SEIR not an SIR model, although in earlier modeling of societal-level lockdown strategy SIR and SEIR models produced very similar results. That does not mean asymptomatic transmission is unimportant, only that in a national model with extended time horizons its effects can reasonably be reflected via a suitable choice of the reproductive number.

In reality, vaccine production capacity increases over time, whereas in our simple model we left it as a fixed constant.

### 7.3. Closing comments

The COVID-19 pandemic is moving beyond the period addressed here, which is the time right after vaccines are first approved. However, there may be future pandemics, and the devastation wrought by COVID-19 suggests society would do well to invest now in creating planning models for the next pandemic. In that regard, we offer a few closing comments.

At least within this model, very different lockdown policies could be optimal for different sets of parameter values that all seem plausible, at least for some countries. That suggests people may want to be humble in their certainty as to what strategy is optimal and to invest more in trying to pin down those parameter values, including the societal value of preventing a premature death from COVID-19. It is not clear that standard values for preventing deaths are necessarily the best for this specific situation.

Also, some parameters are at least partially controllable within a higher-level meta-optimization. It is clear that larger values of  $b$  can be beneficial, and even billion-dollar vaccine production facilities are cheap compared to the costs of COVID-19 deaths and/or the costs of locking down the economy. Nations spend a great deal maintaining militaries during peacetime as precaution and prevention measure. It may be equally prudent for nations to spend substantial sums in vaccine research and in maintaining idle vaccine production capacity even though pandemics come along only every few decades.

### Appendix A. Derivations

Hamiltonian:

$$\mathcal{H} = V_I(W, \gamma) + V_h(I) + V_u(u) + \lambda_S \left( \nu N - \beta(\gamma) \frac{SI}{N} - \mu S - b \frac{S}{S+R_1} \right)$$

$$+ \lambda_I \left( \beta(\gamma) \frac{SI}{N} - (\mu + \mu_I + \alpha) I \right) + \lambda_{R_1} \left( \alpha I - \mu R_1 - b \frac{R_1}{S+R_1} \right) + \lambda_{R_2} (-\mu R_2 + b) + \lambda_\gamma u \tag{14}$$

First order condition:

$$\frac{\partial \mathcal{H}}{\partial u} = \frac{\partial V_u(u)}{\partial u} + \lambda_\gamma \tag{15}$$

Adjoint equations:

$$\begin{aligned} \dot{\lambda}_S &= \mu \lambda_S - \frac{\partial V_I(W, \gamma)}{\partial W} - (\lambda_I - \lambda_S) \beta(\gamma) \frac{I}{N} - (\lambda_{R_1} - \lambda_S) \frac{b R_1}{(S+R_1)^2} \\ \dot{\lambda}_I &= (\mu + \mu_I + \alpha) \lambda_I - \frac{\partial V_h(I)}{\partial I} - (\lambda_I - \lambda_S) \beta(\gamma) \frac{S}{N} - \alpha \lambda_{R_1} \\ \dot{\lambda}_{R_1} &= \mu \lambda_{R_1} - \frac{\partial V_I(W, \gamma)}{\partial W} - (\lambda_S - \lambda_{R_1}) \frac{b S}{(S+R_1)^2} \\ \dot{\lambda}_{R_2} &= \mu \lambda_{R_2} - \frac{\partial V_I(W, \gamma)}{\partial W} \\ \dot{\lambda}_\gamma &= - \frac{\partial V_I(W, \gamma)}{\partial \gamma} - (\lambda_I - \lambda_S) \frac{\partial \beta(\gamma)}{\partial \gamma} \frac{SI}{N} \end{aligned} \tag{16}$$

If we set  $\mu = 0$  this implies (together with the transversality conditions, i.e.  $\lambda_x(T) = \frac{\partial X(T)}{\partial x}$ )

$$\begin{aligned} \lambda_S(t) &= \lambda_S(T) + \int_t^T \left( \frac{\partial V_I(W, \gamma)}{\partial W} + (\lambda_I - \lambda_S) \beta(\gamma) \frac{I}{N} \right. \\ &\quad \left. + (\lambda_{R_1} - \lambda_S) \frac{b R_1}{(S+R_1)^2} \right) ds \\ &= \lambda_S(T) + \int_t^T e^{\int_s^t -\beta(\gamma) \frac{I}{N} - \frac{b R_1}{(S+R_1)^2} ds'} \left( \frac{\partial V_I(W, \gamma)}{\partial W} + \lambda_I \beta(\gamma) \frac{I}{N} \right. \\ &\quad \left. + \lambda_{R_1} \frac{b R_1}{(S+R_1)^2} \right) ds \\ \lambda_I(t) &= \int_t^T \left( \frac{\partial V_h(I)}{\partial I} + (\lambda_I - \lambda_S) \beta(\gamma) \frac{S}{N} + \alpha (\lambda_{R_1} - \lambda_I) - \mu_I \lambda_I \right) ds \\ &= \int_t^T e^{\int_s^t \beta(\gamma) \frac{S}{N} ds' - (\alpha + \mu_I)(s-t)} \left( \frac{\partial V_h(I)}{\partial I} - \lambda_S \beta(\gamma) \frac{S}{N} + \alpha \lambda_{R_1} \right) ds \\ \lambda_{R_1}(t) &= \lambda_S(T) + \int_t^T \left( \frac{\partial V_I(W, \gamma)}{\partial W} + (\lambda_S - \lambda_{R_1}) \frac{b S}{(S+R_1)^2} \right) ds \\ &= \lambda_S(T) + \int_t^T e^{\int_s^t \frac{b S}{(S+R_1)^2} ds'} \left( \frac{\partial V_I(W, \gamma)}{\partial W} + \lambda_S \frac{b S}{(S+R_1)^2} \right) ds \\ \lambda_{R_2}(t) &= \lambda_S(T) + \int_t^T \frac{\partial V_I(W, \gamma)}{\partial W} ds \\ \lambda_\gamma(t) &= \lambda_\gamma(T) + \int_t^T \left( \frac{\partial V_I(W, \gamma)}{\partial \gamma} + (\lambda_I - \lambda_S) \frac{\partial \beta(\gamma)}{\partial \gamma} \frac{SI}{N} \right) ds \end{aligned} \tag{17}$$

where we used  $\lambda_I(T) = 0$  and  $\lambda_S(T) = \lambda_{R_1}(T) = \lambda_{R_2}(T)$ .

### Appendix B. Understanding the dynamics of the lockdown intensity

In this section we study the non-monotonic behavior of the amount of locking down as shown in Fig. 2. To do so, we focus on the cases in which the lockdown intensity reaches a local maximum or minimum.

From the dynamics of  $\gamma$  and the first-order condition on  $u$  (see Eq. (15) in Section Appendix A), we have

$$\dot{\gamma}(t) = u(t) = -k \lambda_\gamma(t) \text{ with } k > 0. \tag{18}$$

Let define the set of times during which there is lockdown but the lockdown intensity does not change as  $\mathcal{T} = \{t \in (0, T) : \gamma(t) < 1 \text{ and } \dot{\gamma}(t) = 0\}$ . From Eq. (18) we have that  $\lambda_\gamma(\tau) = 0$  or

$$\lambda_\gamma(T) + \int_\tau^T \frac{\partial V_I(W, \gamma)}{\partial \gamma} ds + \int_\tau^T (\lambda_I - \lambda_S) \frac{\partial \beta(\gamma)}{\partial \gamma} \frac{SI}{N} ds = 0, \tag{19}$$

for any  $\tau \in \mathcal{T}$ . The first two terms on the LHS of (19) represent the total reduction in the economic cost (i.e. economic benefit) from relaxing the lockdown from time  $\tau$  until the end of the pandemic  $T$ . From the transversality conditions for  $\gamma$  we have that  $\lambda_\gamma(T) < 0$ . The last integral term on the LHS of (19) is the total cost of marginally increasing the number of infected people from time  $\tau$  until the end of the pandemic.

For the base case parameter values reported in Table 1, which implies a moderate basic reproduction number, the set  $\mathcal{T}$  can have zero, one, or two elements. Fig. 2 shows one case in which  $\mathcal{T}$  has one element (see the solid line) and another case in which  $\mathcal{T}$  has two elements (see the dashed line). The number of elements of the set  $\mathcal{T}$  depends on two exogenous parameters:  $b$  (vaccination rate) and  $M$  (value of a life lost). This is because  $b$  and  $M$  affect the total cost of marginally increasing the number of infected people. In particular, a higher  $b$  reduces the number of susceptible people who can get infected, which diminishes the probability that one infected infects more people and hence also the total cost of marginally increasing the number of infected people. In contrast, a higher  $M$  increases the value associated to each death, increasing the total cost of having infected people. We distinguish the following three cases:

- a) Lockdown relaxation (i.e.  $\mathcal{T} = \{\emptyset\}$ ):
  - a.1) If  $b$  is sufficiently large to prevent infections ( $I \approx 0$ ), the economic benefit from relaxing the lockdown will always dominate the total cost of increasing the number of infected until the end of the pandemic.
  - a.2) If  $b$  is not large enough to prevent infections ( $I \gg 0$ ) and  $M$  is sufficiently small, it is expected that the economic benefit from relaxing the lockdown will also dominate the total cost of increasing the number of infected until the end of the pandemic. In this scenario the economy is prioritized over health.
- b) Stringent lockdown followed by lockdown relaxation (i.e.  $\mathcal{T} = \{\tau\}$ ):
  - b.1) If  $b$  is not large enough to prevent infections ( $I \gg 0$ ),  $M$  is sufficiently large, and the initial lockdown intensity  $\gamma(0)$  is sufficiently close to one, the total cost of increasing the number of infected until the end of the pandemic will exceed the economic benefit from relaxing the lockdown until time  $\tau \in \mathcal{T}$ . Given the transversality condition for  $\gamma$  (i.e.,  $\lambda_\gamma(T) < 0$ ), it is a necessary that  $\dot{\lambda}_\gamma(\tau) < 0$ , for  $\tau \in \mathcal{T}$ , which implies that the reduction in the economic cost (i.e., economic benefit) from relaxing the lockdown at time  $\tau$  is greater than the future cost of having more infected people at time  $\tau$  or, equivalently,

$$\frac{\partial V_I(W(\tau), \gamma(\tau))}{\partial \gamma(\tau)} + (\lambda_I(\tau) - \lambda_S(\tau)) \frac{\partial \beta(\gamma(\tau)) S(\tau) I(\tau)}{\partial \gamma(\tau) N(\tau)} > 0. \quad (20)$$

Therefore, it becomes optimal to reduce the intensity of the lockdown after time  $\tau$ . Note that since  $\frac{\partial V_I(W(\tau), \gamma(\tau))}{\partial \gamma(\tau)} < 0$ , the reduction in the lockdown at  $\tau$ ,  $\dot{\lambda}_\gamma(\tau) < 0$ , can only occur in the middle of the pandemic when both  $S(\tau)$  and  $I(\tau)$  differ from zero.

- c) Lockdown relaxation-stringent lockdown-lockdown relaxation (i.e.  $\mathcal{T} = \{\tau', \tau\}$ ):
  - c.1) If  $b$  is not large enough to prevent infections ( $I \gg 0$ ),  $M$  is sufficiently large to prioritize health over the economy, and the initial lockdown intensity is sufficiently strong (i.e.  $\gamma(0) \ll 1$ ), the economic benefit from relaxing the lockdown will initially dominate the total cost of increasing the number of infected until the end of the pandemic. However, if this policy leads to an increasing number of infected peo-

ple (due to a low vaccination rate), there will be a  $\tau' \in \mathcal{T}$  for which  $\dot{\lambda}_\gamma(\tau') > 0$ . At time  $\tau'$  the future cost of having more infected people is greater than the reduction in the economic cost (i.e., economic benefit) from relaxing the lockdown or, equivalently,

$$\frac{\partial V_I(W(\tau'), \gamma(\tau'))}{\partial \gamma(\tau')} + (\lambda_I(\tau') - \lambda_S(\tau')) \frac{\partial \beta(\gamma(\tau')) S(\tau') I(\tau')}{\partial \gamma(\tau') N(\tau')} < 0. \quad (21)$$

Hence, it becomes optimal to increase the intensity of the lockdown after time  $\tau' \in \mathcal{T}$ . Notice that since the transversality condition imposes that  $\lambda_\gamma(T) < 0$ , and at  $\tau' \in \mathcal{T}$ ,  $\dot{\lambda}_\gamma(\tau') > 0$ , the *intermediate value theorem* guarantees that there exists an additional  $\tau \in \mathcal{T}$  for which  $0 < \tau' < \tau < T$ . The intermediate value theorem guarantees that after  $\tau'$ , the lockdown intensity will follow case b.1.

### Appendix C. Proofs of results in Section 4

#### Proof of Lemma 2

We find the derivative along the contour line of  $\Gamma^*(b, M)$  from the total derivative. Since the value of  $\Gamma^*(b, M)$  is constant along the contour line we get

$$d\Gamma^*(b, M) = \frac{\partial \Gamma^*(b, M)}{\partial b} db + \frac{\partial \Gamma^*(b, M)}{\partial M} dM = 0. \quad (22)$$

Using the Leibnitz rule for the partial derivative, using the definition of a shadow price and manipulation leads to the assertion of the Lemma.

The derivative along the contour lines of  $\gamma_l^*(b, M)$  and  $\gamma_h^*(b, M)$  is obtained analogously. □

**Proof of Theorem 3.** In the following we distinguish two sets  $U$  and  $S$  defined as

$U := \{(b, M) : \text{the optimal solution is unique at } (b, M)\}$

$S := \{(b, M) : \text{two optimal solutions exist at } (b, M)\}$ .

For  $(b_0, M_0) \in U$  the function  $\Gamma^*(b, M_0)$  is continuously differentiable at  $b_0$  and for  $(b_0, M_0) \in S$  the function  $\Gamma^*(b, M_0)$  is continuously differentiable from below and above at  $b_0$ .

Specifically, the derivative at  $(0, M)$  from above exist for all  $M$ . Thus, we set

$$\bar{M} := \sup_M \left\{ M : \frac{d\Gamma^*(0, M)}{db} \geq 0 \right\}. \quad (23)$$

Additionally we define

$$\underline{b} := \sup_b \left\{ \frac{d\Gamma^*(b, M)}{db^-} > 0 \right\} \quad \text{and} \quad \bar{b} := \inf_b \left\{ \frac{d\Gamma^*(b, M)}{db^+} < 0 \right\}. \quad (24)$$

To get an explicit expression for the derivative at  $(b_0, M_0) \in U$

$$\frac{d\Gamma^*(b_0, M_0)}{db}$$

we consider the total derivative of  $\Gamma^*(b_0, M_0)$ , which equals

$$d\Gamma^*(b_0, M_0) = \frac{\partial \Gamma^*(b_0, M_0)}{\partial b} db + \frac{\partial \Gamma^*(b_0, M_0)}{\partial M} dM \quad (25)$$

we obtain the derivative with respect to the parameter  $b$  by setting  $dM = 0$  (variation in direction of  $M$  is nil), i.e.

$$\begin{aligned} d\Gamma^*(b_0, M_0) &= \frac{\partial \Gamma^*(b_0, M_0)}{\partial b} db \\ &= \left( \frac{\partial T_1}{\partial b} \underbrace{(1 - \gamma^*(T_1))}_{=0} - \int_0^{T_1} \frac{\partial \gamma^*(t)}{\partial b} dt \right) db \end{aligned}$$



$$\begin{aligned}
 &= \left( - \int_0^{T_1} \frac{\partial \gamma^*(t)}{\partial \gamma} \frac{\partial \gamma}{\partial b} dt \right) db \\
 &= \left( - \int_0^{T_1} \frac{\lambda_b}{\lambda_\gamma} dt \right) db. \tag{26}
 \end{aligned}$$

Thus,

$$\frac{d\Gamma^*(b_0, M_0)}{db} = - \int_0^{T_1} \frac{\lambda_b}{\lambda_\gamma} dt. \tag{27}$$

In the following we show that  $\bar{M} > 0$  exists and that  $\frac{d\Gamma^*(b, M)}{db}$  is negative for  $b \rightarrow +\infty$ , which proves the existence of a global maximum of  $\Gamma^*(b, M)$  which is in general non-unique.

Starting with the latter case (i.e.  $b \rightarrow +\infty$ ) we check the fraction of (27). The numerator can be derived analogous to a usual adjoint variable (deriving the adjoint equation and solving backwards). We obtain

$$\lambda_b(t) = \int_t^T \frac{1}{S + R_1} (S(\lambda_{R_2} - \lambda_S) + R_1(\lambda_{R_2} - \lambda_{R_1})) ds. \tag{28}$$

As straightforwardly can be shown  $\lambda_{R_2}(s) < \lambda_S(s)$  and  $\lambda_{R_2}(s) < \lambda_{R_1}(s)$  holds for all  $s$ ,  $\lambda_b(t)$  is negative for all  $t$ . Considering the denominator of the integrand  $\lambda_\gamma(t)$  the analytic expression follows from the canonical system, i.e.

$$\lambda_\gamma(t) = \lambda_\gamma(T) + \int_t^T \left( \frac{\partial V_1(W, \gamma)}{\partial \gamma} + (\lambda_I - \lambda_S) \frac{\partial \beta(\gamma) IS}{\partial \gamma N} \right) ds. \tag{29}$$

From the definition of  $S(X(T))$  and  $V_1(W(t), \gamma(t))$  it follows that  $\lambda_\gamma(T)$  and  $\frac{\partial V_1(W, \gamma)}{\partial \gamma}$  are both negative. For  $b$  tending to infinity  $S(t)$  goes to 0 immediately and thus  $\lim_{b \rightarrow +\infty} \lambda_\gamma(t) < 0$  for all  $t$ . This implies  $\lim_{b \rightarrow +\infty} \frac{d\Gamma^*(b, M)}{db} < 0$  for all  $M$ .

For the case  $b = 0$  we know by the assumption that the optimal solution at  $M = 0$  satisfies  $(\gamma^*(t), u^*(t)) = (1, 0)$  for all  $t$ . Therefore  $\Gamma^*(0, 0)$  is minimal, i.e. equal to zero. Therefore  $\frac{d\Gamma^*(0, 0)}{db} \geq 0$ . This proves that  $\bar{M} > 0$ , defined in (23), exists and that the result holds for all  $M$  with  $0 \leq M < \bar{M}$ .

By the definition (24) of  $\underline{b}$  and  $\bar{b}$  we find that a maximum is achieved at  $b^*$  satisfying  $\underline{b} \leq b^* \leq \bar{b}$  and that  $\underline{b} = b^* = \bar{b}$  if  $b^*$  is unique. □

**Proof of Theorem 5.**

For sufficiently high  $M$  the second term in the integral of (29) dominates the first term after  $t = 0$ , such that  $\lambda_\gamma(t) < 0$ . Following the same steps as in the proof of Theorem 3 proves the assertion. □

**Proof of Corollary 6**

$\dot{\gamma}(t) \geq 0$  for  $t \in [0, T]$  implies  $u(t) \geq 0$  for  $t \in [0, T]$ . From the first order condition (15) or (13) it follows that  $\lambda_\gamma(t) < 0$  whenever  $u(t) > 0$ . From the proof of Theorem 3 we follow  $\lambda_b(t) < 0$ . Then from (26) it follows that  $\frac{d\Gamma^*(b, M)}{db} < 0$  which proves the assertion of the Lemma. □

**Appendix D. Adaptations for the numerical calculations**

To handle “critical” ratio

$$\frac{S}{S + R_1}, \quad \text{with } S, S + R_1 \approx 0$$

we introduce a parameter  $\tau$  such that the state dynamics writes as

$$\dot{S} = \nu N - \beta(\gamma) \frac{SI}{N} - \frac{bS}{S + R_1 + \tau} - \mu S$$

$$\dot{I} = \beta(\gamma) \frac{SI}{N} - (\mu + \mu_I + \alpha) I$$

$$\dot{R}_1 = \alpha I - \frac{bR_1}{S + R_1 + \tau} - \mu R_1$$

$$\dot{R}_2 = b \frac{S + R_1}{S + R_1 + \tau} - \mu R_2.$$

For the calculations we use  $\tau = 10^{-3,4}$

We assume a constant normalized population size  $N(t) = S(t) + I(t) + R_1(t) + R_2(t) = 1$  for all  $t$  with initial conditions

$$S(0) = 1 - I(0) - R_1(0) - R_2(0) = \frac{53}{60}$$

with

$$I(0) = \frac{1}{60}$$

$$R_1(0) = 0.1, R_2(0) = 0$$

and

$$\gamma(0) = 0.8.$$

**References**

Aborode, A. T., David, K. B., Uwishema, O., Nathaniel, A. L., Imisioluwa, J. O., Onigbinde, S. B., & Farooq, F. (2020). Fighting COVID-19 at the expense of malaria in africa: The consequences and policy options. *The American Journal of Tropical Medicine and Hygiene*, 104(1), 26–29.

Acemoglu, D., Chernozhukov, V., Werning, I., & Whinston, M. D. (2020). A multi-risk SIR model with optimally targeted lockdown, volume 2020. In *National bureau of economic research cambridge, MA.*

Acemoglu, D., Chernozhukov, V., Werning, I., & Whinston, M. D. (2021). A multi-risk SIR model with optimally targeted lockdown. *American Economic Review: Insights*, 3, 487–502.

Adam, D. (2020). A guide to r—the pandemic’s misunderstood metric. *Nature*, 583(7816), 346–348.

Alfaro, L., Faia, E., Lammersdorf, N., & Saidi, F. (2020). Social interactions in pandemics: Fear, altruism, and reciprocity. In *Technical report*. National Bureau of Economic Research.

Alvarez, F. E., Argente, D., & Lippi, F. (2021). A simple planning problem for COVID-19 lockdown. *American Economic Review: Insights*, 3, 367–382.

Arolas, H. P., Acosta, E., López-Casasnovas, G., Lo, A., Nicodemo, C., Riffe, T., & Myrskylä, M. (2021). Years of life lost to COVID-19 in 81 countries. *Scientific Reports*, 11(1), 1–6.

Aspri, A., Beretta, E., Gandolfi, A., & Wasmer, E. (2021). Mortality containment vs. economics opening: Optimal policies in a SEIARD model. *Journal of Mathematical Economics*, 93, 102490.

Auld, C., & Toxvaerd, F. (2021). *The great COVID-19 vaccine rollout: Behavioral and policy responses*. CEPR Discussion Paper No. DP16070.

Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science*, 15, 215–227. Control strategies in sir epidemic models. *Mathematical Biosciences*, 292, 86–96

Bosi, S., Camacho, C., & Desmarchelier, D. (2021). Optimal lockdown in altruistic economies. *Journal of Mathematical Economics*, 93, 102488.

Buratto, A., Muttoni, M., Wrzaczek, S., & Freiburger, M. (2022). Should the COVID-19 lockdown be relaxed or intensified in case a vaccine becomes available? *PLoS ONE*, 17(9), E0273557

Caulkins, J. P., Grass, D., Feichtinger, G., Hartl, R., Kort, P. M., Prskawetz, A., ... Wrzaczek, S. (2020). How long should the COVID-19 lockdown continue? *PLoS ONE*, 15(2), E0243413

Caulkins, J. P., Grass, D., Feichtinger, G., Hartl, R., Kort, P. M., Prskawetz, A., ... Wrzaczek, S. (2021). The optimal lockdown intensity for COVID-19. *Journal of Mathematical Economics*, 93, 102489.

Caulkins, J. P., Grass, D., Feichtinger, G., Hartl, R. F., Kort, P. M., Prskawetz, A., ... Wrzaczek, S. (2022). Covid-19 and optimal lockdown strategies: The effect of new and more virulent strains. In *In pandemics: Insurance and social protection* (pp. 163–190). Cham: Springer.

Charpentier, A., Elie, R., Laurière, M., & Tran, V. C. (2020). COVID-19 Pandemic control: Balancing detection policy and lockdown intervention under ICU sustainability. *Mathematical Modelling of Natural Phenomena*, 15, 57.

Eichenbaum, M. S., Rebelo, S., & Trabandt, M. (2022). The macroeconomics of testing and quarantining journal of economic dynamics and control. *Elsevier*, 2022(138), 104337.

Farboodi, M., Jarosch, G., & Shimer, R. (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, 196, 105293.

Federico, S., & Ferrari, G. (2021). Taming the spread of an epidemic by lockdown policies. *Journal of Mathematical Economics*, 93, 102453.

Federico, S., Ferrari, G., & Torrente, M. (2022). Optimal vaccination in a SIRS epidemic model. *arXiv:2206.03284*.

<sup>4</sup> A smaller value of  $\tau$  does not change the numerical results. The chosen value is more convenient for the calculations than e.g.  $\tau = 10^{-5}$ .

- Freiberger, M., Grass, D., Kuhn, M., Seidl, A., & Wrzaczek, S. (2022). Chasing up and locking down the virus: Optimal pandemic interventions within a network. *Journal of Public Economic Theory*, 24, 1182–1217.
- Fu, Y., Jin, H., Xiang, H., & Wang, N. (2022). Optimal lockdown policy for vaccination during COVID-19 pandemic. *Finance Research Letters*, 45, 102123.
- Fu, Y., Xiang, H., Jin, H., & Wang, N. (2021). Mathematical modelling of lockdown policy for covid-19. *Procedia Computer Science*, 187, 447–457.
- Garriga, C., Manuelli, R., & Sanghi, S. (2022). Optimal management of an epidemic: Lockdown, vaccine and value of life. *Journal of Economic Dynamics and Control*, 140, 104351.
- Giagheddu, M., & Papetti, A. (2023). The macroeconomics of age-varying epidemics. *European Economic Review*, 151, 104346.
- Goldstein, J. R., & Lee, R. D. (2020). Demographic perspectives on the mortality of COVID-19 and other epidemics. *Proceedings of the National Academy of Sciences*, 117(36), 22035–22041.
- Gonzalez-Eiras, M., & Niepelt, D. (2020). On the optimal 'lockdown' during an epidemic. *CEPR discussion paper 14612*. Swiss National Bank, Study Center Gerzensee.
- Gozzi, M., Bajardi, P., & Perra, N. (2021). The importance of non-pharmaceutical interventions during the COVID-19 vaccine rollout. *medRxiv*.
- Grass, D., Caulkins, J. P., Feichtinger, G., Tragler, G., & Behrens, D. (2008). *Optimal control of nonlinear processes: With applications in drugs, corruptions, and terror*. Berlin: Springer-Verlag.
- Hammit, J. K. (2020). *Valuing mortality risk in the time of COVID-19*. SSRN 3615314.
- Hannah, R., Ortiz-Ospina, E., Beltekian, D., Mathieu, J. H., Macdonald, B., Giattino, C., Appel, C., Rodes-Guirao, L., & Roser, M. (2020). Coronavirus pandemic (COVID-19). Published online at OurWorldInData.org. Retrieved from: <https://ourworldindata.org/coronavirus>.
- Hsu, W.-T., Lin, H.-C., & Yang, H. (2020). *Between lives and economy: Optimal COVID-19 containment policy in open economies*. Available at SSRN 3705800.
- Huberts, N. F., & Thijssen, J. J. (2023). Optimal timing of non-pharmaceutical interventions during an epidemic. *European Journal of Operational Research*, 305(3), 1366–1389.
- International Labour Organization (2021). COVID-19 and the world of work, 7<sup>th</sup> edition. [https://www.ilo.org/wcmsp5/groups/public/@dgreports/@dcomm/documents/briefingnote/wcms\\_767028.pdf](https://www.ilo.org/wcmsp5/groups/public/@dgreports/@dcomm/documents/briefingnote/wcms_767028.pdf).
- Jackson, K. J., Weiss, A. M., Schwarzenberg, B. A., & Nelson, M. R. (2021). *Global economic effects of COVID-19*. Congressional Research Service.
- Jones, C., Philippon, T., & Venkateswaran, V. (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies*, 34(11), 5188–5223.
- Kaplan, G., Moll, B., & Violante, G. L. (2020). *The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the us*. Technical Report, National Bureau of Economic Research.
- Kermack, W. O., & McKendrick, A. G. (1927). Contributions to the mathematical theory of epidemics - i. *Proceedings of the Royal Society*, 115(772), 700–721.
- Kniesner, T. J., Kip Viscusi, W., Woock, C., & Ziliak, J. P. (2012). The value of a statistical life: Evidence from panel data. *The Review of Economics and Statistics*, 94(1), 74–87.
- Libotte, G. B., Lobato, F. S., Platt, G. M., & Silva Neto, A. J. (2020). Determination of an optimal control strategy for vaccine administration in COVID-19 pandemic treatment. *Computer Methods and Programs in Biomedicine*, 196, 105664.
- Mak, H.-Y., Tiglong, D., & Tang, S. T. (2022). Managing two-dose COVID-19 vaccine rollouts with limited supply. *Production and Operations Management*, 31, 4424–4442.
- Makris, M. (2021). Covid and social distancing with a heterogenous population. *Economic Theory*, 1–50.
- Pigullem, F., & Shi, L. (2022). Optimal covid-19 quarantine and testing policies. *The Economic Journal*, 132(647), 2534–2562.
- Rao, I. J., & Brandeau, M. L. (2021). Optimal allocation of limited vaccine to control an infectious disease: Simple analytical conditions. *Mathematical Biosciences*, 337, 108621.
- Rao, I. J., & Brandeau, M. L. (2022). Sequential allocation of vaccine to control an infectious disease. *Mathematical Biosciences*, 351, 108879.
- Tsai, T. C., Jacobson, B. H., & Jha, A. K. (2020). American hospital capacity and projected need for COVID-19 patient care. *Health Affairs Blog*.
- World Health Organization (2022). Coronavirus (COVID-19) dashboard. <https://covid19.who.int/>.