Journal Pre-proof

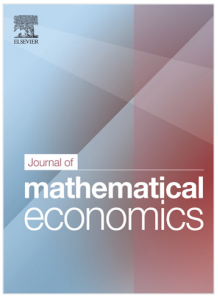

The optimal lockdown intensity for COVID-19

Jonathan P. Caulkins, Dieter Grass, Gustav Feichtinger, Richard F. Hartl, Peter M. Kort, Alexia Prskawetz, Andrea Seidl, Stefan Wrzaczek

PII: S0304-4068(21)00027-6
DOI: https://doi.org/10.1016/j.jmateco.2021.102489
Reference: MATECO 102489

To appear in: *Journal of Mathematical Economics*

Please cite this article as: J.P. Caulkins, D. Grass, G. Feichtinger et al., The optimal lockdown intensity for COVID-19. *Journal of Mathematical Economics* (2021), doi: https://doi.org/10.1016/j.jmateco.2021.102489.

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, 4800 Forbes Avenue, Pittsburgh PA 15213, USA*
[b] *International Institute for Applied Systems Analysis (IIASA), Schlossplatz 1, 2361 Laxenburg, Austria*
[c] *Department for Operations Research and Control Systems, Institute of Statistics and Mathematical Methods in Economics, TU Wien, Wiedner Hauptstraße 8, 1040 Vienna, Austria*
[d] *Wittgenstein Centre for Demography and Global Human Capital (IIASA, OeAW, University of Vienna),Vordere Zollamtsstraße 3, 1030 Vienna, Vienna, Austria*
[e] *Department of Business Decisions and Analytics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria*
[f] *Tilburg School of Economics and Management, Tilburg University, Warandelaan 2, 5037 AB Tilburg, Netherlands*
[g] *Department of Economics, University of Antwerp, Prinsstraat 13, 2000 Antwerp, Belgium*
[h] *Research Group Economics, Institute of Statistics and Mathematical Methods in Economics, TU Wien, Wiedner Hauptstraße 8, 1040, 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 emerge, 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 prevails. Indeed, there are constellations of parameters for which two or even three of these distinct strategies can all perform equally well for the same set of initial conditions; these correspond to so-called triple Skiba points. The performance of trajectories can be highly

*Corresponding author
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)

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 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. Re-opening has been prompted both by progress in pushing down infection rates and also "lockdown fatigue", in which the public's cooperation wanes when lockdowns are perceived of as having gone on too long. Some regions and countries have then seen infection rates rebound and returned to a renewed lockdown, sometimes more severe sometimes less severe than the first. Some places have suffered 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 several ways. For example, the lethality of the infection varies depending on whether there are so many simultaneous infections that critical care capacity has been swamped. The most fundamental extension, though, is creating an objective function that balances three 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 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 modulate 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 very 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, and there are constellations of parameters for which multiple qualitatively different strategies may perform equally well, even though those strategies 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 for sharp disagreements. Hence, 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. Sometimes it appears that the optimal solution involves locking down, ending the lockdown and reinstituting it, sometimes with the second lockdown being more severe than the first. 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 related to extended periods of economic lockdown, although so far, only a minority of these papers have

investigated the optimal timing, length and extent of the lockdown itself. We mention a few of the exceptions. 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. 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 be 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 the lockdown to differ for different age-groups, distinguishing between "young", "middle-aged" and "old" populations, in a SIR model. In that model, 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) consider a SEIRD model, where the population is divided into susceptibles, exposed but asymptomatic, infected, recovered and deceased, and obtain multiple lockdowns as well as Skiba points.

Caulkins et al. (2020) adds to this literature by considering the limited capacity constraint of intensive care units within the health care system. If the number of infected people needing intensive care exceeds the constraint the death rate of these patients increases relative to similar individuals who are able to receive appropriate care. However, the modeling of lockdowns was very simple, optimizing only over the lockdown's start and end time, not over its intensity.

The present paper preserves Caulkins et al. (2020) modeling of intensive care capacity and extends it by allowing for multiple lockdowns with different intensities and lengths. We also explicitly model "lockdown fatigue" by an additional state variable that accumulates the intensity

and length of the lockdown. Greater fatigue undermines the lockdown's effectiveness as people become less compliant with restrictions, so this "memory of lockdowns" affects the efficiency of the lockdown. Furthermore, we assume that adjusting the lockdown is costly. In particular, we allow for an asymmetry in the costs of strengthening and weakening the lockdown.

The analysis also contributes to the celebrated history of papers exploring Skiba thresholds (see Grass et al. (2008), Sethi (2019)). In particular, we find triple Skiba points in a finite time horizon problem, and even multiple triple Skiba points for specific parameter constellations, using bifurcation analyses comparable to those found in Grass (2012) and Kiseleva and Wagener (2010, 2015). 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 papers consider infinite time horizon models, 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 as well.

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-19 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 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, 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 reason 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 economic 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 with 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,

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

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 $z_{\max} = \kappa_1/\kappa_2$.

2.3. Epidemic dynamics

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.

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

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:

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

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 a 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 $f = 0.45$, so this attenuation by lockdown fatigue generally has a somewhat modest force in the analysis below. Nonetheless, we believe it is important to acknowledge this human dimension of how a population responds to extended lockdowns.

2.4. State dynamics

The state dynamics can then be written as

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

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

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

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

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

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

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

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

Note: We write these equations with greater generality than we need or use in this paper. In particular, 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.

That means deaths play a prominent role in the objective function, but not in the state equations. That may seem odd, but it is a reasonable and expedient modeling approximation. COVID-deaths are a central reason why the pandemic is a crisis; they cannot be ignored, so in the objective function, deaths are modeled as a realistic and hence somewhat complicated function of the infection rate. Including those deaths in the state equation would considerably complicate the model, and to little avail. Even though COVID-19 is a horrible pandemic, the infection fatality rate is on the order of 1 percent, so even if everyone were to be infected that

would only reduce the population by 1 percent. Furthermore, the deaths are very highly concentrated among older people who are retired and past the age of having children. So omitting deaths from the state equations is a small discrepancy compared to other approximations and uncertainties in the model.

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. Note: This parameter does influence the character of the optimal solutions, suggesting that figuring out ways of estimating it rapidly for new pandemics could be important for effective policy making.

2.5. Objective function

The other essential part of an optimal 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. A 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, for those who cannot receive critical care, 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:

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

with

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

The label "max" with a subscript s is meant to denote a smoothed version of the maximum function.

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 function but capital is assumed to be fixed because the time horizon is short. 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 $L(t) = S(t) + R(t)$. Without loss of generality capital K 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-19 ($TKL(0)^{\sigma}\gamma(0)^{\sigma}$) – 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-19.

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 recovery 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

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

Putting all of these objective function elements together with the state dynamics, the resulting optimal control model will be the following:

$$V(X_0,u(\cdot)) \coloneqq \int_0^T \left(V_l(L(t),\gamma(t)) - V_h(I(t),\gamma(t)) - V_u(u(t),\gamma(t))\right)\mathrm{d}t \tag{2a}$$
$$- TKL(0)^\sigma\gamma(0)^\sigma - \Gamma\left(KL(0)^\sigma\gamma(0)^\sigma - KL(T)^\sigma\gamma(T)^\sigma\right)$$

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

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

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

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

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

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

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

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

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

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

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

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

2.6. Necessary Optimality Conditions

The Hamiltonian[1] is

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

$$= V_l(L,\gamma) - V_h(I,\gamma) - V_u(u,\gamma) + \Lambda_1\left(\nu N - \beta(\gamma)\frac{SI}{N} - \mu S + \varphi R\right) \tag{3b}$$

$$+ \Lambda_2\left(\beta(\gamma)\frac{SI}{N} - (\alpha+\mu+\mu_I)I\right) + \Lambda_3\left(\alpha I - \mu R - \varphi R\right)$$

$$+ \Lambda_4 u + \Lambda_5\left(\kappa_1(1-\gamma) - \kappa_2 z\right)$$

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

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

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

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

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

For $z(t) > 0$ the second order derivative is strictly positive and the Hamiltonian is regular. For $z(t) = 0$ we find from the state dynamics (1e) that these properties only hold true if $\gamma(t) = 1$. Due to the initial condition $z(0) = 0$, it holds that $z(t) = 0$ can only be satisfied for $t \in [0, T_s]$ with some $T_s \ge 0$, which necessarily implies $\gamma(t) = 1, t \in [0, T_s]$ 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.

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

For the Lagrangian multiplier ψ we formally solve

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

yielding

$$\psi = -\Lambda_4$$

and

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

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

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

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

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

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 $\tau_0 \coloneqq 0$ and $\tau_{n+1} \coloneqq T$. 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

$$X(\tau_i^-) = X(\tau_i^+), \quad i = 1,\dots,n \tag{5a}$$

$$\Lambda(\tau_j^-) = \Lambda(\tau_j^+), \quad j \in I_s \cup I_x \tag{5b}$$

$$\Lambda_{1,2,3,5}(\tau_j^-) = \Lambda_{1,2,3,5}(\tau_j^+), \quad j \in I_e \tag{5c}$$

$$\Lambda_4(\tau_j^-) = \Lambda_4(\tau_j^+) - \chi_j, \quad j \in I_e \tag{5d}$$

with

$$\chi_j \geq 0, \quad j \in I_e. \tag{5e}$$

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

$$\psi(t)(1-\gamma(t)) = 0 \tag{5f}$$

and

$$\psi(t) \geq 0, \quad \tau_j \leq t \leq \tau_{j+1},\ j \in I_e \text{ or } j+1 \in I_x \tag{5g}$$

$$\psi(\tau_j^+) = \chi_j, \quad j \in I_e. \tag{5h}$$

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

$$\dot{\psi}(t) \leq 0, \quad \tau_j \leq t \leq \tau_{j+1},\ j \in I_e \text{ or } j+1 \in I_x. \tag{5i}$$

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

$$\Lambda(T) = -K\frac{\partial}{\partial X}L(T)^\sigma\gamma(T)^\sigma \tag{5j}$$

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

$$\Lambda_4(T) = -K\frac{\partial}{\partial \gamma}L(T)^\sigma\gamma(T)^\sigma + \chi_T \tag{5k}$$

with

$$\chi_T \geq 0. \tag{5l}$$

Since the state space and control region are bounded, the conditions for the Fillipov-Cesari existence theorem hold, (see e.g. Seierstad and Sydsæter, 1987). But no sufficiency condition guarantees the uniqueness of the optimal solution and in fact it is one of the features of this model that multiple optimal solutions occur. In Appendix B the numerical approach is explained in detail that allows us to detect these solutions. In the sequel we refer to these solutions as 'optimal' since they are locally optimal and the numerical approach attempts to systematically consider all other candidate solutions, but there are in fact no sufficiency conditions or formal proof. So, we use the word 'optimal' to mean superior to any other solutions detected via this systematic search. The challenge of formally establishing optimality in non-convex dynamic optimization problems is considerable, and the particular challenges posed by SIR epidemic models are now an active area of research. See, for example, the recent contributions of Goenka et al. (2021).

Parameterization

Table 1 shows the base case parameter values. Most have been discussed already, but a few merit more explanation, with additional details available in Caulkins et al. (2020).

We set α equal to $\frac{1}{15}$ per day, corresponding to an average dwell time in the infected state of fifteen days.

Since the average length of stay in hospital is shorter for regular vs. critical care patients, about the proportion of hospitalized COVID-19 patients requiring critical care is greater than the proportion of all hospital beds are critical care beds, the constraint will be on critical care beds, not total hospital beds. So we make them the basis for $H_{\max}$. Tsai et al. (2020) estimate that in the U.S., 58,166 of the existing 84,750 ICU beds could be made available for treating COVID-19 patients. Given the U.S. population is about 330 million, that is 0.176 per 1,000 people. The model acts as if patients who need critical care at some point need that care throughout their 15-day dwell time in the I state, but CDC data suggest that the average time in hospital for those needing critical care is actually only about 12 days. So we set $H_{\max} = 0.0002$, which is approximately equal to $(15/12) * 176$ per 1,000.

There is not truly consensus about any of the key parameters, but the two for which the widest range of values seem plausible are the probability an infected individual needs critical care, p, and the social cost of a death, M.

Based on early CDC guidance and the literature generally, our sense was that the probability of needing hospitalization given a *detected* infection was around 15%, about 30% of those entering the hospital required critical care beds, and about 45% of those needing critical care died even if they received that care.

ξ_1 is the death rate per day for infected people who need critical care and receive it. If the death rate for such individuals over an entire infection is 45% and the average dwell time in the I state is 15 days, then the death rate per day is $\xi_1 = \alpha 45\%$, or about 3%.

ξ_2 is the additional, incremental death rate per day for infected people who need critical care but do not receive it. If the death rate for such individuals over an entire infection is 100% and the average dwell time in the I state is 15 days, then the incremental death rate per day is $\xi_2 = \alpha(1 - 45\%)$, or about 3.67%.

At one point it appeared that about half of all infections were detected, implying that the probability of needing a critical care bed given infection, p, might be about $50\% \times 15\% \times 30\% = 2.25\%$. We take that as our base case value.

In Alvarez et al. (2020) a premature death is valued at 20×GDP per capita. Kniesner et al.

(2012) use a much greater value of 150×GDP per capita. Hammitt (2020) provide multiple reasons why lower values may be preferred for analysis of COVID-19 in particular. For example, lower values would apply if one focused on years-of-life-lost, since most deaths are among the elderly, especially those with other pre-existing conditions. E.g., Richardson et al. (2020) report that the vast majority of those hospitalized for COVID-19 had prior serious comorbidities such as hypertension, obesity, and diabetes, to the extent that their estimated 10-year survival rate absent COVID-19 was only 53%. So we consider a range from 10× to 150×GDP per capita.

We set $\sigma = \frac{2}{3}$. The term $K(\gamma L)^{2/3}$ measures GDP per day - K is the constant that we assume to capture everything except labor. Therefore, $365K(\gamma L)^{2/3}$ equals the nation's GDP. Without loss of generality we set $K = 1$ and consider a wide range of values for M to study the relation between the values of lost work and lost lives.

3. Results

3.1. Results with base case parameters

For the base case parameters in Table 1 three qualitatively different lockdown strategies can be optimal depending on the value of M, which denotes the value of preventing a death due to COVID-19. Typical trajectories for γ, the level of employment, are shown in Panels (a)–(c) of Figure 1.

Regime I applies for smaller values of M; it has only one relatively brief lockdown early on to dampen the intensity of the epidemic (Panel (a)). In Regime II, for intermediate values of M, it is optimal to have two separate lockdowns, one early and one later, shortly before the vaccine gets widely deployed (Panel (b)). In Regime III, with larger values of M, there is just one lockdown, but it is sustained (Panel (c)). In this case, that 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 Figure 2, respectively. The lower three panels show the time evolution of two key components of the objective function, health costs from premature deaths and economic costs from unemployment. Both are inverted to show them as costs (so large values are bad). Economic costs are relatively small with the short lockdown (Panel (g)), but it is the health costs that are massive. So the problem would appear to those living through it to be primarily a health crisis.

Skipping over to the far right, Panel (i) shows that with a sustained lockdown, economic losses are very large (a 40% or greater loss of output for more than one year), and the health costs are nonetheless still substantial, but mostly constrained to the first year. People living

variable	value	description
α	$\frac{1}{15}$	reciprocal of average duration of the infection
β_1	0	minimum level of infection risk
β_2	0.2	increment in the level of infection risk
$H_{\max}$	2×10^{-4}	capacity of intensive care units
p	2.25×10^{-2}	probability that infected person needs critical care
M	$*$	social cost of a premature death due to COVID-19
K	1	coefficient on economic activity
Γ	365	reflects time required to return to full employment
f	0.45	impact of lockdown fatigue on infection risk
κ_1	0.15	rate of accumulation of fatigue
κ_2	0.2	rate of exponential decay of fatigue
σ	$\frac{2}{3}$	labor elasticity in Cobb-Douglas production function
ν	0	birth rate
μ	0	death rate (not caused by COVID-19)
μ_I	0	COVID-19 death rate
ζ	5000	parameter in the approximation of the max-function
ξ_1	0.03	death rate of infected individual in critial care
ξ_2	0.55/15	incremental death rate if IC capacity is exceeded
φ	0.001	rate by which recovered get susceptible again
c_l	1000	parameter in business shutting down costs
c_r	5000	parameter in business reopening costs
θ	2	exponent in the proportionality function $\beta(t)$
$S(0)$	0.999	initial susceptible population
$L(0)$	0.001	initial infected population
$R(0)$	0.001	initial recovered population
$\gamma(0)$	1	initial employment level
$z(0)$	0	initial lockdown fatigue

Table 1: Base case parameter values and initial state variable values.

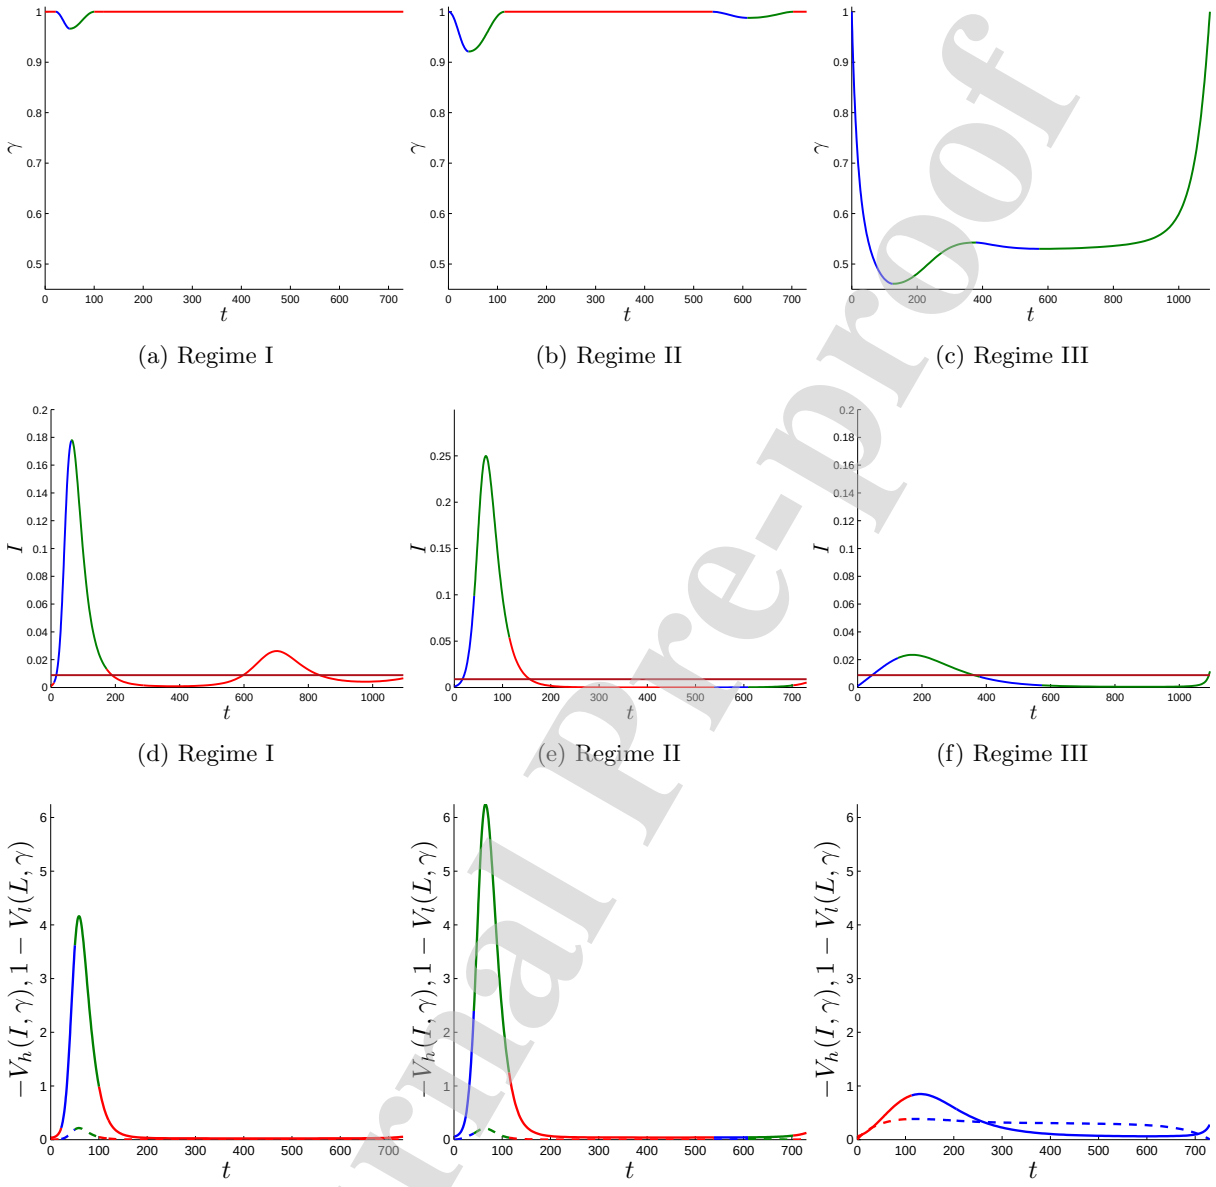

(a) Regime I (b) Regime II (c) Regime III

(d) Regime I (e) Regime II (f) Regime III

(g) Regime I (h) Regime II (i) 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. Panels (d), (e) and (f) show the number of infected and denote the ICU capacity as a horizontal line. Panels (g), (h) and (i) show the health (solid) and economic (dashed) costs.

through COVID times under that strategy would experience it as an acute health crisis that trigger a sustained economic slump.

With the particular double lockdown illustrated in Panel (e) there is only one (early) spike in infection and health costs; that is, the second lockdown is timed to preempt a resurgence in infections, not as a response to it. So in this model, a double lockdown can be optimal, but with these parameters it would not look like lockdowns reinstated in Europe in Fall of 2020, which were imposed grudgingly, only after infection rates had become quite high.

There are also notable differences in the duration of the costs. With a short lockdown, the pain reaches excruciating levels but is largely over within months. Conversely, the sustained lockdown imposes sustained (economic) pain, more or less right up to the time that a vaccine is widely deployed.

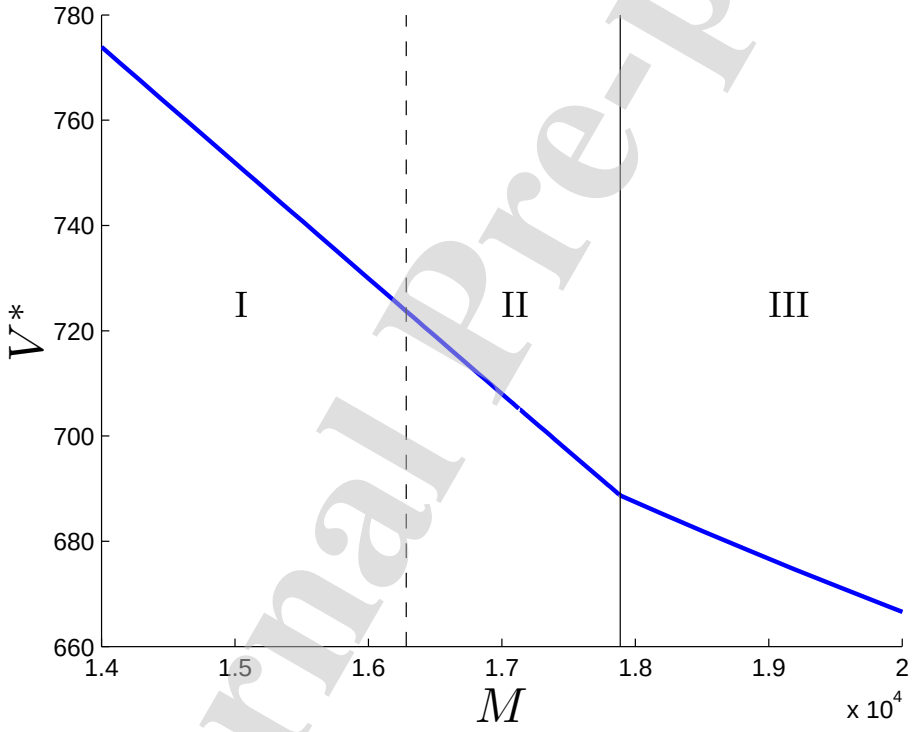

Figure 2: Dependence of the value function on the social cost of a death M for the base case parameters given in Table 1. There are three regimes which differ by the duration, intensity, and number of lockdowns of the optimal solutions. For the value of M highlighted by a solid vertical black line ($M = 1.7888 \times 10^4$) two different solution paths are optimal. At the dashed vertical black line the transition from region I to region II is continuous and the optimal solution is unique.

Naturally, as Figure 2 shows, the value function 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 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 does the dependence of V on M become less steep.

The solid vertical line in Figure 2 passing through the kink in $V(M)$ is a point at which two different strategies perform equally well; their lockdown intensities and corresponding health and economic costs are illustrated in Figure 3. When $M = 1.7888 \times 10^4$ the solid and dashed trajectories perform equally well overall even though they represent very different policies. The solid line is a double lockdown strategy that is very similar to the double lockdown strategy in Figure 1b; the dashed line shows a sustained and aggressive lockdown that suffers unemployment around 40% for more than a year but greatly reduces infections and deaths.

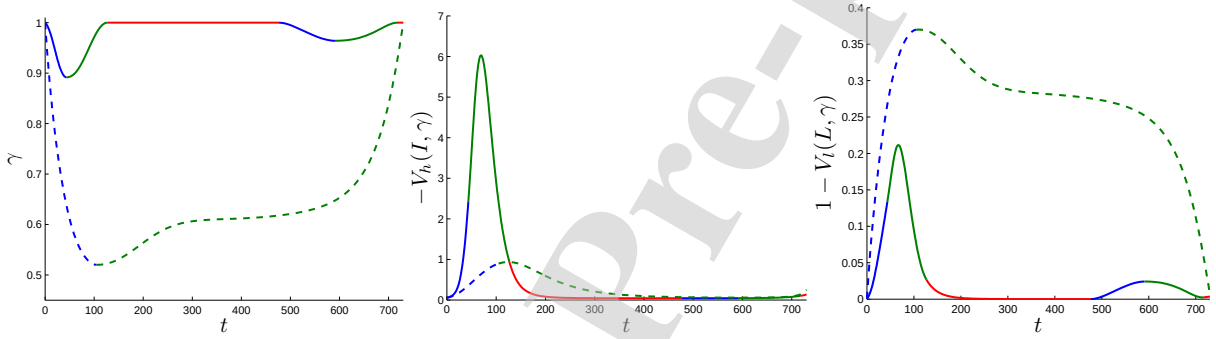

(a) (b) (c)

Figure 3: Optimal time paths for the Skiba solutions highlighted by the vertical black line in Figure 2. Panel (a) shows the proportion of employed people. Panel (b) and (c) show the health and economic costs.

Such points at which there are alternate optimal strategies are Skiba points (for an exact definition see Definition 1 in Appendix B). From the same initial point, two different optimal trajectories emerge.

Figure 3 shows the trajectory of the health (Panel (b)) and economic costs (Panel (c)). A political challenge of implementing the double lockdown strategy would be extremely intense health costs early on; literally, there would be people dying in the streets. Also Panel (c) shows that under this optimal double lockdown, there would be a significant second wave of economic costs (bump in Panel (c)) without a corresponding bump up in infection. That is, the second lockdown should be implemented before there is a second wave of infection in order to prevent that second wave. By contrast, what has been observed in reality in many countries is second lockdowns coming after there is already a significant second wave of infection. It is always hard

politically to implement painful measures to prevent something because the public does not ever see or experience that which has successfully been prevented. So people living through that second lockdown might be highly critical of the government imposing that second round of economic hardship, seemingly without cause.

There are also two conspicuous political challenges with the sustained lockdown strategy. First, there are still many infections and deaths, so the public would suffer severe economic hardship and also see large numbers of infections and deaths. Second, there would be significant lockdown fatigue, as indicated in Figure 4.

In particular, Figure 4 shows additional consequences of following those two very different strategies that are both optimal when $M = 1.7888 \times 10^4$. The double lockdown strategy (solid line) starts with a modest initial lockdown; that flattens the infection curve only moderately relative to the no-control scenario shown by the faint gray line). Relative to no control, at the epidemic's peak, the number who are infected at one time is about 25 percent lower, but that would still completely swamp hospital's treatment capacity. That is, not only does that 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.

Quite a few people still become infected with the sustained lockdown strategy, as can be seen in the dashed lines by the decline in the number of susceptibles (Panel (b)) and increase in the number of recovered individuals (Panel (d)), but the infections are spread out more over time.

It is interesting to contrast the two 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. The sustained lockdown strategy keeps this parameter value R_{eff} close to 1 throughout most of the time horizon, through a combination of economic shutdown and the roughly 30 percent reduction in the number of susceptibles produced by the initial wave of infections. In particular, because θ is 2, shutting down 40 percent of the economy would reduce the reproductive rate to $(1 - 0.4)^2$ or about one-third of its original value of 3.0, but because of lockdown fatigue, the decline is only to a little less than half. However, since the number of susceptibles is also about 30 percent lower, that leaves R_{eff} quite close to 1.0 because $3.0 * (1 - 0.3)/2$ is close to one.

With the double lockdown strategy, the infection rate never really stabilizes. Initially R_{eff} falls well below 1.0 primarily because of depletion of the stock of susceptibles, i.e., through herd immunity. Because of the backflow from the Recovered to the Susceptible state as infection-

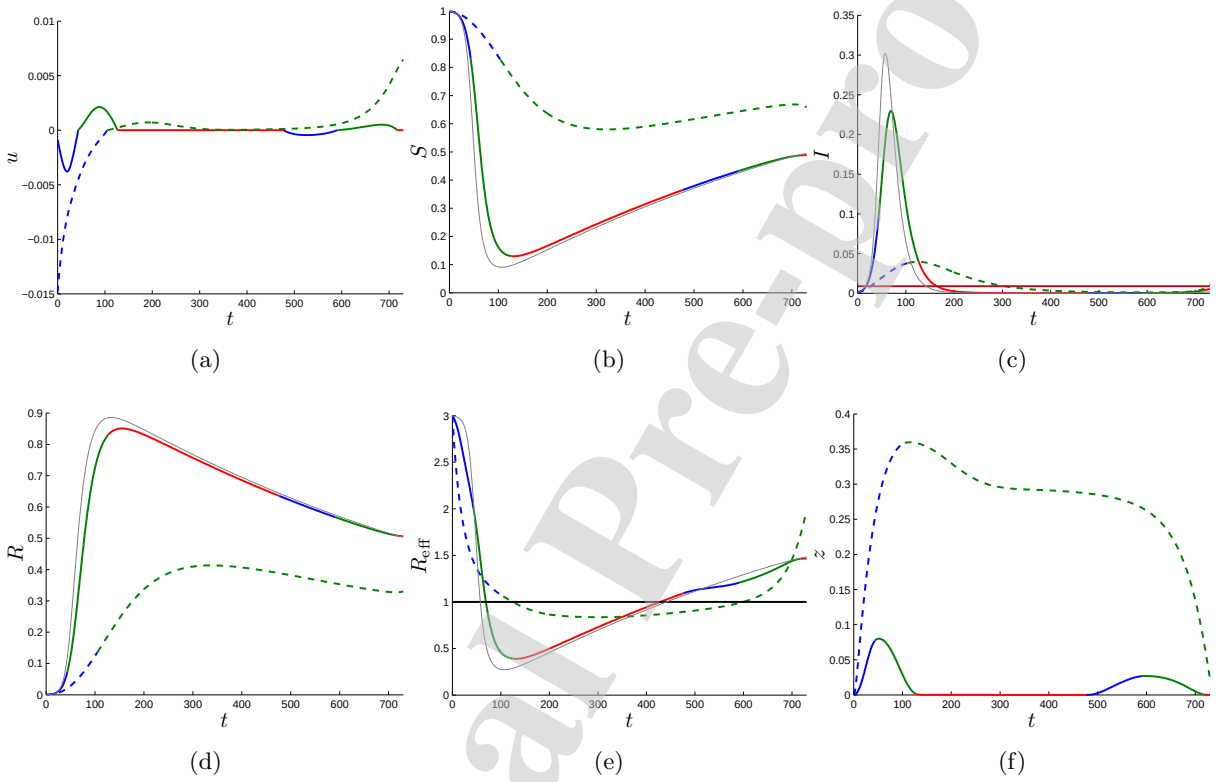

Figure 4: Time paths for the Skiba solutions at $M = 1.7888 \times 10^4$. Panel (a) depicts the control governing lockdown intensity. Panel (b) shows the number of susceptibles, Panel (c) the number of infected (above the red horizontal line hospital capacity is exceeded), Panel (d) the number of recovered patients. In Panel (e) the effective reproduction number can be seen and in Panel (f) the perceived lockdown intensity z. The gray line shows the uncontrolled epidemic's time path.

generated immunity wears off, the reproductive rate recovers to above one, but a second severe wave of infections is preempted, first by the second lockdown and then by the arrival of the vaccine (i.e., the end of the planning horizon of this problem).

The sustained lockdown strategy also 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 or approved, there is a lag between that and its mass production and widespread deployment. The production and distribution stages are reasonably well-understood, 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 the end of infectiousness 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.

One of the unique aspects of this model is its treatment of lockdown fatigue, meaning that over time a sustained lockdown loses its effectiveness as people become less compliant. Figure 5 shows how increasing or reducing the lockdown fatigue parameter from its base case value of $f = 0.45$ alters the threshold value of preventing a COVID-19 death that is necessary to make sustained lockdowns optimal. The upward slope of the line separating Region III from the other Regions shows that the weaker the lockdown fatigue effect, the more appealing sustained lockdowns become. In particular, eliminating the lockdown fatigue effect (setting $f = 0$) almost halves the threshold valuation in a death at which sustained lockdowns become optimal.

3.2. Triple Skiba points

The previous discussion focused on sensitivity analysis with respect to the social cost per premature death (parameter M), because there is not agreement as to the value of that parameter. Here we continue to vary M but also allow the virus' speed of spread, or infectivity, to vary from its base case value of $\beta_2 = 0.2$. We show that this produces a variety of interesting and complex behaviors.

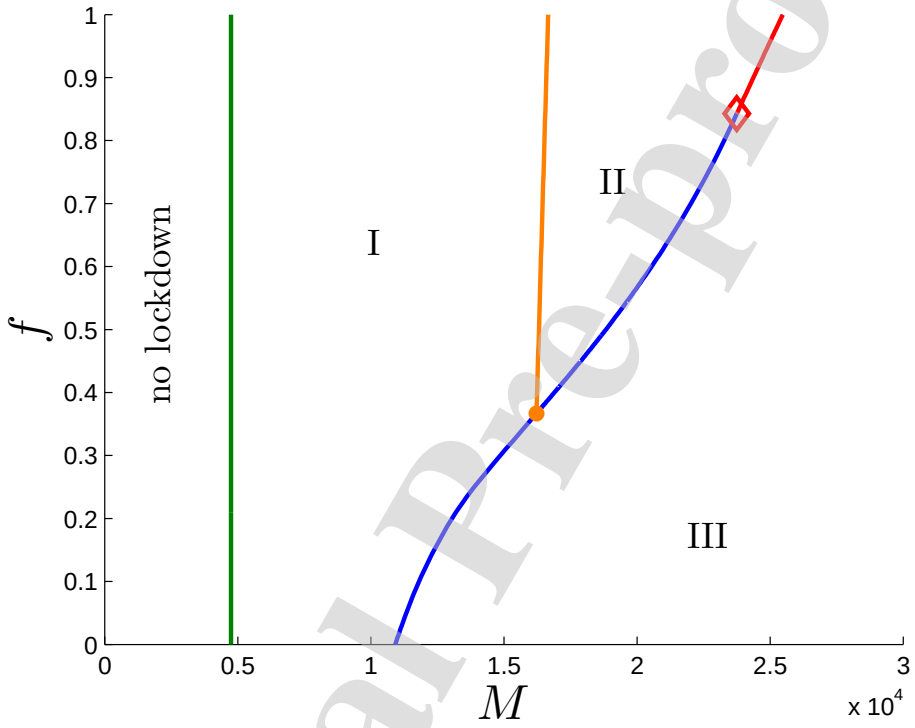

Figure 5: This figure shows the different regions in the $M - f$ space. The first region corresponds to optimal solutions without a lockdown. The green and orange lines denote continuous transitions from region "no lockdown" to I and region I to II. The blue curve is a Skiba curve, where the transition from region II to III is discontinuous and at the Skiba curve two optimal solutions exist. The Skiba curve switch to a continuous transition curve (red) at the red diamond.

One potential source of complexity is lockdown fatigue which might favor intermittent or pulsed lockdowns. However, modeling lockdown fatigue is not necessary to obtain interesting behaviors. To underscore that fact, in this section we set the fatigue parameter f equal to 0.

Figure 6 is a bifurcation diagram over the two parameters M and β_2. The Regions labeled 0, I, and II correspond to solutions with 0, 1, or 2 lockdowns, and Region III corresponds to one deep and sustained lockdown. The interleaving of the regions is much more complicated than in Figure 5.

To connect the two figures, note that the bottom of Figure 5 (when $f = 0$) corresponds to $\beta_2 = 0.2$ in Figure 6. With increasing M the solutions involve zero, one brief, or one sustained lockdown with transition points around M = 5,000 and 12,000, respectively.

For a more complex case, consider, for example, what happens with $\beta_2 = 0.256$. As M increases and one moves from left to right across the figure, one traverses successively through regions with one lockdown, then two, then one, then two again, and finally one long, sustained lockdown. That can happen because of the distinction between early lockdowns that address the initial explosive situation when nearly everyone is susceptible and late lockdowns that prevent a resurgence. In particular, Figure 7 shows that this sequence can be described in greater detail as: (1) One early lockdown to soften slightly that initial severe spike, (2) Adding a second, later lockdown to address resurgence, (3) Expanding the later lockdown but forgoing the initial lockdown, (4) Having both a forceful later lockdown and also an early lockdown, and finally (5) Two separate lockdowns are replaced by one continuous and substantial lockdown. The last strategy is the only one that prevents a goodly share of the population from becoming infected at some point.

There are conventional Skiba points throughout Figure 6, everywhere the curves are blue. In addition, there are four triple Skiba points (labelled T_1–T_4) where there are three distinct optimal solutions that produce the same objective function value.

Three of the triple Skiba points involve β_2 parameter values that are even greater than what is believed to describe COVID-19. Since COVID-19 has not such an unusually high reproductive rate, those are perhaps mostly of mathematical interest in the present context[2]. The first separates solutions involving one or two brief or one sustained lockdown. The second and third involve varying numbers of brief lockdowns, either early or late, but no sustained lockdown. Note: The segment of blue line separating two areas both labeled Region II indicates

[2]Note, however, that these points may be relevant when considering a more contagious disease

that it is possible to have points with two distinct optimal solutions, both of which involve two lockdowns.

We show the control trajectories for the last triple Skiba, T_4, because it occurs where $\beta_2 = 0.1$ which seems more likely to be seen in some future pandemic. As Figure 6 Panel (b) shows, all three optimal solutions emanating from that point involve one sustained lockdown, but they vary in their intensity. The mildest lasts a little over a year and peaks at about 10 percent forced unemployment (meaning $\gamma = 0.9$). The other two last for almost the entire time horizon and peak with closer to 20 percent forced unemployment. None involve cutting γ nearly as sharply as in the solutions discussed above because with a smaller $\beta_2 = 0.1$ the lockdown does not need to be as severe in order to push the reproductive rate down to 1.0.

This points to a quite interesting observation. When the virus is more virulent (higher β_2) one is less not more likely to want to pursue what amounts to "eradication" strategies because achieving that would require lockdowns that are too severe to sustain until a vaccine arrives. Sustained lockdowns are more appealing for less virulent pandemics when they involve laying off 5 or 10 percent of workers, not 40 or 50 percent. Lockdowns are just too blunt a tool to prevent a highly contagious condition from spreading throughout the population, at least if its infection fatality rate (IFR) is akin to that of COVID-19. (A higher IFR is effectively the same as a higher cost per death M in this model.)

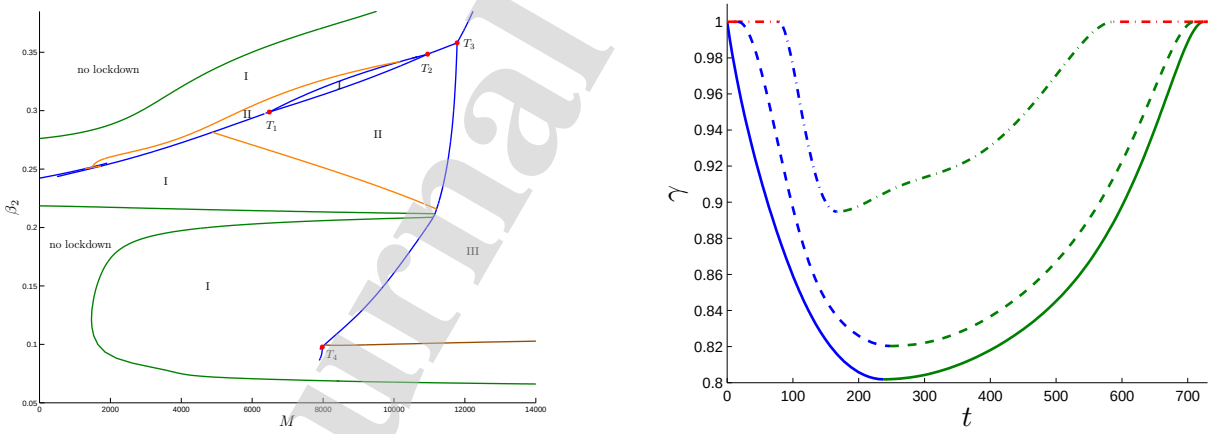

(a) (b) $M = 7959, \beta_2 = 0.0971$

Figure 6: Bifurcation diagram in the M–β_2 space, for $f = 0$ (Panel (a)). The blue curves show the Skiba curves and the dots denote the triple Skiba points (red), continuous transition of solutions (green) and boundaries of the different regions (black). Panel (b) shows the time paths of γ starting at the triple Skiba point T_4.

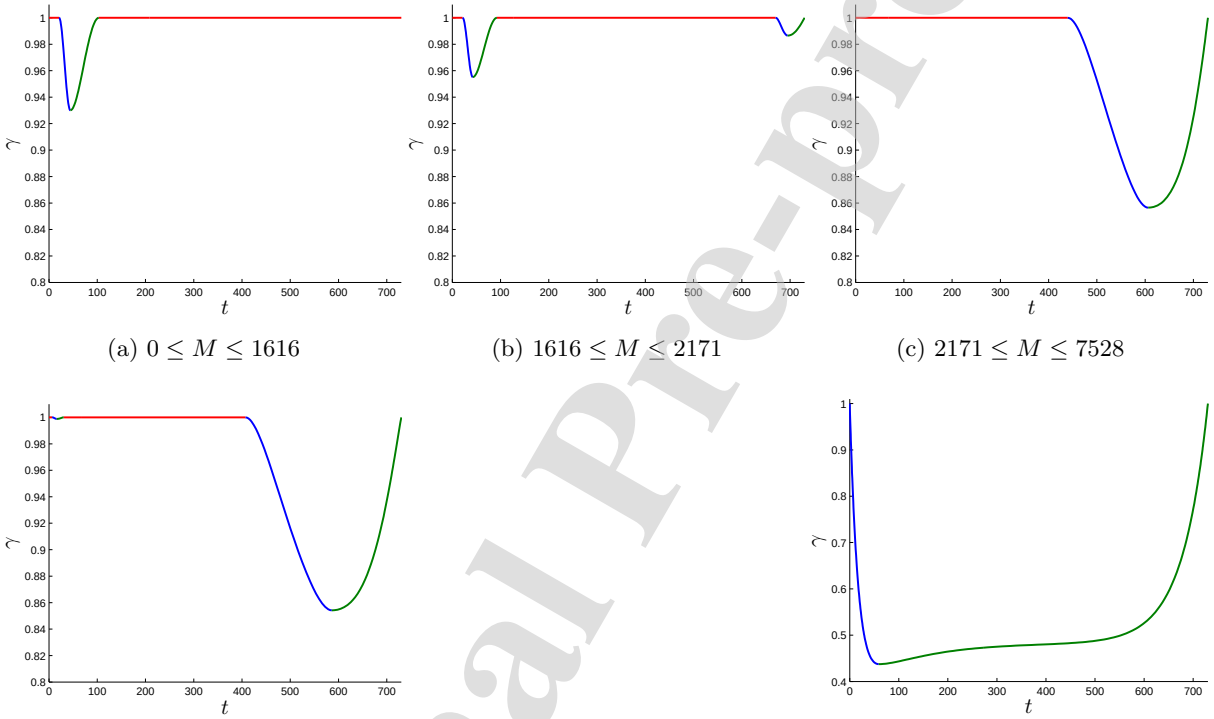

(a) $0 \leq M \leq 1616$ (b) $1616 \leq M \leq 2171$ (c) $2171 \leq M \leq 7528$

(d) $7528 \leq M \leq 11514$ (e) $11514 \leq M \leq \infty$

Figure 7: Showing the different transitions of the regimes for $\beta_2 = 0.256$ (a highly contagious epidemic) in terms of the time paths of γ.

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 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).

The literature suggests values for M ranging between 10 and 150 times annual GDP per capita, which translates to 3,650 up to 54,750 since we denominate the objective function in terms of GDP per day. Figure 5 shows that for our base case value of parameter $f = 0.45$ (standing for a modest degree of lockdown fatigue), varying parameter M much less than this, indeed only by a factor of 4 (from slightly below 5,000 up to 20,000), carries one all the way across the bifurcation diagram. When M is a bit smaller than 500, one is in Regime 0 where it is optimal to more or less let the epidemic run its course. When M is a bit larger than 5,000, it is optimal to have one lockdown. A second (also relatively brief) lockdown is added when M reaches about 16,000. And by the time M reaches 20,000, 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 and 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 had bounced back up into the top 5 as of late summer, with about 19 new confirmed infections per 100,000 per day. That appears to be a policy

disaster and, indeed, may indicate policy failure; certainly Prime Minister Netanyahu faced strident protests. But the model shows that the mere presence of a second lockdown is not in and of itself proof of error. A double lockdown can be an optimal strategy. That said, most of the second lockdowns that appear in optimal solutions to this model preempt a resurgence, not come after it.

A fourth observation concerns the Skiba points. Skiba points separate distinct optimal solution trajectories that spread out from a common initial condition. 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 Figure 6b, 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 between roughly days 100 an 550, we have the following odd situation in Figure 6b. 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 most times 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 multiple state variables, so projections onto a single dimension can be deceiving, and the objective function is a highly nonlinear function of the state variables. On the other hand, all of that nonlinearity arises 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.

5. Conclusion

In sum, this relatively simple model produces a wide range of interesting behaviors that are 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 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 extension would 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 as recently as mid 2019, 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 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 are 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 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 $0 \le f \le 1, \theta \ge 1$.

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

yields

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

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

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

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

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

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

The (in)equalitie signs in Eqs. (A.1d) to (A.1g) follow from

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

B. Numerical algorithm

Since the necessary conditions of problem (2) are not sufficient, the computed solutions that satisfy these conditions are just candidates for optimality. Ascertaining that one of these is actually optimal would require a strategy that both enumerates and considers all possible candidates. Hence, the challenge of the numerical procedure described here is the detection of all possible candidates such that the optimal solution can be determined among these candidates. The strategy we employ for attempting to enumerate all possibly relevant solutions is

to recognize that the number of possible solutions is restricted by the number of turning points detected during the numerical calculations.

Before we give a detailed explanation of the numerical procedure we have to introduce some notation. The state and costate variables are denoted as[3]

$$X \coloneqq (S, I, R, \gamma, z)', \quad \text{and} \quad Y \coloneqq \begin{pmatrix} X \\ \Lambda \end{pmatrix}.$$

If we refer to the state or costate values only we write

$$Y_X \coloneqq X \quad \text{or} \quad Y_\Lambda \coloneqq \Lambda.$$

The canonical system is given by

$$\dot{Y}(t) = \mathcal{C}(Y(t)), \quad t \in [0, T],$$

and the salvage value is abbreviated as $\mathcal{S}(X)$.

Let a solution path $Y(\cdot)$ consist of n arcs defined on the intervals $\tau_0 \coloneqq 0 < \tau_1 < \ldots < \tau_{n-1} < \tau_n \coloneqq T$, where τ_i, $i = 1, \ldots, n-1$, are called switching times.[4] This yields a multipoint BVP for the n arcs

$$\dot{Y}^{(i)}(t) = \mathcal{C}^{(i)}(Y^{(i)}(t)), \quad t \in [\tau_{i-1}, \tau_i], \quad i = 1, \ldots, n \tag{B.1a}$$

$$\mathrm{BC}(Y(\tau_0), \ldots, Y(\tau_i), \ldots, Y(\tau_n)). \tag{B.1b}$$

For the actual computation, the switching times τ_i, $i = 1, \ldots n-1$ have to be handled as free parameter values. In this way, the ODEs are transformed to the fixed time intervals $0 < 1 < \ldots < n-1 < n$. Therefore, BVP (B.1) is transformed into a BVP on fixed time intervals[5]

$$\dot{Y}^{(i)}(s) = \Delta^{(i)} \mathcal{C}^{(i)}(Y^{(i)}(s)), \quad s \in [i-1, i], \quad i = 1, \ldots, n \tag{B.2a}$$

$$\mathrm{BC}(Y(0), \ldots, Y(i), \ldots, Y(n)). \tag{B.2b}$$

with

$$t = \upsilon^{(i)}(s) \coloneqq \tau_{i-1} + (\tau_i - \tau_{i-1})(s - i + 1), \quad s \in [i-1, i], \quad i = 1, \ldots, n$$

[3] Note that we use a slightly different notation in this section for the state vector X.

[4] Due to structural changes of the solution this condition has to be relaxed since switching times can coincide. These changes are detected by our algorithm and handled accordingly, i.e. the BVP will be reformulated to follow the new solution structure.

[5] To keep notation simple we use the variable Y also for the time transformed variable.

and

$$\Delta^{(i)} := \frac{\mathrm{d}}{\mathrm{d}s}\upsilon^{(i)}(s) = \tau_i - \tau_{i-1}.$$

Let $(\tau_{i-1}, \tau_i), i = 1, \ldots, n$, be a specific interval. Then we note that this arc corresponds to one of the following three cases, with $t \in (\tau_{i-1}, \tau_i)$:

$$\begin{cases} u(t) < 0,\ \gamma(t) < 1 \\ u(t) > 0,\ \gamma(t) < 1 \\ u(t) = 0,\ \gamma(t) = 1 \end{cases}.$$

To distinguish between these cases we define a function arcid

$$\text{arcid}\colon \{1, \ldots, n\} \to \{0, 1, 2\}, \quad \text{arcid}(i) \in \{0, 1, 2\}.$$

This function assigns to each arc its specific type, i.e. negative control, positive control or active state constraint. This is necessary since each of these cases yields a different relation between costates and control. Moreover, the boundary conditions depend on the specific types and the structure of the arcs, as is explained in the next paragraphs.

In Table B.2 we formulate all possible and admissible boundary conditions for the switch between two arcs. Essentially, these conditions say that the state, most costate variables and the Hamiltonian are continuous at switching times. Only the fourth costate jumps if the state constraint becomes active on the second arc. The extent of the jump is χ. Obviously, the state constraint becoming active means that at the end of arc $\gamma = 1$. In Table B.3 the possible transversality conditions are stated. We formulate these conditions for the time transformed ODEs, i.e. on the fixed time intervals with $n \geq 1$. For a compact notation in Table B.2 the time argument $j \in \{1, \ldots, n-1\}$ is omitted in all expressions and $\mathcal{H}^{(l)} := \mathcal{H}^{(l)}(Y^{(l)}(j)),\ l = j, j+1$.

arcid(j)	arcid$(j+1)$	BC_X	BC_Λ	BC_τ	BC_χ
0/1/2	1/0/1	$Y_X^{(j)} = Y_X^{(j+1)}$	$Y_\Lambda^{(j)} = Y_\Lambda^{(j+1)}$	$\mathcal{H}^{(j)} = \mathcal{H}^{(j+1)}$	-
0	2	$Y_X^{(j)} = Y_X^{(j+1)}$	$Y_\Lambda^{(j)} = Y_\Lambda^{(j+1)} + \chi\eta$	$\mathcal{H}^{(j)} = \mathcal{H}^{(j+1)}$	$\gamma^{(j)} = 1$

Table B.2: Boundary conditions for the possible switches from arc j to $j+1$ with $\eta := (0, 0, 0, -1, 0)'$.

The admissibility conditions that a solution has to satisfy are summarized in Table B.4

For the actual computations we use a numerical continuation approach, see. e.g. Allgower and Georg (1990) and Kuznetsov (2004). Hence, instead of computing *the solution* for the

arcid(n)	$\gamma^{(n)}(n)$	BC_Λ	BC_χ
0/2	-	$Y_\Lambda^{(n)}(n) = \dfrac{\partial}{\partial X}\mathcal{S}(X^{(n)}(n))$	-
1	< 1	$Y_\Lambda^{(n)}(n) = \dfrac{\partial}{\partial X}\mathcal{S}(X^{(n)}(n))$	-
1	$= 1$	$Y_\Lambda^{(n)}(n) = \dfrac{\partial}{\partial X}\mathcal{S}(X^{(n)}(n)) - \chi\eta$	$\gamma^{(n)}(n) = 1$

Table B.3: Transversality conditions for the different possible cases.

arcid(j)	$u(s)$	$\gamma(s)$	$\psi(s)$	$\dot{\psi}(s)$	χ
0	< 0	< 1	-	-	-
1	> 0	< 1	-	-	-
2	-	-	≥ 0	≤ 0	≥ 0

Table B.4: Conditions that are monitored during the continuation process for every arc $Y^{(j)}(\cdot),\ j = 1, \ldots, n$ and $s \neq \tau_i, i = 0, \ldots, n$.

specific data we compute a *family of solutions* $Y(\cdot, \omega)$ satisfying the BVP

$$\dot{Y}^{(i)}(s, \omega) = \Delta^{(i)}(\omega)\mathcal{C}^{(i)}(Y^{(i)}(s, \omega), \omega), \quad s \in [i-1, i], \quad i = 1, \ldots, n \tag{B.3a}$$

$$\mathrm{BC}(Y(0, \omega), \ldots, Y(i, \omega), \ldots, Y(k, \omega)), \quad i = 1, \ldots, n. \tag{B.3b}$$

for a continuation parameter $\omega \in \mathbb{R}$.

One of the strengths of the continuation approach is its feature that the switching structure is revealed during the continuation process so that no prior information about this structure is needed. This is realized by checking the admissibility of the solution. If one of the conditions in Table B.4 is violated, the step width of the continuation process is reduced until a minimal step size is reached. Then the continuation process stops and the BVP (B.3) is changed to cover the new solution structure revealed from the type of violation. The numerical algorithm is implemented in the MATLAB toolbox `OCMat`[6], see also Grass et al. (2008).

B.1. How to find these solution

In this section we present the procedure for finding such a solution for the base case parameter values, specified in Table 1 with $M = 10000$, time horizon $T = 730$ and initial values $X_0 = (0.999, 0.001, 0, 1, 0)$. Whatever numerical approach we apply, we have to resolve two

[6] An older version of this toolbox can be downloaded from `http://orcos.tuwien.ac.at/research/ocmat_software`.

difficulties that are related. Solutions of the BVP (B.3) need not be unique and solutions that satisfy the necessary optimality conditions need not be optimal. The continuation approach allows us to develop a branch of solutions and therefore reveal the structural changes in the course of the parameter change.

To start the continuation we have to provide an initial solution. To do so we choose the time horizon T as continuation parameter, i.e. $T(\omega) \coloneqq \omega$ and start with $\omega = 0$. In the first instance we set $\gamma_0 = 0.5$ in order to avoid numerical difficulties with the state constraint. For this trivial problem, i.e. $Y(\cdot) \equiv (X_0, \Lambda_0)$ the costate values satisfy $\Lambda(\cdot) \equiv \Lambda_0$ and can therefore be derived from the transversality condition Eq. (5j)

$$\Lambda_0 = (153.3415, 0, 153.3415, 306.3764, 0)'.$$

This implies that the corresponding $u(\cdot) \equiv u_0 > 0$, see Eq. (3f) and hence the solution path starts with $\text{arcid}(1) = 1$. Summing up we get that the initial BVP is represented by

$$\dot{Y}^{(1)}(s) = T(\omega)\mathcal{C}^{(1)}(Y^{(1)}(s)), \quad s \in [0,1] \tag{B.4a}$$

$$Y_X^{(1)}(0) = (0.999, 0.001, 0, 0.5, 0)' \tag{B.4b}$$

$$Y_\Lambda^{(1)}(1) = \frac{\partial}{\partial X}\mathcal{S}(X^{(1)}(1)) \tag{B.4c}$$

$$\text{arcid}(1) = 1.$$

BVP (B.4) is regular for $\omega = 0$, since the linearization reduces to the non-singular Jacobian of the ODEs. Hence, a continuation process can be started guaranteeing for some $\varepsilon > 0$ the existence of a solution branch $(Y(\cdot, \omega), \omega)$ with $0 \leq \omega < \varepsilon$ and $T(0) = 0$.

The calculations show that the continuation process is successful until $T(\omega) = 11.0882$. Then the state $\gamma(T(\omega))$ hits the upper bound $\gamma(T(\omega)) \leq 1$. This means that BVP (B.4) does not describe an admissible solution for $T > 11.0882$ and we have to change the BVP taking care of the new solution structure.

For $T > 11.0882$ essentially there are two different possibilities. First, the solution path consists of two arcs, $i = 1, 2$ with $\text{arcid}(1) = 1$ and $\text{arcid}(2) = 2$, meaning that there exists $0 < \tau_1 < T(\omega)$, such that for $t \geq \tau_1$ the state constraint is binding. Second, the solution path behaves such that the state constraint becomes binding exactly at the endtime $T(\omega)$. In this case the solution consists of one arc and the transversality conditions have to be changed, see Figures B.8a and B.8b.

We note that the control value at $T(\omega) = 11.0882$ is larger than zero, cf. Figure B.8b. Since the control has to be continuous in the transition to an arc with active state constraint, this is

in contradiction to the first solution and implies that the solution for larger values of $T > T(\omega)$ will hit the state constraint at the end-time.

Therefore, the transversality condition Eq. (B.4c) of BVP (B.4) has to be changed, see Table B.3, into

$$\Lambda(1) = \frac{\partial}{\partial \gamma} \mathcal{S}(X(1)) + \chi \eta \tag{B.5a}$$

with

$$\chi \geq 0, \quad \eta \coloneqq (0, 0, 0, -1, 0)',$$

and

$$\gamma(1) = 1. \tag{B.5b}$$

Importantly, the last admissible solution of BVP (B.4) for $T = 11.0882$ is a solution of BVP (B.4) and transversality conditions (B.5) with $\chi = 0$.

Starting the continuation process of the end-time with this solution, it turns out that this solution type is admissible until $T = 730$ is reached.[7] Thus we finally computed a solution for the end-time $T = 730$ and the initial condition with $\gamma(0) = 0.5$, see Figures B.8c and B.8d.

To determine the solution with $\gamma(0) = 1$ we employ the continuation process with $\gamma(0)$ as the continuation parameter. The calculations show that for $\gamma(0) = 0.5388$ the initial control becomes zero, see Figures B.8e and B.8f. Therefore, a new BVP has to be formulated, where the solution path consists of two arcs $i = 1, 2$ such that arcid(1) = 0 and arcid(2) = 1. Thus,

[7] We omit that in the interval $T = (162.4571, 248.4864)$ the solution path consists of two arcs, where the first arc has arcid(1) = 0, i.e. the control is negative. Anyhow, for $T > 373.4438$ this arc vanishes and the solution consists of one arc satisfying BVP(B.4) with transversality conditions (B.5).

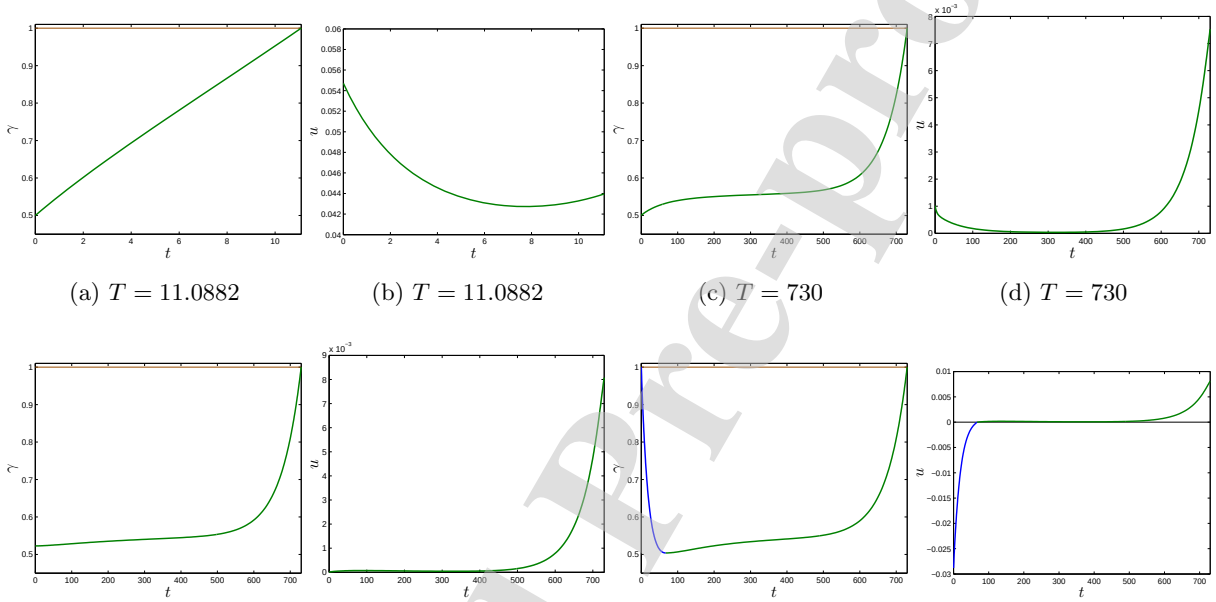

(a) $T = 11.0882$ (b) $T = 11.0882$ (c) $T = 730$ (d) $T = 730$

(e) $\gamma(0) = 0.5388$ (f) $\gamma(0) = 0.5388$ (g) $\gamma(0) = 1$ (h) $\gamma(0) = 1$

Figure B.8: This figure shows the different phases for the state path $\gamma(\cdot)$ and control path $u(\cdot)$ of the continuation with respect to T (in (a)–(d)) and the continuation with respect to the initial state of γ (in (e)–(h)). The first continuation with respect to T stops, when the state $\gamma(T)$ hits the state constraint, see (a) exhibiting a positive control value, see (b). The continuation process ends at $T = 730$, see (c)–(d). For the continuation with respect to the initial γ value the control hits zero at $\gamma(0) = 0.5227$, see (f) and the BVP has to be adapted accordingly. The final solution is reached for $\gamma(0) = 1$ in (g)–(h).

the new BVP writes as

$$\dot{Y}^{(1)}(t) = \tau_1 \mathcal{C}^{(1)}(Y^{(1)}(t)) \tag{B.6a}$$

$$\dot{Y}^{(2)}(t) = (720 - \tau_1)\mathcal{C}^{(2)}(Y^{(2)}(t)) \tag{B.6b}$$

$$Y_X(0) = (0.999, 0.001, 0, \omega, 0)' \tag{B.6c}$$

$$\Lambda_4(2) = \frac{\partial}{\partial \gamma}\mathcal{S}(X(2)) + \chi \tag{B.6d}$$

$$Y^{(1)}(1) = Y^{(2)}(1) \tag{B.6e}$$

$$\mathcal{H}(Y^{(1)}(1)) = \mathcal{H}(Y^{(2)}(1)) \tag{B.6f}$$

$$\gamma(2) = 1 \tag{B.6g}$$

$$\text{arcid}(1) = 0, \quad \text{arcid}(2) = 1. \tag{B.6h}$$

BVP (B.6) consists of two arcs that are connected at the free parameter τ_1 and state γ hits the state constraint at the final time $T = 730$. Since the state constraint does not become binding at the switching time τ_1, the states and costates are continuous, yielding Eq. (B.6e). Moreover, the Hamiltonian is continuous, yielding the boundary condition Eq. (B.6f). This equation is needed for the free switching time τ_1. The other boundary conditions are those from the previous BVP (B.4) with transversality conditions (B.5).

To start the continuation process with BVP (B.6) we have to transform the last admissible solution of the previous continuation. Let $Y_0^{(1)}(s),\ 0 \le s \le 1$ be this last admissible solution. Then we define for $\tau_1 = 0$

$$Y_1^{(1)}(s) := \begin{cases} Y_0^{(1)}(0) & 0 \le s \le 1 \\ Y_0^{(1)}(s) & 1 \le s \le 2. \end{cases}$$

This path trivially satisfies BVP (B.6) for $\tau_1 = 0$. The continuation process finally yields an admissible solution for $\gamma(0) = 1$, see Figures B.8g and B.8h.

We demonstrated that after a few continuation steps we were able to determine an admissible path satisfying the necessary optimality conditions starting from the trivial case $T = 0$. However, so far there is no assurance that this calculated path is the optimal solution.

To seek such assurance we use the same procedure and apply the continuation process with respect to the parameter M. We omit the numerical details, which are analogous to the previously described steps, in the sense that we check admissibility of the computed solution and perform necessary adaptations of the BVP to the changing solution structure if a violation occurs. Continuing with respect to M we find a solution branch that consists of several segments. Each segment corresponds to a different BVP, i.e. paths with a different number of arcs

and structure, see Figure B.9. During the continuation process four turning point bifurcations occur,[8] where the continuation process changes its direction. Thus, the solution branch intersects at several M values. One of these intersection points corresponds to a Skiba point, see Definition 1, depicted by the black vertical line in Figure B.9; other intersection points can be excluded as being optimal (gray vertical lines). Moreover, it is revealed that the solution found by the process depicted in Figures B.8g to B.8h is not optimal (gray dot). The solution that would appear to be optimal is found by the M-continuation (black dot), cf. Figure B.9.

Definition 1 (k-tuple Skiba point)**.** *A point* $(X_0, p_0) \in \mathbb{R}^n \times \mathbb{R}^p$ *(state-parameter-space*[9]*) is called a* k-tuple Skiba point *iff there exist* k *solution paths* $(X_i^*(\cdot, p_0), u_i^*(\cdot, p_0)),\ i = 1, \dots, k$ *of problem* (2) *satisfying*

$$\int_0^T \|(X_i^*(t, p_0) - X_j^*(t, p_0), u_i^*(t, p_0) - u_j^*(t, p_0))\|_2 \,\mathrm{d}t > 0, \quad i \neq j \tag{B.7a}$$

and

$$V(X_0, u_i(\cdot), p_0) = V(X_0, u_j(\cdot), p_0) = V^*(X_0, p_0), \quad i = 1, \dots, k. \tag{B.7b}$$

If $k = 2$ *the point is simply called a* Skiba point *and if* $k = 3$ *this point is called a* triple Skiba point*. The* k *solutions are called* Skiba solutions*.*

A similar definition can be found in Caulkins et al. (2015), where Skiba points are analyzed in free endtime problems.

Remark 1. *Condition* (B.7a) *states that the solutions have to be essentially different. From condition* (B.7b) *the optimality of the* k *different solutions follows. Note that for infinite time horizon problems, where Skiba points are usually described, condition* (B.7a) *is formulated as the condition that the solutions converge to different limit-sets (equilibria).*

For the calculation of a k-tuple Skiba point and its solutions let $Y_i(\cdot),\ i = 1, \dots, k$ be the k paths. Note that here the index i does not refer to a specific arc of the path but is used to distinguish different solution paths. Each of these paths may consist of multiple arcs,

[8] At these bifurcations, also called fold bifurcations, the tangent on the solution curve becomes "vertical", i.e. the change of the solution with respect to the continuation parameter becomes zero, see e.g. Kielhöfer (2012).

[9] Usually Skiba points are defined in the state space. In our model the initial states are fixed and the Skiba point appears in the parameter space. Therefore, we slightly generalize the definition of a Skiba point. The extension to a parameter dependent notation is straight forward, e.g. if p_0 is the parameter value $V(X_0, u(\cdot))$ becomes $V(X_0, u(\cdot), p_0)$.

i.e. $Y_i^{(j)},\ j = 1, \dots, n_i$. Each of these paths satisfy a BVP of the form

$$\dot{Y}_i(s) = \Delta_i \mathcal{C}(Y_i(s)), \quad s \in [0, n_i], \quad i = 1, \dots, k \tag{B.8a}$$

where Δ_i denotes the time transformation. Each solution has to satisfy the boundary conditions

$$\mathrm{BC}_i(Y_i(0), \dots, Y_i(n_i)) = 0. \tag{B.8b}$$

consisting of initial, transversality, and switching conditions. Condition (B.7b) implies that the following additional boundary conditions have to be satisfied

$$V(Y_l(0)) = V(Y_m(0)), \quad l \neq m,\ 1 \leq l, m \leq k, \tag{B.8c}$$

where $V(Y_i(0))$ is the objective value of the solution $Y_i(\cdot)$.[10]

This yields $k-1$ additional conditions. Therefore, $k-1$ model parameters have to be used as free parameter values. The formulation as a BVP allows the application of the continuation approach and hence e.g. the continuation of Skiba points, yielding branches of Skiba curves in the parameter space, cf. Figure 6a.

Remark 2. *For the continuation of a* regular[11] *k-tuple Skiba point the model has to include at least k parameter values. Given regularity in a $1D$ bifurcation diagram, where solution branches are computed for one free parameter value, at most (double) Skiba points can appear, see Figure B.9. For a $2D$ bifurcation diagram, where solution branches of Skiba solutions are computed, at most a triple Skiba point can appear at the intersection of these branches, see Figure 6a.*

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.

E. L. Allgower and K. Georg. *Numerical Continuation Methods: An Introduction.* Springer-Verlag, Philadelphia, PA, 1990.

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.

[10] The objective value can be calculated in parallel with the time paths by adding an additional ODE for the objective function. With this extension conditions (B.8c) become usual boundary conditions.

[11] A k-tuple Skiba point is regular if the corresponding BVP (B.8) is regular.

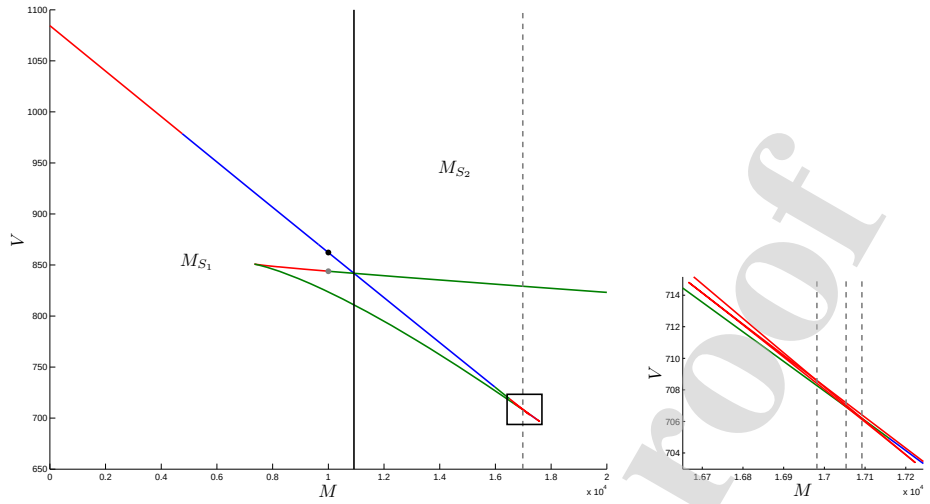

Figure B.9: This figure illustrates the result of the continuation process of extremal solutions with respect to M. It shows the objective value evaluated at each solution of the continuation process. The sequence of colors accentuates the different segments, where each segment of the branch correspond to a different BVP, due to a change in the solution structure. Same colors do not necessarily corresponds to the same solution structure. During the continuation process four turning points occur. Therefore, the curve representing the objective value intersects several times. At two of these intersection points (solid black and dashed gray line) the objective value is equal. The intersection point at the solid black line is a Skiba point and the intersection point at the dashed gray line can be excluded to be optimal. The gray dot denotes the initially calculated solution, which turned out to be non optimal. The black dot corresponds to the solution found during the continuation process with respect to M.

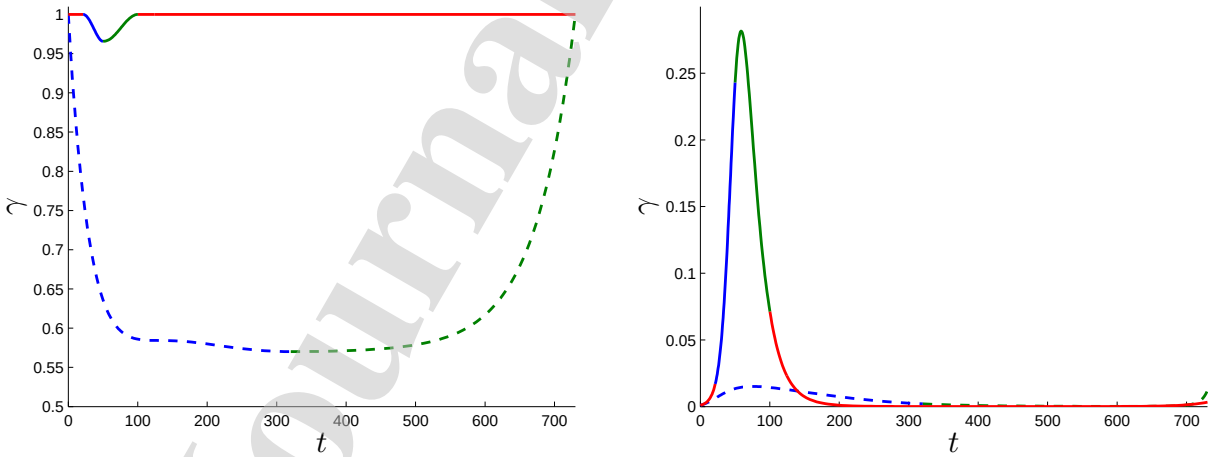

(a) $T = 11.0882$ (b) $T = 11.0882$

Figure B.10: This figure shows the time paths for the lockdown intensity ($\gamma(\cdot)$) and number of infected ($I(\cdot)$) of the extremal solution that has been found by the time and initial state continuation process, shown in Figure B.8 (dashed lines) and for the solution (solid lines) found by the continuation with respect to parameter M, see Figure B.9. The colors refer to the number of the according arcid, blue: 0, green: 1, red: 2.

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. Hartl, P. M. Kort, A. Prskawetz, A. Seidl, and S. Wrzaczek. How long should the COVID-19 lockdown continue? *PLOS ONE*, 15(12), 2020. doi: 10.1371/journal.pone.0243413.

A. Goenka, L. Liu, and M.-H. Nguyen. SIR economic epidemiological models with disease induced mortality. *forthcoming in Journal of Mathematical Economics*, 2021.

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.

J. K. Hammitt. Valuing mortality risk in the time of COVID-19. *SSRN Electronic Journal*, 2020. doi: 10.2139/ssrn.3615314.

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.

H. Kielhöfer. *Bifurcation Theory: An Introduction with Applications to Partial Differential Equations*. Springer-Verlag, 2012. doi: 10.1007/978-1-4614-0502-3.

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.

T. J. Kniesner, W. K. Viscusi, C. Woock, and J. P. Ziliak. The value of a statistical life: Evidence from panel data. *The Review of Economics and Statistics*, 94(1):74–87, 2012. doi: 10.1162/REST_a_00229.

J. Köhler, L. Schwenkel, A. Koch, J. Berberich, P. Pauli, and F. Allgöwer. Robust and optimal predictive control of the COVID-19 outbreak. *forthcoming in Annual Reviews in Control*, 50, 2020.

Y. A. Kuznetsov. *Elements of Applied Bifurcation Theory*. Springer-Verlag, New York, third edition, 2004.

S. Richardson, J. S. Hirsch, M. Narasimhan, J. M. Crawford, T. McGinn, K. W. Davidson, , and the Northwell COVID-19 Research Consortium. Presenting characteristics, comorbidities, and outcomes among 5700 patients hospitalized with COVID-19 in the new york city area. *JAMA*, 2020. doi: 10.1001/jama.2020.6775.

A. Seierstad and K. Sydsæter. *Optimal Control Theory with Economic Applications*. North-Holland, Amsterdam, 1987.

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.

T. C. Tsai, B. H. Jacobson, and A. K. Jha. American hospital capacity and projected need for COVID-19 patient care. *Health Affairs Blog*, March 2020. doi: 10.1377/hblog20200317.457910.

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.

Declarations of interest: none