TECHNISCHE UNIVERSITÄT WIEN

Operations Research and Control Systems

SWM ORCOS

The optimal lockdown intensity for COVID-19

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

Research Report 2020-11

August 2020

ISSN 2521-313X

Operations Research and Control Systems
Institute of Statistics and Mathematical Methods in Economics
Vienna University of Technology

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

The optimal lockdown intensity for COVID-19

Jonathan P. Caulkins[a], Dieter Grass[b], Gustav Feichtinger[c,d], Richard F. Hartl[e], Peter M. Kort[f,g], Alexia Prskawetz[h,d], Andrea Seidl[e], Stefan Wrzaczek[d]

[a] *Heinz College, Carnegie Mellon University, Pittsburgh, USA*
[b] *International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria*
[c] *Department for Operations Research and Control Systems, Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria*
[d] *Wittgenstein Centre for Demography and Global Human Capital (IIASA, OeAW, University of Vienna)*
[e] *Department of Business Decisions and Analytics, University of Vienna, Vienna, Austria*
[f] *Tilburg School of Economics and Management, Tilburg University, Tilburg, Netherlands*
[g] *Department of Economics, University of Antwerp, Antwerp, Belgium*
[h] *Research Group Economics, Institute of Statistics and Mathematical Methods in Economics, Vienna University of Technology, Vienna, Austria*

Abstract

One of the principal ways nations are responding to the COVID-19 pandemic is by locking down portions of their economies to reduce infectious spread. This is expensive in terms of lost jobs, lost economic productivity, and lost freedoms. So it is of interest to ask: What is the optimal intensity with which to lockdown, and how should that intensity vary dynamically over the course of an epidemic? This paper explores such questions with an optimal control model that recognizes the particular risks when infection rates surge beyond the healthcare system's capacity to deliver appropriate care. The analysis shows that four broad strategies can be optimal, ranging from brief lockdowns that only "smooth the curve" to sustained lockdowns that prevent infections from spiking beyond the healthcare system's capacity. Within this model, it can be optimal to have two separate periods of locking down, so returning to a lockdown after initial restrictions have been lifted is not necessarily a sign of failure. Relatively small changes in judgments about how to balance health and economic harms can alter dramatically which strategy is optimal. Indeed, there are constellations of parameters for which two or even three of these distinct strategies can all be optimal for the same set of initial conditions; these correspond to so-called triple Skiba points. The performance of trajectories can be highly nonlinear in the state variables, such that for various times t, the optimal unemployment rate could be low, medium, or high, but not anywhere in between. These complex dynamics emerge naturally

Email addresses: caulkins@andrew.cmu.edu (Jonathan P. Caulkins), dieter.grass@tuwien.ac.at (Dieter Grass), gustav@eos.tuwien.ac.at (Gustav Feichtinger), richard.hartl@univie.ac.at (Richard F. Hartl), Kort@uvt.nl (Peter M. Kort), afp@econ.tuwien.ac.at (Alexia Prskawetz), andrea.seidl@univie.ac.at (Andrea Seidl), stefan.wrzaczek@oeaw.ac.at (Stefan Wrzaczek)

from modeling the COVID-19 epidemic and suggest a degree of humility in policy debates. Even people who share a common understanding of the problem's economics and epidemiology can prefer dramatically different policies. Conversely, favoring very different policies is not evidence that there are fundamental disagreements.

Keywords: COVID-19, Lockdown, Skiba threshold, SIR model, optimal control

JEL codes: C61, I15

1. Introduction

A central strategy for responding to the COVID-19 pandemic is "locking down" parts of the economy to reduce social interaction and, hence, contagious transmission. Multiple countries have started aggressively, locking down all but essential services such as healthcare and public safety, and then gradually re-opened increasing shares of the economy. Some have then seen infection rates rebound and returned to a more stringent lockdown. Some places have also seen such widespread infection that a nontrivial proportion of the population has passed through infection to reach a "recovered state", although there is uncertainty as to whether the resulting immunity is brief (as with seasonal flu) or long-lasting (as with chicken pox).

All of these considerations raise the challenging question of what is the optimal degree to which a country should lock down, and how that intensity should vary as the state of the epidemic evolves. We try to address that problem with an optimal control model. The heart of the model is a classic SIR or Susceptible-Infected-Recovered differential equation model, but it is enhanced in multiple ways. For example, the lethality of the infection varies depending on whether there are so many infections that critical care capacity has been swamped. The most fundamental extension, though, is creating an objective function that balances four considerations: (1) Health harms (primarily COVID-related deaths), (2) Economic harm (primarily from unemployment), and (3) Adjustment costs, meaning that sharp oscillations in the intensity of the lockdown are costly because it becomes hard for people and businesses to constantly have to adapt to changing rules.

Although businesses can be shut down quickly, re-opening is not as easy; policy makers cannot just order by fiat all businesses to return to their previous levels of employment. So the level of employment or economic activity is treated as a state variable, and the control is adjustments to that level, with asymmetric costs reflecting that it is easier to destroy than to create jobs. Another innovation is that public discontent with the duration and intensity of the lockdown is represented by a fifth state variable that can enter the objective function directly and also modulates compliance with social distancing demands and, hence, the rate of infection.

The solutions are complex and span a range of qualitatively different strategies, such as locking down sufficiently long and forcefully to drive infection rates down to low levels and, at the other extreme, locking down only sparingly to merely soften the peak of infections, without truly sparing most of the public from infection. Which strategy wins – in the sense of delivering the lowest overall total cost – depends on the various parameter values in predictable ways, but there are constellations of parameters for which two qualitatively different strategies may perform equally well, even though they are very different. These tipping points have been variously called Skiba, Sethi-Skiba, DNS, and DNSS points to celebrate the contributions of various pioneers in the field.

Interestingly in this model there are not only conventional Skiba points separating two alternate optimal strategies, but also "triple Skiba points" separating three different equally appealing strategies, and even instances in which there are multiple triple Skiba points in the same bifurcation diagram.

Importantly, there are Skiba thresholds depending on parameters that are either not known scientifically or that reflect value judgments (such as how to trade off saving lives with creating jobs). Hence, one meta-message of this analysis is that when two countries or two people favor sharply different policies, that does not imply that they must have sharply different understandings of the disease, its contagious spread, or even the extent of economic dislocation lockdowns create. Preferences for sharply different policies does not imply there need be sharp disagreements. Conversely, a degree of humility and generosity may be appropriate when talking with people who favor very different policies.

This even extends to the number of lockdowns. There are optimal solutions that involve locking down, ending the lockdown and reinstituting it. Hence, if a country endures a second lockdown, that cannot be taken as proof that the first lockdown "failed", or that policymakers made mistakes.

There is now a growing literature on COVID-19 and its economic consequences as it is related to extended periods of economic lockdown. So far, only a few papers have investigated the optimal timing, length and extent of the lockdown itself. Starting from the simple epidemiological SIR model, Gonzalez-Eiras and Niepelt (2020) investigate the optimal lockdown intensity and duration taking into account the tradeoff between health and economic consequences of the lockdown itself. Alvarez et al. (2020) similarly employ a standard SIR model where they control the fraction of the population going into lockdown. The model is derived with and without testing as a control variable. If testing is included, the optimal lockdown in the US should be started one week after the outbreak of the virus and relaxed after one month. The absence of testing shortens the optimal length of the lockdown, which is due to the dynamics of the epidemiology, i.e. the fraction of recovered people over time increases, implying that the efficiency of the lockdown decreases since also recovered people are locked down.

Köhler et al. (2020) analyze the impact of measures like social distancing which reduce the infection rate. The paper distinguishes between different groups of infected, and assumes that the mortality rate depends on the capacity of the health system. The objective is to minimize the number of fatalities, but the authors take the societal and economic costs of the policy measures into account by means of requiring these costs not to exceed the costs of some baseline policy. To handle uncertainties, they promote a model predictive control based feedback strategy where the policy measures are updated at discrete points in time.

Acemoglu et al. (2020) allow the intensity of lockdown to differ by different age-groups, distinguishing between "young", "middle-aged" and "old" populations, in a SIR model. It can be shown that differentiated policy measures significantly outperform optimal uniform policies. The gains can be realized by having stricter policies on the oldest age-group. Aspri et al. (2020) extends a SEIRD model, where the population is divided into susceptibles, exposed, infected, recovered and deceased, by an asymptomatic population compartment and obtain multiple lockdowns as well as Skiba points.

We expand these previous attempts to model the optimal strategy of lockdowns by three novel features. First, in the balance of economic and health consequences during a lockdown, we consider a capacity constraint in the intensive care units of the health care system. If the number of infected needing intensive care exceeds the constraint the death rate of these patients increases. Second, we explicitly model the "memory of lockdowns" by an additional state variable that accumulates the intensity and length of the lockdown. This "memory of lockdowns" affects the efficiency of the lockdown and accounts for the fact that people get fed up from long and drastic lockdowns. Third, we assume that adjusting the lockdown is costly. In particular, we allow for an asymmetry in the costs for strengthening and weakening the lockdown.

There is likewise a celebrated history of papers exploring Skiba thresholds (see Grass et al. (2008), Sethi (2019)). Our paper belongs to this stream of literature, because in different scenarios Skiba points occur. Comparable bifurcation analyses as in our paper can be found in Grass (2012) and Kiseleva and Wagener (2010, 2015). However, in addition to these contributions, we found triple Skiba points and even multiple triple Skiba points for specific parameter constellations. The first triple Skiba point was found when solving the two-state intensity splitting production/inventory model in Steindl and Feichtinger (2004). Zeiler et al. (2011) is another example where a solution with a triple Skiba point occurs. However, both of these contributions consider optimal control models with infinite time horizon, whereas in our framework the time horizon is finite. In that sense the model of Caulkins et al. (2015) is more related, but there just Skiba points in the usual sense, i.e. separating "only" two different solutions with equal objective value, occur.

We proceed by introducing the model. Section 3 presents the numerical results for the base case parameters and provides an in-depth discussion of the implications of triple Skiba points. In Section 4 the results are discussed and Section 5 concludes.

2. The Model

2.1. Lockdowns

A lockdown reduces interaction among people by closing down businesses and restricting social interaction (e.g., preventing families from visiting loved ones in nursing homes). We do not distinguish between business-related and non-business restrictions and so effectively assume that they move together. If the rate of infection and other factors point to severe [mild] restrictions on business, then one would expect greater [lesser] restrictions on personal social interactions.

We define $\gamma(t)$ to be the actual number of people working as a proportion of those who would normally be working, so apart from COVID we would have $\gamma(t) = 1$. As soon as the lockdown starts, $\gamma(t)$ will drop below 1, which hurts the economy, but reduces social interactions and, hence, the rate of new infections, in a manner described below.

Note that $\gamma(t)$ is modeled as a state variable, not a control, for three reasons. First, outside of a command-and-control state-run economy, policy makers do not get to choose directly the level of employment. Second, adjusting the level of employment takes time and is costly. If a country that has shut down its auto manufacturing supply chain permits that supply chain to reopen, it will take time to reestablish connections (e.g., because some suppliers may have gone bankrupt) and could even require some sort of fiscal stimulus to "prime the pump" in the Keynesian sense of the term. We allow these costs to be asymmetric; it may well be easier to shut down industries than to restart them.

Third, and related, the final value of $\gamma(t)$ at the model's terminal time T (when a vaccine renders lockdowns moot) enters into the salvage value function. The reasons is that if two solutions rack up identical costs over the time period $(0, T)$ but one reaches time T with its economy intact (i.e., $\gamma(T)$ is close to 1) and the other reaches time T in the midst of a deep recession ($\gamma(T)$ well below 1), then the first solution should be preferred. This salvage function reflects the hang-over effect of economic damage that extends beyond the period when the infection's dynamics are relevant. If $\gamma(t)$ and, hence, $\gamma(T)$, were a control variable, then the optimal solution would always choose to discontinuously jump $\gamma(t)$ to 1 at time T to magically make the long-run costs of the lockdown-induced dislocation disappear.

Hence, we let the change in the employment ratio $u(t)$ be a control variable that has adjustment costs, and add a state equation

$$\dot{\gamma}(t) = u(t), \quad \gamma(0) = 1,$$

which reflects a pre-COVID situation with $\gamma(0) = 1$.

We include a state constraint that

$$\gamma(t) \leq 1, \quad 0 \leq t \leq T,$$

since an economy having more than 100% employment makes no sense.

2.2. Lockdown fatigue

People are not robots, and the effectiveness of policies restricting activities depends, in part, on the public's level of cooperation and their dedication to public health protocols. A country could restrict restaurants to take-out service, but if the kitchen workers refuse to wear masks, wash hands frequently, or maintain social distancing during break times then some of the potential benefits will not be realized.

Our sense is that in many jurisdictions the public's tolerance for restrictions begins to wane the more restrictive is the lockdown, and the longer it lasts. So the lockdown's effect on virus transmission depends not only on the instantaneous value of $\gamma(t)$, but also on some accumulated memory of how burdensome the lockdown has been up until time t.

The state variable $z(t)$ captures this "lockdown fatigue" through a standard accumulation stock dynamic that is driven by the rate of COVID-induced unemployment. Since $\gamma(t)$ measures the proportion who are employed, $1 - \gamma(t)$ is the proportion who are unemployed. Hence,

$$\dot{z}(t) = \kappa_1 (1 - \gamma(t)) - \kappa_2 z(t).$$

where κ_1 governs the rate of accumulation of fatigue and κ_2 measures its rate of exponential decay. Note that if the worst imaginable lockdown ($\gamma(t) = 0$) lasted forever then z would grow to its maximum possible value of $z_{\max} = \kappa_1/\kappa_2$.

The foundation of our epidemic model is the standard SIR or Susceptible-Infected-Recovered structure. In it, new infections are proportional to the number of susceptible people, the proportion of people they meet who are infectious, and a proportionality factor $\beta(t)$, which encompasses both the number of interactions and the likelihood that an interaction produces an infection. Numbers of interactions can be reduced by shutting down business and by adaptations on the consumer side; e.g., only going to the grocery store once every two weeks instead of every week. The likelihood of infection given an interaction is affected by things like mask wearing, hand washing, and remaining at least two meters apart during an interaction.

The function $\beta(z(t), \gamma(t))$ should be convex in $\gamma(t)$ because the first businesses that are closed are the ones whose activities generate the most infections per unit of employment or economic value. E.g., a society could be expected to first forbid concerts and other large public gatherings, then socializing in bars and dine-in restaurants, and then, if the need is great enough, to shut down manufacturing, construction, and other non-essential workplaces that do not involve direct interaction with the public.

$$\beta := \beta(z(t), \gamma(t)), \quad \beta_\gamma > 0, \quad \beta_{\gamma\gamma} \geq 0, \quad \beta_z > 0, \quad \beta(1, 0) = \bar{\beta},$$

where $\bar{\beta}$ stands for the rate of social interaction in pre-COVID times.

In the absence of lockdown fatigue, we might model β as some minimum level of infection risk β_1 that is produced just by essential activities (providing healthcare, food, and emergency services) plus an increment β_2 that is proportional to $\gamma(t)$ raised to an exponent θ that is greater than one to achieve the convexity.

We model the dependence of β on $z(t)$ and $\gamma(t)$ as follows:

$$\beta(\gamma(t), z(t)) := \beta_1 + \beta_2 \left(\gamma(t)^\theta + f \frac{\kappa_2}{\kappa_1} z(t)(1 - \gamma(t)^\theta) \right).$$

For its properties see Appendix A.

This expression can be interpreted as follows. The term $\frac{\kappa_2}{\kappa_1} z(t)$ is the lockdown fatigue expressed as percentage of its maximum possible value. So if $f = 1$ and $z(t)$ reached its maximum value, then all of the potential benefits of locking down and pushing $\gamma(t)$ below 1.0 would be negated. In reality, the lockdown fatigue will not reach its maximum and we choose a relatively small value of $f = 0.05$, so this attenuation of the lockdowns benefit by lockdown fatigue has a quite modest force in the analysis below. Nonetheless, we believe it is important to at least acknowledge this human dimension of how a population responds to extended lockdowns.

2.3. State dynamics

The state dynamics can then be written as

$$\dot{S}(t) = \nu N(t) - \beta(\gamma(t))\frac{S(t)I(t)}{N(t)} - \mu S(t) + \varphi R(t) \tag{1a}$$

$$\dot{I}(t) = \beta(\gamma(t))\frac{S(t)I(t)}{N(t)} - (\alpha + \mu + \mu_I)I(t) \tag{1b}$$

$$\dot{R}(t) = \alpha I(t) - \mu R(t) - \varphi R(t) \tag{1c}$$

$$\dot{\gamma}(t) = u(t), \quad \gamma(0) = 1 \tag{1d}$$

$$\dot{z}(t) = \kappa_1(1 - \gamma(t)) - \kappa_2 z(t), \quad z(0) = 0 \tag{1e}$$

$$\gamma(t) \leq 1, \quad 0 \leq t \leq T \tag{1f}$$

$$\beta(\gamma(t), z(t)) := \beta_1 + \beta_2\left(\gamma(t)^\theta + f\frac{\kappa_2}{\kappa_1}z(t)(1 - \gamma(t)^\theta)\right) \tag{1g}$$

where $N(t) = S(t) + I(t) + R(t)$ is the total population.

These equations allow for births at rate ν, deaths from COVID-19 at rate μ_I, and deaths from other causes at rate μ but we set those three parameters to zero because the COVID-19 epidemic is playing out over a time horizon that is short enough that births and deaths are not greatly affecting the total population.

The equations also allow a backflow of recovered individuals back into the susceptible state at a rate φ. How long acquired immunity lasts varies by disease. Immunity to smallpox was once thought to be relatively brief (3-5 years), but is now understood to be longer. Immunity to any specific cold rhinovirus is prolonged, but there are so many rhinoviruses that we can keep getting colds year after year. How long immunity will last with SARS-CoV-2 virus is not known at this time, but immunity to other corona viruses often lasts 3-5 years, so we set φ to 0.001 per day in our base case, which corresponds to a mean duration of immunity of $1000/365 = 2.74$ years.

2.4. Objective function

The other essential part of an optimal dynamic control model is the objective. Optimally responding to COVID-19 requires juggling three to five key considerations, depending on whether one lumps all economic considerations together or breaks them out.

Of course the primary consideration is health which we model as in an earlier paper, see Caulkins et al. (2020). Deaths dominate health costs because the duration of sickness is relatively short compared with diseases such as cancer, let alone dementia. An important contribution of Caulkins et al. (2020) that we also include here is making the risk of death for an infected individual depend on the population-prevalence because the healthcare system can become swamped. In particular, if the number of infected individuals $I(t)$ times the probability that an infected person needs critical care p is less than the healthcare system's capacity ($H_{\max}$) then the death rate has one value (ξ_1); otherwise it gets bumped up by an additional increment (ξ_2). Implementing that literally would require a function with a discontinuous derivative, but as Caulkins et al. (2020) explain, it is possible to find a continuously differentiable function which very closely approximates it. Hence, the health care cost component of the objective function is:

$$V_h(I,\gamma) := M\left(\xi_1 pI(t) + \xi_2 \max_{\mathrm{s}}(\{0, pI(t) - H_{\max}\}, \zeta)\right)$$

with

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

The only difference relative to our previous paper is that we have reduced the value of ξ_1 since the healthcare system has developed better ways of caring for COVID-19 patients (when its capacity to treat is not overwhelmed).

Two of the economic costs are the same as in Caulkins et al. (2020). The first is the reduction in economic activity up until time T, when a vaccine is widely deployed. Economic activity is modeled with a standard Cobb-Douglas form so output is proportional to the number of workers $L(t)$ times the proportion who are working $\gamma(t)$ raised to an exponent σ that is less than one (2/3 in our base case parameter set). Infected individuals are assumed to be too sick to work, so $L(t) = S(t) + R(t)$. Since the time horizon is relatively short, capital K is assumed to be fixed, and without loss of generality is set equal to 1, meaning the units of the objective function are a day's economic output at full employment pre-COVID. The economic loss to be minimized is the difference between what production would have been through time T in the absence of COVID ($TKL(0)^\sigma\gamma(0)^\sigma$) – which sits outside the integral over time since it is a constant – minus the equivalent term with $L(t)$ and $\gamma(t)$ varying over time due to COVID.

The second that is the same as in Caulkins et al. (2020) is the residual loss in economic activity after the vaccine is deployed, because it takes time for full employment to be restored. This is the difference between economic output at time T versus time 0 multiplied by a constant

Γ representing the restoration time. For example, if residual unemployment declined linearly to zero over two years, then Γ would be one year (or 365 days) taking into account that over these two years, on average residual unemployment equals half of the amount of unemployment at time T. We use that as our base case parameter value, but note that it does not imply a linear recovery; any shape of decay that integrated out to the equivalent of one year would be equivalent.

The third economic term is the cost of adjusting employment $\gamma(t)$. This is not the cost of people being unemployed but rather the cost of opening or closing businesses, such as loss of perishable inventory upon shut down and start-up costs when re-opening. As is customary we make these quadratic in the control $u(t)$ and allow for them to be asymmetric with different constants for shutting down businesses c_l and reopening them c_r, with an extra penalty for reopening after an extended shut down so that

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

Putting all of these elements together, the resulting optimal control model will be the following:

$$V(X_0,u(\cdot)) := \int_0^T \left(V_l(L(t),\gamma(t)) - V_h(I(t),\gamma(t)) - V_u(u(t),\gamma(t))\right)\mathrm{d}t \tag{2a}$$

$$- TKL(0)^\sigma\gamma(0)^\sigma - \Gamma\left(KL(0)^\sigma\gamma(0)^\sigma - KL(T)^\sigma\gamma(T)^\sigma\right)$$

$$V^*(X_0) := \max_{u(\cdot)} V(X_0,u(\cdot)) \tag{2b}$$

$$X(t) := (S(t),I(t),R(t),\gamma(t),z(t)),\quad L(t) := S(t)+R(t),\quad N(t) := S(t)+I(t)+R(t). \tag{2c}$$

$$\text{s.t.}\quad \dot{S}(t) = \nu N(t) - \beta(\gamma(t))\frac{S(t)I(t)}{N(t)} - \mu S(t) + \varphi R(t) \tag{2d}$$

$$\dot{I}(t) = \beta(\gamma(t))\frac{S(t)I(t)}{N(t)} - (\alpha+\mu+\mu_I)I(t) \tag{2e}$$

$$\dot{R}(t) = \alpha I(t) - \mu R(t) - \varphi R(t) \tag{2f}$$

$$\dot{\gamma}(t) = u(t),\quad \gamma(0)=1 \tag{2g}$$

$$\dot{z}(t) = \kappa_1(1-\gamma(t)) - \kappa_2 z(t),\quad z(0)=0 \tag{2h}$$

$$\gamma(t)\le 1,\quad 0\le t\le T \tag{2i}$$

$$\beta(\gamma(t),z(t)) := \beta_1 + \beta_2\left(\gamma(t)^\theta + f\frac{\kappa_2}{\kappa_1}z(t)(1-\gamma(t)^\theta)\right) \tag{2j}$$

$$V_l(L(t),\gamma(t)) := K\gamma(t)^\sigma L(t)^\sigma \tag{2k}$$

$$V_h(I(t),\gamma(t)) := M\left(\xi_1 pI(t) + \xi_2 \max\nolimits_{\mathrm{s}}(\{0, pI(t)-H_{\max}\},\zeta)\right) \tag{2l}$$

$$V_u(u(t),\gamma(t)) := \begin{cases} c_l u(t)^2 & u(t)\le 0 \\ c_r(z(t)+1)u(t)^2 & u(t)>0 \end{cases} \tag{2m}$$

2.5. Necessary Optimality Conditions

The Hamiltonian[1] is

$$\mathcal{H}(X,u,\Lambda) = V_l(L,\gamma) - V_h(I,\gamma) - V_u(u,\gamma) + \Lambda'\dot{X}, \tag{3a}$$

$$\begin{aligned} &= V_l(L,\gamma) - V_h(I,\gamma) - V_u(u,\gamma) + \Lambda_1\left(\nu N - \beta(\gamma)\frac{SI}{N} - \mu S + \varphi R\right) \\ &+ \Lambda_2\left(\beta(\gamma)\frac{SI}{N} - (\alpha+\mu+\mu_I)I\right) + \Lambda_3\left(\alpha I - \mu R - \varphi R\right) \\ &+ \Lambda_4 u + \Lambda_5\left(\kappa_1(1-\gamma) - \kappa_2 z\right) \end{aligned} \tag{3b}$$

[1]In the subsequent we omit time argument t unless needed.

with $\Lambda := (\Lambda_1, \Lambda_2, \Lambda_3, \Lambda_4, \Lambda_5)$ denoting the costate variables. We use the indirect adjoining approach for the pure state constraint (2i), see Hartl et al. (1995). Therefore we define the Lagrangian

$$\mathcal{L}(X, u, \Lambda, \psi) := \mathcal{H}(X, u, \Lambda) + \psi u. \tag{3c}$$

For the derivatives we find

$$\frac{\partial}{\partial u}\mathcal{H}(X, u, \Lambda) = \begin{cases} 2c_l u + \Lambda_4 & u \leq 0 \\ 2c_r u(z+1) + \Lambda_4 & u > 0 \end{cases} \tag{3d}$$

$$\frac{\partial^2}{\partial u^2}\mathcal{H}(X, u, \Lambda) = \begin{cases} 2c_l & u \leq 0 \\ 2c_r z & u > 0 \end{cases} \tag{3e}$$

Let $(X^*(\cdot), u^*(\cdot))$ be an optimal solution. Then the Hamiltonian maximizing condition yields for $\gamma^*(t) < 1$

$$u^*(t) = \operatorname*{argmax}_u \mathcal{H}(X^*(t), u, \Lambda(t)) = \begin{cases} -\dfrac{\Lambda_4(t)}{2c_l} & \Lambda_4(t) \geq 0 \\ -\dfrac{\Lambda_4(t)}{2c_r(z^*(t)+1)} & \Lambda_4(t) < 0 \end{cases} \tag{3f}$$

For $z(t) > 0$ the second order derivative is strictly positive and the Hamiltonian is regular. Due to the state dynamics (1e) and the initial condition $z(0) = 0$, it holds that $z(t) = 0$ for $t \in [0, T_s]$ iff $\gamma(t) = 1, t \in [0, T_s]$ and either $T = T_s$ or $\gamma(t) < 1$ for $T_s < t < T_s + \varepsilon$ with some $\varepsilon > 0$. Therefore, in order to have $z(t) = 0$ for $t \in [0, T_s]$, it has to hold that $u(t) = 0$. Thus, the control value is unique and hence, the control $u(\cdot)$ continuous.

For the Lagrangian multiplier ψ we formally solve

$$\frac{\partial}{\partial u}\mathcal{L}(X, u, \Lambda, \psi) := \frac{\partial}{\partial u}\mathcal{H}(X, u, \Lambda)|_{u=0} + \psi = 0$$

yielding

$$\psi = -\Lambda_4$$

and

$$\dot{\psi} = -\dot{\Lambda}_4.$$

Let $(X^*(\cdot), u^*(\cdot))$ be an optimal solution. Let $\tau_i,\ i = 1, \ldots n$ be connecting times $0 < \tau_1 < \ldots < \tau_n < T$ and I_s, I_e and I_x three pairwise disjoint sets with $I_s \cup I_e \cup I_x = \{1, \ldots, n\}$. These sets are defined as

$$j \in I_s \quad \text{iff for some } \varepsilon > 0 \quad u(t) \begin{cases} < 0 & \tau_j - \varepsilon < t < \tau_j \\ = 0 & t = \tau_j \\ > 0 & \tau_j < t < \tau_j + \varepsilon \end{cases}$$

$$j \in I_e \quad \text{iff for some } \varepsilon > 0 \quad \gamma(t) \begin{cases} = 1 & \tau_j \leq t < \tau_j + \varepsilon \\ < 1 & \tau_j - \varepsilon < t < \tau_j \end{cases}$$

$$j \in I_x \quad \text{iff for some } \varepsilon > 0 \quad \gamma(t) \begin{cases} = 1 & \tau_j - \varepsilon < t \leq \tau_j \\ < 1 & \tau_j < t < \tau_j + \varepsilon. \end{cases}$$

The set I_s contains the switching times for the control from being strictly positive to strictly negative. I_e is the set of entry times and I_x the set of exit times for the state constraint.

Then there exists a costate $\Lambda(\cdot)$ being continuously differentiable for $t \in (\tau_i, \tau_{i+1}),\ i = 0, \dots n$ with $\tau_0 := 0$ and $\tau_{n+1} := T$. The Lagrangian multiplier $\psi(\cdot)$ is piecewise continuously differentiable. For each $i \in I_e$ there exists $\chi_i \in \mathbb{R}$. In each interval $(\tau_i, \tau_{i+1}),\ i = 0, \dots n$ the costates $\Lambda(\cdot)$ satisfy the adjoint ODEs

$$\dot{\Lambda}(t) = -\frac{\partial}{\partial X}\mathcal{H}(X^*(t), u^*(t), \Lambda(t)), \quad t \in (\tau_i, \tau_{i+1}),\ i = 0, \dots n. \tag{4}$$

At the connecting times for the state, costates and Lagrangian multiplier it holds that

$$\begin{aligned} &X(\tau_i^-) = X(\tau_i^+), i = 1, \dots, n \\ &\Lambda(\tau_j^-) = \Lambda(\tau_j^+), \quad j \in I_s \cup I_x \\ &\Lambda_{1,2,3,5}(\tau_j^-) = \Lambda_{1,2,3,5}(\tau_j^+), \quad j \in I_e \\ &\Lambda_4(\tau_j^-) = \Lambda_4(\tau_j^+) - \chi_j, \quad j \in I_e \end{aligned}$$

with

$$\chi_j \geq 0, \quad j \in I_e.$$

The Lagrangian multiplier $\psi(\cdot)$ satisfies the complementary slackness condition

$$\psi(t)(1 - \gamma(t)) = 0$$

and

$$\begin{aligned} &\psi(t) \geq 0, \quad \tau_j \leq t \leq \tau_{j+1},\ j \in I_e \text{ or } j + 1 \in I_x \\ &\psi(\tau_j^+) = \chi_j, \quad j \in I_e. \end{aligned}$$

Additionally $\dot{\psi}(\cdot)$ has to satisfy

$$\dot{\psi}(t) \leq 0, \quad \tau_j \leq t \leq \tau_{j+1},\ j \in I_e \text{ or } j+1 \in I_x.$$

For $\gamma(T) < 1$ the costates satisfy the transversality conditions

$$\Lambda(T) = -K \frac{\partial}{\partial X} L(T)^\sigma \gamma(T)^\sigma$$

and for $\gamma(T) = 1$ the costate Λ_4 has to satisfy

$$\Lambda_4(T) = -K \frac{\partial}{\partial \gamma} L(T)^\sigma \gamma(T)^\sigma + \chi_T$$

with

$$\chi_T \geq 0.$$

Table 1 shows the base case parameter values; compare also Caulkins et al. (2020).

α	$\beta_{1,2}$	$H_{\max}$	p	M	K	Γ	f	κ_1
$\frac{1}{15}$	$0, 0.2$	2×10^{-4}	2.25×10^{-2}	$*$	1	365	0.05	0.15

κ_2	σ	μ	ν	μ_I	ζ	ξ_1	ξ_2	φ	c_l	c_r	θ
0.2	$\frac{2}{3}$	0	0	0	5000	0.03	1	0.001	1000	5000	2

Table 1: Base case parameter values.

3. Results

3.1. Results with base case parameters

For the base case parameters in Table 1 three qualitatively different solution strategies can be optimal depending on the value of M, which denotes the value of preventing a death due to COVID-19. Trajectories for γ, the level of employment, are shown in panels (a), (b), and (c) of Fig. 1.

The first regime applies for small values of M; it has only one relatively brief lockdown early on to dampen the intensity of the epidemic (Panel (a)). In the second regime, for intermediate values of M, it is optimal to have two separate lockdowns, one early and another – often smaller – later, shortly before the vaccine gets widely deployed (Panel(b)). In the third regime, with larger values of M, there is just one lockdown, but it is sustained (Panel (c)). In this case, that

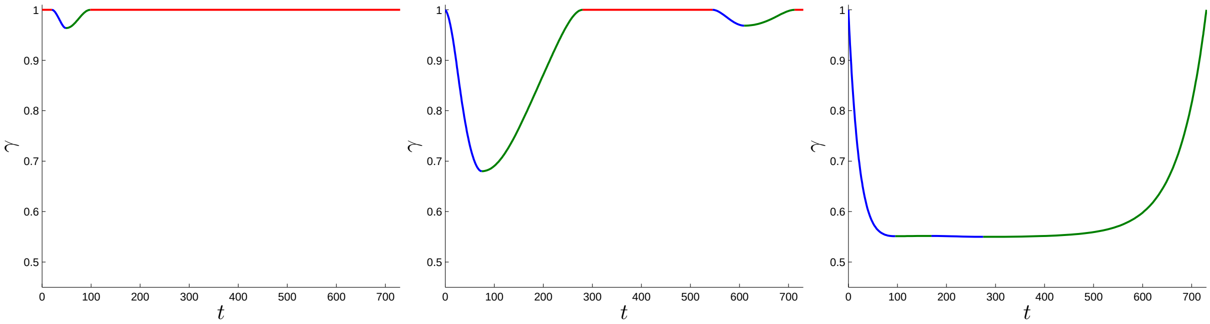

(a) Regime I (b) Regime II (c) Regime III

Figure 1: Panels (a), (b) and (c) show solution paths for each of the three regimes which differ with respect to the size of the social cost of a death M. On the blue part of the solution paths the value of the control u is negative, on the green part it is positive and on the red part it is zero with the constraint $\gamma \leq 1$ being active.

effectively drives the epidemic down to minimal levels for an extended time. We call these the "short lockdown", "double lockdown" and "sustained" strategies; they correspond to Regimes I, II, and III in Fig. 2, respectively.

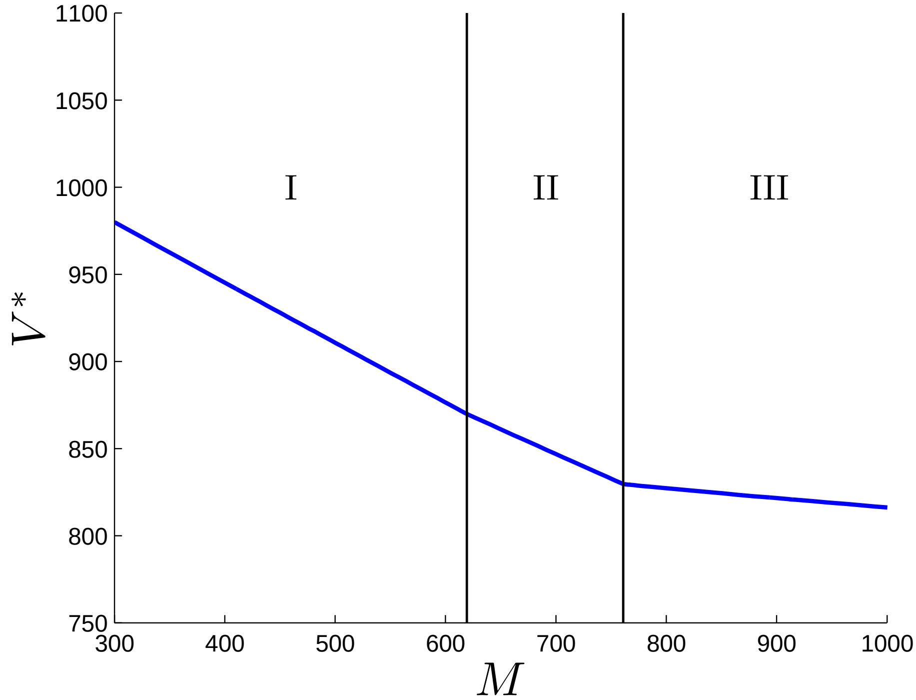

Figure 2: Objective value of the optimal solution path depending on the social cost of a death M for the base case parameters given in Table 1. There are three regimes which differ by the duration, intensity as well as the number of lockdowns of the optimal solutions. For the values of M highlighted by a vertical black line ($M = 619.4$ and $M = 761$) two different solution paths are optimal.

Naturally, as Fig. 2 shows, the objective function value is decreasing in M; the more costly a death, the less well the social planner can do. The slope is initially steep because with only a brief initial lockdown, there are many infections, and so many deaths. Increasing the cost per death reduces the objective function value at a steep rate. That is also true in the second regime that has two lockdowns, implying that the total number who become infected is rather large for that strategy as well. Only when M becomes large and it is optimal to sustain a strong lockdown that sharply reduces deaths does the dependence of V on M become less steep.

The vertical lines in Fig. 2 passing through the kinks in $V(M)$ are points at which two different strategies perform equally well, which is illustrated in Fig. 3. For example, when M = 619.4 the solid and dashed trajectories perform equally well overall, even though the Regime I strategy limits unemployment to less than 10% whereas the Regime II approach allows unemployment to exceed 30% at one point. Likewise, Fig. 3 Panel (b) shows two very different trajectories that perform equally well when $M = 761$. The solid line is a double lockdown strategy that is very similar to the double lockdown strategy in Fig. 3a; the dashed line shows a sustained and radically more aggressive lockdown that suffers unemployment over 40% for more than a year but – as noted – greatly reduces infections and deaths.

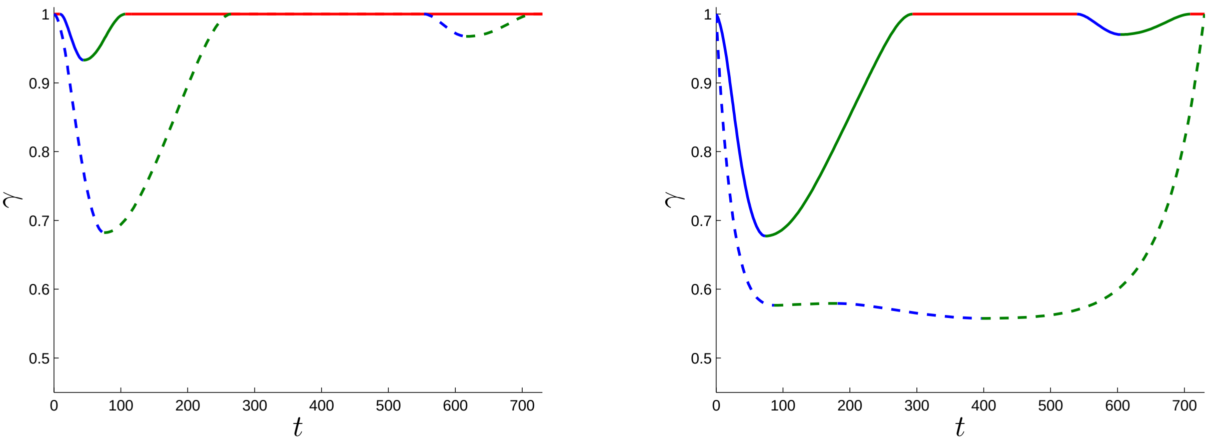

(a) (b)

Figure 3: Optimal time paths of the the proportion of working people corresponding to the different regions shown in Fig. 2. Panel (a) and (b) return the optimal solutions for the values of M highlighted by the vertical black line (for (a) $M = 619.4$ and (b) $M = 761$).

These points at which there are alternate optimal strategies are Skiba points. From the same initial point, two different trajectories emerge but which produce the same optimal objective function value.

Fig. 4 shows in greater detail the consequences of following the two strategies that are optimal when $M = 619.4$. Note that both strategies involve roughly the same number of people getting infected, but the double lockdown strategy (dashed line) starts with a larger and longer initial lockdown; that flattens the curve considerably. With the single, small lockdown (solid dark lines), at the epidemic's peak a quarter of the population is infected at one time, which would completely swamp hospital's treatment capacity. That is, when M is small, not only does the optimal strategy allow many people to become infected, it lets many of them get infected at the same time, so many who need critical care cannot receive it, increasing the number of deaths. Indeed, the epidemic's trajectory with that strategy is not so different than it would be with no lockdown (the light lines).

Quite a few people still become infected with the double lockdown strategy, as can be seen by the decline in the number of susceptibles (Panel (c)) and increase in the number of Recovered individuals (Panel (e)), but the infections are spread out over time.

Fig. 5 gives further detail contrasting the two solutions that are optimal when $M = 761$. These strategies differ even more markedly, as is perhaps best seen in Panel (b), which shows that with this sustained strategy, the infection rate never exceeds the hospital's capacity which is indicated by the horizontal solid line.

It is very interesting to contrast the different strategies' variation over time in the epidemic's effective reproductive number (R_{eff}), meaning the raw reproductive number modified by both the control intervention and also the accumulation of people in the Recovered state. Both, the two-lockdown and the sustained strategies for these parameter values keep R_{eff} close to 1 throughout most of the time horizon.

The only exception is that the sustained strategy allows the effective reproductive rate to increase just before the vaccine is distributed. At that point the number of infections is so low, that even a month or two of spread does not push the absolute number of infections up very high.

Note that strategies involving a change in policy a month or two before the vaccine is widely deployed are not unrealistic. Although it is not possible to predict when a vaccine will be invented, there is a lag between invention and widespread deployment during mass production and distribution. The production and distribution stages are reasonably well-understood processes, so their duration is fairly predictable. That means a strategy that calls for a change 30 or 60 days before the vaccine has been fully deployed is feasible.

The speed of the epidemic's spread requires this hovering of R_{eff} near 1.0 for any "interior" solution with a substantial pool of susceptibles. The time from infection to end of infectiousness

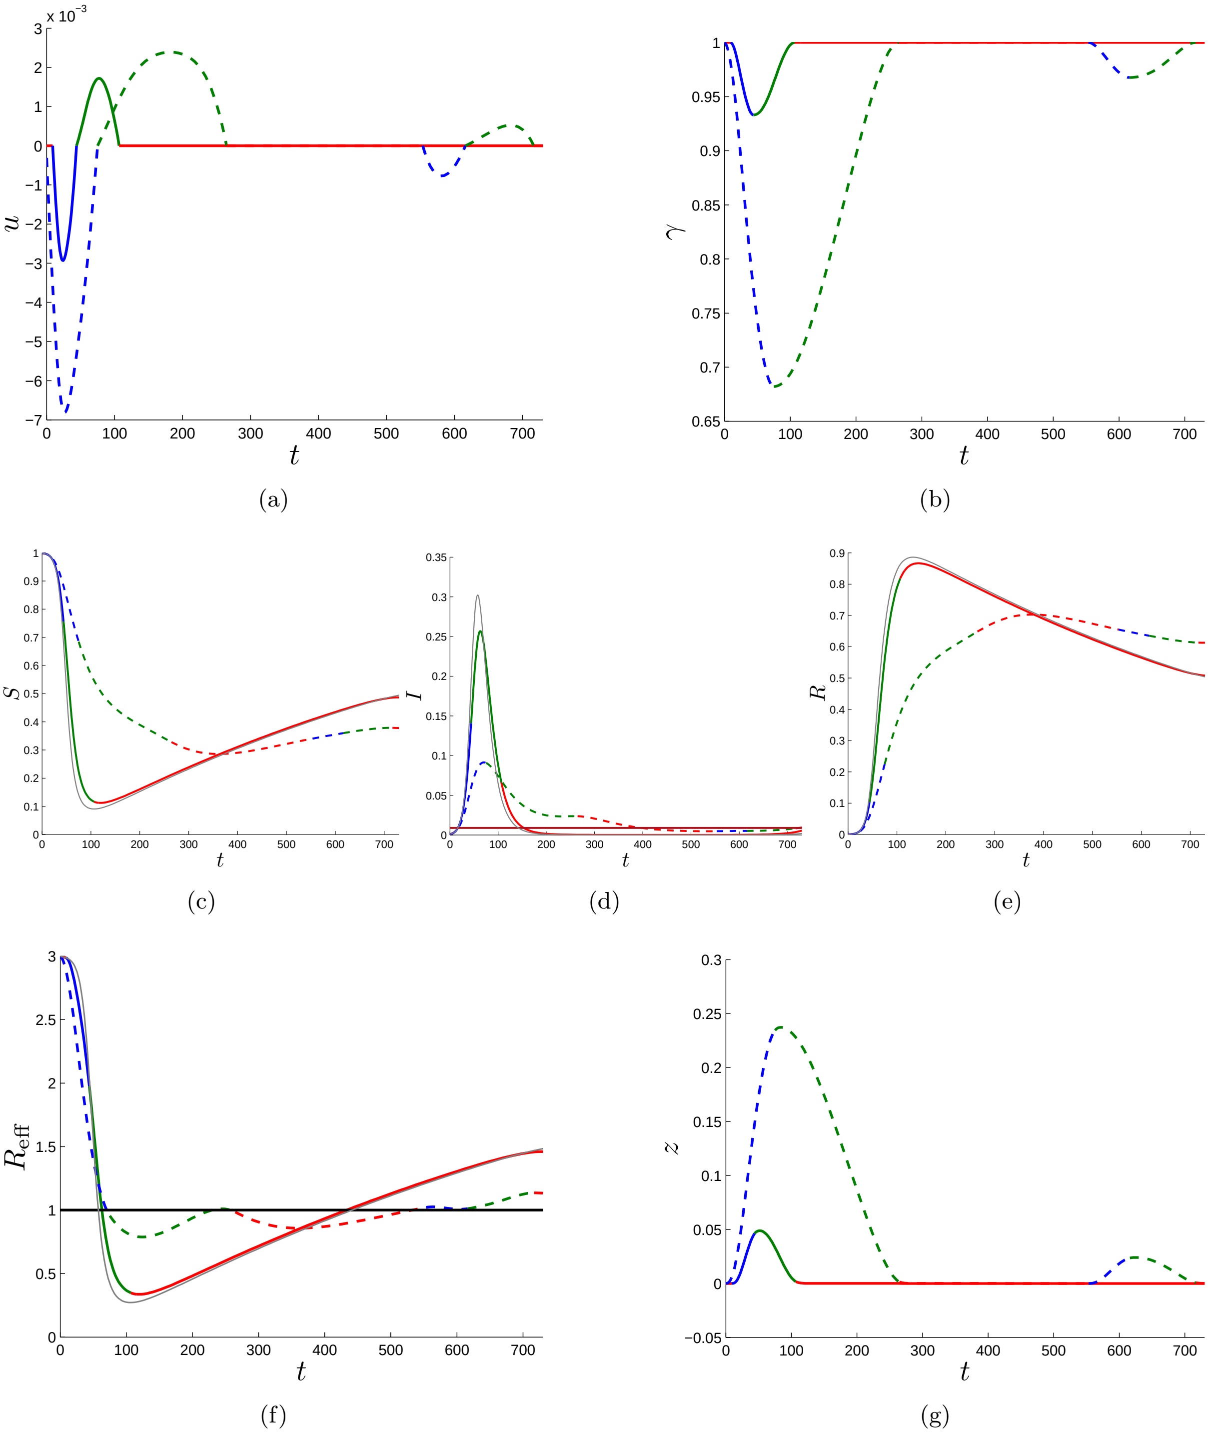

Figure 4: Time paths for the Skiba solutions at $M = 619.4$. Panel (a) depicts the control and panel (b) the proportion of working people. Panel (c) shows the number of susceptibles, panel (d) the number of infected (above the red horizontal line hospital capacity is exceeded), panel (e) the number of recovered patients. In panel (f) the effective reproduction number can be seen and in panel (g) the perceived lockdown intensity z. The gray line shows the uncontrolled epidemic's time path.

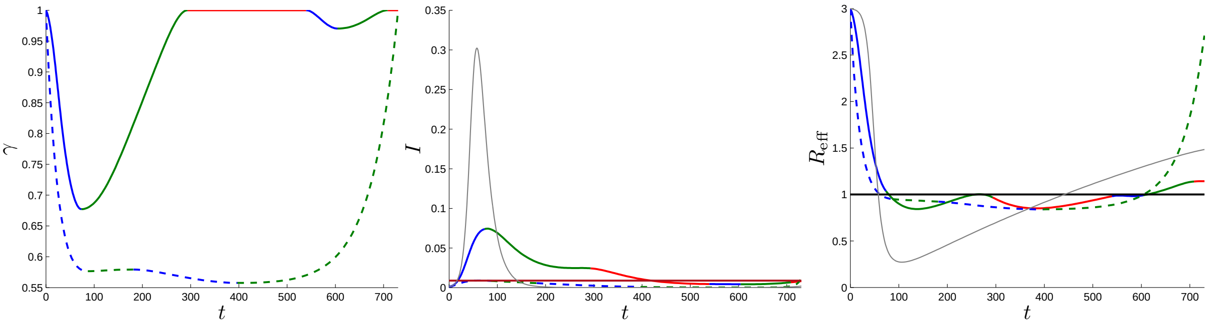

(a) (b) (c)

Figure 5: Time paths for the Skiba solution at $M = 761$. Panel (a) depicts the proportion of working people, panel (b) the number of infected and panel (c) the effective reproduction number.

is short; about two weeks. So within a 52-week year, that reproductive rate can effectively get raised to the 26th power. If it is anything other than about 1, that will cause the number of infected individuals to vary rapidly. Regime I strategies dispense with that stability, with R_{eff} swinging from 3 to one-third over just three months, before rebounding to well above 1, and that variation is driven by variation in the numbers in the susceptible and recovered states.

So in a sense, one can think of the three strategies as follows. Regime I: Let the epidemic run its course, more or less. Regime II: Smooth out the epidemic curve with intermittent pulses of locking down. Regime III: Use a sustained, deep lockdown to forestall the epidemic.

3.2. Triple Skiba points

The model produces rich behavior, including the possibility of a triple Skiba point from which three distinct optimal trajectories emerge. The additional strategy at these points involves a double pulse lockdown but without ending it completely in between. That is in some sense intermediate between Regimes II and III because some degree of lockdown is maintained throughout (as in Regime III) but there are two distinct waves of locking down (as in Regime II).

Fig. 7 Panel (d) illustrates a typical trajectory with this new Regime IV strategy.

Fig. 6 Panel (a) shows two triple-Skiba points in a bifurcation diagram in the space of parameters M (cost of a death) and c_r, which governs the adjustment costs of reopening businesses that have been shut down. In Panel (c), which corresponds to the left-hand triple Skiba point at $c_r = 75,672.5$ and $M = 869$, there is a standard Regime I strategy (solid line), a standard Regime II strategy (dot-dashed line) and a new (Regime IV) type strategy (dashed line). At this particular point, the Regime IV strategy appears only very slightly different than the Regime

II strategy because γ comes very close to 1.0 near the end of year 1 and beginning of year 2.

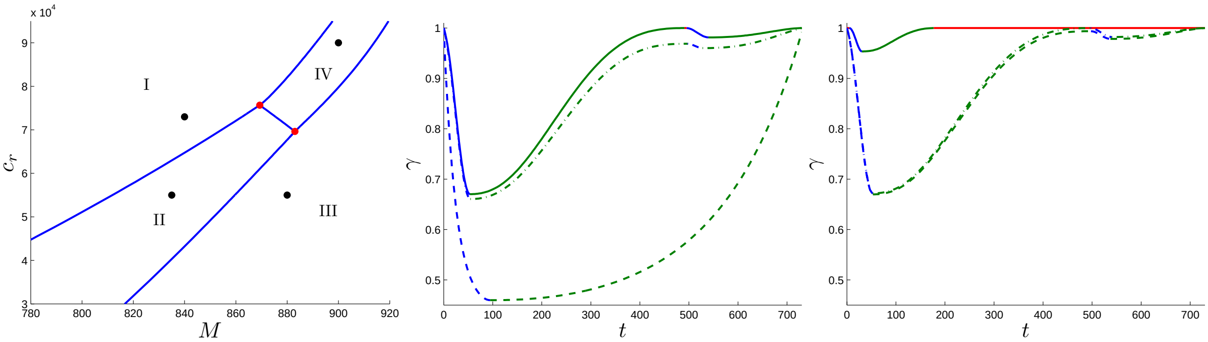

(a) (b) $c_r = 69632.7, M = 883$ (c) $c_r = 75672.5, M = 869$

Figure 6: Bifurcation diagram in the M–c_r space, for $\varphi = 0.001, \beta_1 = 0, \beta_2 = 0.2$ and $\kappa_1 = 0.15, \kappa_2 = 0.2, c_l = 1000, \xi_1 = 0.03, \xi_2 = 1$ (panel (a)). The blue curves show the Skiba curves and the dots denote the triple Skiba points (red), examples for the Skiba solutions (green) and examples for the different regions (black). Panels (b) and (c) show the time paths of γ starting at the triple Skiba points occurring for $c_r = 69,632.7, M = 883$ and $c_r = 75,672.5, M = 869$.

Note that it makes sense that the line dividing Regions I and II/IV slopes upward in this bifurcation diagram. Larger values of M place greater value on reducing infection, justifying more locking down, and smaller values of c_r mean that two waves of locking down is not prohibitive because the cost of businesses rebuilding after a lockdown is not so high (small c_r).

Panel (b), which corresponds to the right-hand triple Skiba point with $c_r = 69,632.7$ and $M = 883$, has a Regime II strategy (solid line), a Regime IV strategy (dot-dashed line), and a Regime III near-eradication strategy (dashed line).

The upward slope of the line separating the II/IV and III regimes again makes sense. Larger M favors a sustained lockdown, so Regime III is to the right. Also, Regime III involves a deeper trough in employment and so a steeper rebuilding, so that option does well when reopening shuttered industries is not too expensive (c_r is small).

Fig. 8 show that it is also possible to have a triple Skiba involving trajectories from Regimes I and III. Triple Skiba involving both Regime II & IV strategies and one other are perhaps unsurprising; both Regimes II and IV are both in some sense double lockdown strategies. They are distinguished in Figure 6 only by whether the interlude between the two pulses was long enough to allow γ to recover all the way to 1.0. That a decision maker could be indifferent between a Regime I and III strategy (as well as one from Regime II or IV) might seem more surprising because the Regime I and III strategies in Fig. 8 appeared to be near opposites. The Regime I trajectory involved very little in the way of locking down and the Regime III strategy

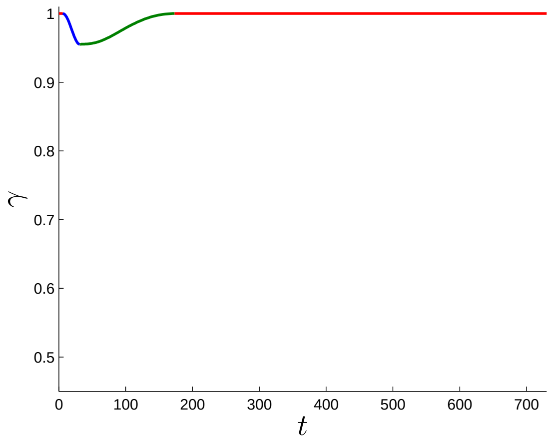

(a) Regime I

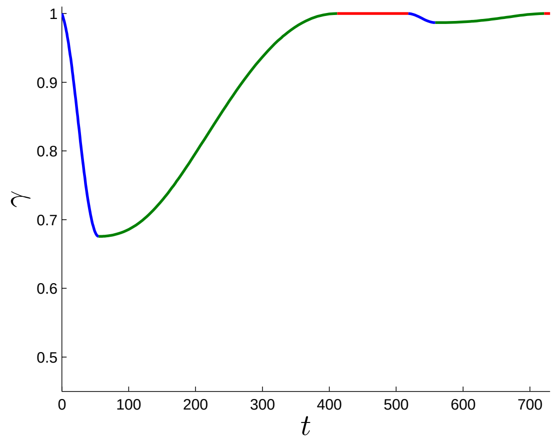

(b) Regime II

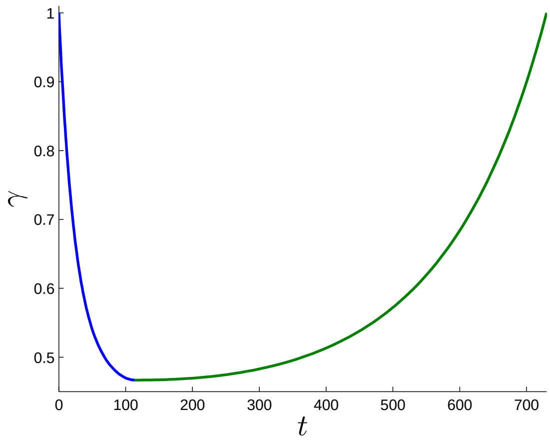

(c) Regime III

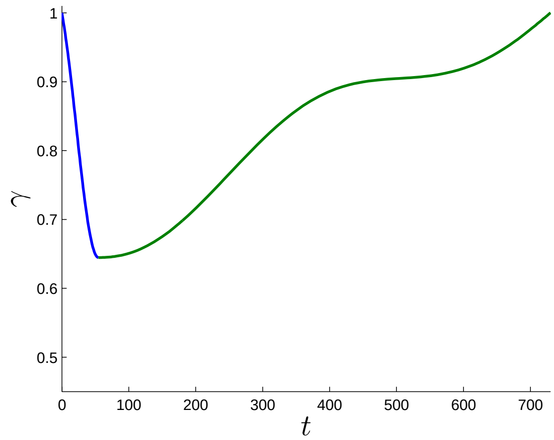

(d) Regime IV

Figure 7: Time paths of γ for the four different regimes depicted in Fig. 6.

was at the other extreme.

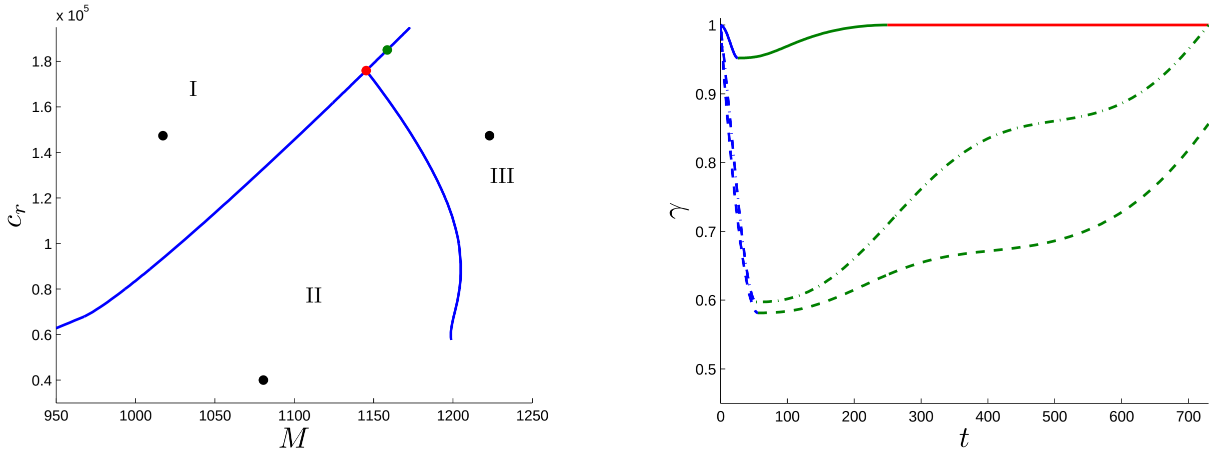

(a)

(b) $M = 1145.3, c_r = 175904.5$

Figure 8: Bifurcation diagram in the M–c_r space, for $f = 0.25$ (panel (a)) and time paths of γ (panel (b)) for solutions starting at the triple Skiba point for $M = 1145.3, c_r = 175904.5$ depicted by a red dot in panel (a).

Fig. 9 shows how results differ if it takes three not two years to develop and deploy an effective vaccine. The qualitative results seem little changed. In particular, the two lockdown strategy continues to employ just two lockdowns; it did not morph into a three lockdown strategy. Likewise, that second lockdown continues to come soon before the vaccine is deployed, not a certain, fixed time after the first lockdown ends.

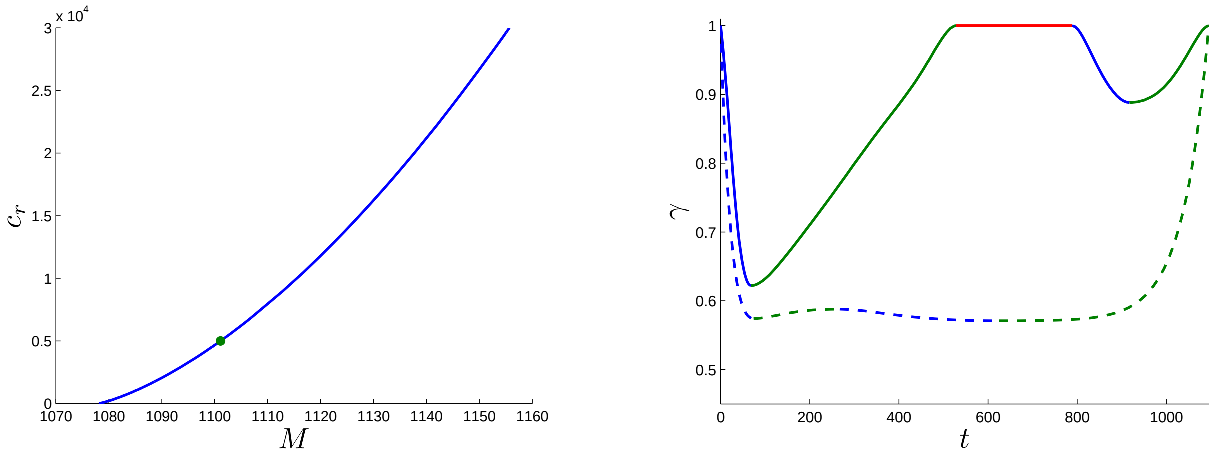

(a) $T = 1095$

(b) $T = 1095, c_r = 5000$

Figure 9: Bifurcation diagram in the M–c_r space for $T = 1,095$.

4. Discussion

Perhaps the most basic conclusion of this analysis is that very different strategies for responding to the COVID-19 pandemic can be optimal with the exact same set of parameter values. Exact equality of performance is a knife-edge case, occurring only exactly at the Skiba point. However, there are neighborhoods around the Skiba points where alternate, very different strategies perform nearly as well.

A second basic conclusion is that even when only a single strategy is optimal, which specific strategy wins can change quickly when certain parameters values vary over a relatively limited range. This is perhaps best illustrated with respect to M, the parameter standing for the cost to the social planner per premature death. There is a long literature discussing what is the appropriate value to use for that parameter in social welfare analysis. There is some common understanding as to the order of magnitude, but considerable debate as to the particular value. That is not surprising inasmuch as it is not an empirical constant akin to the atomic mass of an element so much as an expression of values, and different people can have different values about how they wish to trade-off life and health with economic outcomes (such as unemployment) and happiness more generally (including freedom of association).

Fig. 6 shows that for our base case value of parameter $c_r = 5,000$ (standing for the cost of reopening shuttered businesses), varying parameter M by less than 10% (from slightly below 800 up to 860) carries one all the way across the bifurcation diagram. When M is (a bit smaller than) 800, one is in Regime I where it is optimal to more or less let the regime run its course. When M is a bit larger than 800, it is optimal to have two distinct lockdowns, both lasting well less than a year. And by the time M reaches 860, it is optimal to have one sustained lockdown that involves a very substantial loss of employment, but also a very substantial reduction in infection and death.

A third observation is simply that strategies involving two lockdowns can be optimal. A number of jurisdictions that locked down then opened up are now having to reinstitute restrictions. For example, Israel was once in the top five highest in the world for new infections per capita. It drove that all the way down to below 0.2 per 100,000 per day and so appeared to have largely eliminated infections, but has recently bounced back up into the top 5 as of this writing, with about 19 new confirmed infections per 100,000 per day. Superficially, that appears to be a policy disaster and, indeed, Israel's resurgence of infections may indicate policy failure in this case; certainly Prime Minister Netanyahu is facing strident protests for his leadership on this issue. But the model shows that the mere presence of a resurgence necessitating a second lockdown is not in and of itself proof of error. A double lockdown can be an optimal strategy and part of an optimal plan from the outset.

A fourth observation concerns the Skiba points. Skiba points separate distinct optimal solution trajectories that spread out from a common initial condition in different directions. In a one-state problem, there would generally be one strategy that moves left and another that moves right from that common initial condition. Yet when plotted in state space, particularly with respect to γ, which stands for the rate of employment still permitted despite the lockdown, the alternative trajectories here do not appear to be so sharply resolved. With respect to several of the triple Skiba points observed here, all three optimal strategies start with a lockdown that drives down γ, albeit with varying intensities. And in Fig. 8, in particular, the three strategies seem all to be in the interior and on a continuum. Implicitly, if two trajectories are both optimal, then all strategies that are "in between" must be worse. So for every point in time t, we have the following odd situation in Fig. 8. A moderate amount of unemployment is ideal. A little more is bad. Still more brings one back to ideal. Yet more is bad again. But still more is back to being ideal. Not only is social welfare not a monotonic function of unemployment, at every time t, it is a triple-peaked function.

It is worth reflecting on how peculiar this is. Imagine there were seven identical countries that all started at the same point, and we stopped them at some time t in the middle of the epidemic and rank ordered them from "best" to "worst" in terms of amounts of unemployment. Having done that, every second country on that rank-ordered list could be following an optimal policy (meaning countries #2, #4, and #6 are optimal), while every other country is not on an optimal trajectory, even though all started in exactly the same place.

In a way, this is not altogether surprising. We have a five-state problem, so projections onto a single state can be deceiving, and the objective function is a highly nonlinear function of the state variables. On the other hand, all of that nonlinearity and all of those state variables arise naturally from a modeling of the problem; this is not an artificial model constructed just to produce curious results. It is a model that makes a good faith effort to capture the most important dynamics of the epidemic.

One contributor to this seemingly odd behavior is the inherent instability of an epidemic whose total time from infection to end of infectiousness is as short as 14-days. It moves through the population very quickly. That speed forces society into one of three broad postures: (1) High employment but also high infections, (2) High unemployment in order to achieve very low infections, and (3) Walking a delicate balance with multiple pulses of locking down that tolerates a moderate amount of infection but nonetheless keeps the epidemic's reproductive rate close to 1.0.

So the instability of the epidemic naturally produces a standard double-Skiba point, akin to the "eradicate vs. accommodate" options seen in past models, see Tragler et al. (2001), Grass et al. (2008). But it is possible, at least in this model, to delicately walk a fine line in between, that avoids the Draconian measures required for eradication but also keeps the epidemic from massively overwhelming hospital capacity. And the feasibility – and indeed potential optimality – of such intermediate paths is what produces the triple Skiba points and the extreme non-monotonicity of overall performance with respect to state variables, including even the level of unemployment.

5. Conclusion

In sum, this relatively simple model produces a wide range of interesting behaviors that are directly interpretable in terms of the policy context. There are, as always, abundant opportunities for further work and refining the model. Among its limitations at present, we mention a few that are salient. One is not modeling and including a control for testing and contact tracing. It may be that once the number of infections has been driven down sufficiently low, that aggressive testing and tracing could keep the number of infections from rebounding even if everyone went back to work. That would open up a strategy that locks down very aggressively and for a moderately long time, but does not need to sustain the lockdown all but up to the point at which the vaccine becomes widely deployed. That approach would enjoy the best of both worlds – but only after a moderately long period of economic pain.

Another realistic extension would be to recognize that there are different geographic regions with at least some degree of movement between regions. When the two regions are out of synch in terms of their epidemics, then that movement might trigger a resurgence in a low prevalence region with migrants from a high prevalence region. That possibility has led to very widespread border closures and restrictions on freedom of movement that would have been unimaginable just twelve months ago, and the likes of which have not been seen since the fall of the Soviet Union. It would be tremendously valuable to determine whether all those border closures are truly needed.

Another class of important extensions would be to recognize heterogeneity along at least two dimensions. One is age. Simply put, the infection fatality rate is much, much higher for older people, and for those with certain preexisting medical conditions, than it is for young healthy people. So the tradeoff between economic loss and health harm involves a very large distributional issue. It is working age people who become unemployed and (for the most part) retirees who reap the majority of the health benefits of that loss of income.

There is also important heterogeneity across people in terms of how active they are socially or, in the jargon of HIV/AIDS models, how many risky acts they pursue. Some people are naturally socially isolated even before quarantine; others are social butterflies who frequent indoor places with much circulation of people and little recirculation of the air. Because of stochastic selectivity, high-rate transmitters will be disproportionately over-represented among those who get infected and recover early. That means the effective amount of herd immunity will be greater than is reflected in this model, which treats all people as homogenous with respect to the number of risky contacts they have per unit time.

Of course many more such extensions would be possible. So we close with a final meta-observation. When a central policy response to a pandemic involves shutting down the economy, there are not only complex value tradeoffs, but also complex state dynamics that provide ample fodder for interesting optimal control modeling. Since COVID-19 is unlikely to be the last important pandemic in our lifetimes, that suggests there may be considerable value in analyzing models now that are inspired by COVID-19, but which do not slavishly model it exactly. Instead, there is value in abstracting somewhat to capture the general tensions and considerations that such pandemics create. That way we can not only deal more effectively with the current crisis, but also be better prepared to respond to the next one.

A. Properties of the function $\beta(\gamma, z)$

We choose β_1 and β_2 such that $\beta_1 + \beta_2 = \bar{\beta}$, where $\bar{\beta}$ is the contact rate of the "uncontrolled" epidemics and $0 \leq f \leq 1, \theta \geq 1$.

$$\beta(\gamma, z) := \beta_1 + \beta_2 \left(\gamma^\theta + f \frac{\kappa_2}{\kappa_1} z (1 - \gamma^\theta) \right) \tag{A.1a}$$

yields

$$\beta(1, z) = \beta_1 + \beta_2 = \bar{\beta} \tag{A.1b}$$

$$\beta(\gamma, z) > \beta_1 + \gamma^\theta \beta_2, \quad \text{for} \quad z > 0,\ 0 \leq \gamma < 1 \tag{A.1c}$$

$$\beta(\gamma, 0) = \beta_1 + \gamma^\theta \beta_2 \tag{A.1d}$$

$$\beta(\gamma, z) < \bar{\beta}, \quad \gamma < 1 \tag{A.1e}$$

$$\frac{\partial}{\partial \gamma}\beta(\gamma, z) = \beta_2 \theta \gamma^{\theta - 1}\left(1 - f\frac{\kappa_2}{\kappa_1}z\right) > 0, \quad \gamma > 0 \tag{A.1f}$$

$$\frac{\partial}{\partial z}\beta(\gamma, z) = \beta_2 f \frac{\kappa_2}{\kappa_1}(1 - \gamma^\theta) > 0, \quad \gamma < 1 \tag{A.1g}$$

Inequalities Eqs. (A.1d) to (A.1g) follow from

$$z(t) < \frac{\kappa_1}{\kappa_2}, \quad \text{for all } t \text{ with} \quad z(0) = 0.$$

References

D. Acemoglu, V. Chernozhukov, I. Werning, and M. D. Whinston. A multi-risk SIR model with optimally targeted lockdown. Working Paper 27102, National Bureau of Economic Research, 2020.

F. E. Alvarez, D. Argente, and F. Lippi. A simple planning problem for COVID-19 lockdown. Working Paper 26981, National Bureau of Economic Research, April 2020.

A. Aspri, E. Beretta, A. Gandolfi, and E. Wasmer. Mortality containment vs. economics opening: optimal policies in a SEIARD model. *arXiv preprint arXiv:2006.00085*, 2020.

J. P. Caulkins, G. Feichtinger, D. Grass, R. F. Hartl, P. M. Kort, and A. Seidl. Skiba points in free end-time problems. *Journal of Economic Dynamics and Control*, 51:404 – 419, 2015. doi: 10.1016/j.jedc.2014.11.003.

J. P. Caulkins, D. Grass, G. Feichtinger, R. F. Hartl, P. M. Kort, A. Prskawetz, A. Seidl, and S. Wrzaczek. How long should the COVID-19 lockdown continue? Technical report, ORCOS Research Report 2020-10, 2020.

M. Gonzalez-Eiras and D. Niepelt. On the optimal 'lockdown' during an epidemic. CEPR Discussion Paper 14612, Swiss National Bank, Study Center Gerzensee, 2020.

D. Grass. Numerical computation of the optimal vector field in a fishery model. *Journal of Economic Dynamics and Control*, 36(10):1626–1658, 2012. doi: 10.1016/j.jedc.2012.04.006.

D. Grass, J. P. Caulkins, G. Feichtinger, G. Tragler, and D. A. Behrens. *Optimal Control of Nonlinear Processes: With Applications in Drugs, Corruption, and Terror*. Springer, Berlin, 2008.

R. F. Hartl, S. P. Sethi, and R. G. Vickson. A survey of the maximum principles for optimal control problems with state constraints. *SIAM Review*, 37(2):181–218, 1995. doi: 10.1137/1037043.

T. Kiseleva and F. O. O. Wagener. Bifurcations of optimal vector fields in the shallow lake system. *Journal of Economic Dynamics and Control*, 34(5):825–843, 2010. doi: 10.1016/j.jedc.2009.11.008.

T. Kiseleva and F. O. O. Wagener. Bifurcations of optimal vector fields. *Mathematics of Operations Research*, 40(1):24–55, 2015. doi: 10.1287/moor.2014.0655.

J. Köhler, L. Schwenkel, A. Koch, J. Berberich, P. Pauli, and F. Allgöwer. Robust and optimal predictive control of the COVID-19 outbreak. *arXiv preprint arXiv:2005.03580*, 2020.

S. P. Sethi. *Optimal Control Theory: Applications to Management Science and Economics*. Springer, Cham, Switzerland, third edition, 2019. doi: 10.1007/978-3-319-98237-3.

A. Steindl and G. Feichtinger. Bifurcations to periodic solutions in a production/inventory model. *Journal of Nonlinear Science*, 14(6):469–503, 2004. doi: 10.1007/s00332-003-0565-x.

G. Tragler, J. P. Caulkins, and G. Feichtinger. Optimal dynamic allocation of treatment and enforcement in illicit drug control. *Operations Research*, 49(3):352–362, 2001. doi: 10.1287/opre.49.3.352.11215.

I. Zeiler, J. P. Caulkins, and G. Tragler. Optimal control of interacting systems with DNSS property: The case of illicit drug use. *Journal of Economic Behavior & Organization*, 78(1-2):60–73, 2011. doi: 10.1016/j.jebo.2010.12.008.