

This is a repository copy of *Optimal Timing of Non-Pharmaceutical Interventions During an Epidemic*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/188120/>

Version: Published Version

Article:

Huberts, Nick Fijbo Dick and Thijssen, Jacco orcid.org/0000-0001-6207-5647 (2023)
Optimal Timing of Non-Pharmaceutical Interventions During an Epidemic. European
Journal of Operational Research. pp. 1366-1389. ISSN 0377-2217

<https://doi.org/10.1016/j.ejor.2022.06.034>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Optimal timing of non-pharmaceutical interventions during an epidemic[☆]Nick F. D. Huberts^{a,*}, Jacco J. J. Thijssen^{a,b}^a Management School, University of York, Heslington, York YO10 5ZF, United Kingdom^b Department of Mathematics, University of York, Heslington, York YO10 5ZF, United Kingdom

ARTICLE INFO

Article history:

Received 9 March 2021

Accepted 15 June 2022

Available online 22 June 2022

Keywords:

Decision analysis

Continuous-time Markov chains

SIRD model

COVID-19

Optimal stopping

ABSTRACT

In response to the recent outbreak of the SARS-CoV-2 virus governments have aimed to reduce the virus's spread through, *inter alia*, non-pharmaceutical intervention. We address the question when such measures should be implemented and, once implemented, when to remove them.

These issues are viewed through a real-options lens and we develop an SIRD-like continuous-time Markov chain model to analyze a sequence of options: the option to intervene and introduce measures and, after intervention has started, the option to remove these. Measures can be imposed multiple times.

We implement our model using estimates from empirical studies and, under fairly general assumptions, our main conclusions are that: (1) measures should be put in place not long after the first infections occur; (2) if the epidemic is discovered when there are many infected individuals already, then it is optimal never to introduce measures; (3) once the decision to introduce measures has been taken, these should stay in place until the number of susceptible or infected members of the population is close to zero; (4) it is never optimal to introduce a tier system to phase-in measures but it is optimal to use a tier system to phase-out measures; (5) a more infectious variant may reduce the duration of measures being in place; (6) the risk of infections being brought in by travelers should be curbed even when no other measures are in place. These results are robust to several variations of our base-case model.

© 2022 The Authors. Published by Elsevier B.V.

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

1. Introduction

Non-pharmaceutical interventions can have a significant impact on the rate at which a virus spreads during an epidemic (“flattening the curve”) as has been shown for previous influenza outbreaks, (see, e.g., Cowling et al., 2009; Hatchett, Mecher, & Lipsitch, 2007; Kamigaki & Oshitani, 2009; Stern & Markel, 2009; Wu et al., 2010).¹ These are, therefore, seen as a vital element of government policy (Ferguson et al., 2005; Leung & Nicoll, 2010). Recent

contributions by, e.g., Acemoglu, Chernozhukov, Werning, & Whinston (2020), show that social and economic interventions can have a significant impact on public health and can significantly reduce economic damage. However, such interventions, ranging from social distancing to complete lockdown, come at a significant cost of their own (Anderson, Heesterbeek, Klinkenberg, & Hollingsworth, 2020; Atkeson, 2020). These costs encompass obvious economic costs associated with a partial or complete lockdown, but also health-related costs linked to, e.g., reduced mental health levels in the population, missed diagnoses due to reduced availability of health-care workers, etc. In addition, there are one-off sunk costs of imposing a lockdown, related to, e.g., individuals’ and institutions’ adoption of homeworking. According to the Financial Times (24 August 2020), office owners can expect millions in extra costs to adhere to government guidelines.² An additional difficulty for policy makers who have to balance the costs and benefits of non-

[☆] The authors like to thank Job Thijssen, the participants of the webinars organized by the York Mathematical Biology Seminar Series (May 2020) and The OR Society (December 2020), and the participants of the workshop “Timing Uncertainty in Economics and Finance” in Bielefeld (June, 2020) for their comments. Finally, we are grateful for the many helpful comments that we received from the Editor and three anonymous referees. This paper is dedicated to the memory of Ulco de Jong; teacher, mentor, and *bon vivant*.

* Corresponding author.

E-mail addresses: nick.huberts@york.ac.uk (N.F.D. Huberts), jacco.thijssen@york.ac.uk (J.J.J. Thijssen).

¹ A similar impact has been reported for COVID-19 lockdowns, see, e.g., Talic et al. (2021), the UK government’s report “Analysis of the health, economic and so-

cial effects of COVID-19 and the approach to tiering” (No. 2020), and <https://www.economist.com/graphic-detail/2020/07/04/how-speedy-lockdowns-save-lives>.

² <https://www.ft.com/content/ece1bda0-8719-4987-8428-0f240748c5bb>

pharmaceutical intervention, is that the evolution of an epidemic is uncertain. Their cost-benefit analysis has to take this uncertainty explicitly into account.

The trade-off between costs and benefits in an uncertain environment makes the *timing* of the introduction and subsequent relaxation of non-pharmaceutical interventions a crucial policy decision. A consequence of the uncertainty over a virus's evolution is that it is not straightforward to determine *when* measures should be implemented and once implemented, *when* they should be relaxed. If measures are relaxed too early, a second outbreak could occur. Interventions that last too long lead to unnecessary economic and (mental) health damage. Our contribution is to develop an appropriate model and to apply dynamic programming techniques to find the optimal intervention policy of a social planner with an application to the COVID-19 epidemic. More specifically, this paper builds a continuous-time Markov chain (CTMC) of an SIRD-like³ model of non-pharmaceutical interventions (which we refer to as “lockdown”) to study the questions of (i) *when* to enter a lockdown and, consequently, (ii) *when* to exit it, under an uncertain evolution of the epidemic, by taking into account the social (health and non-health) costs of both the epidemic and the lockdown. We, therefore, not only contribute to the literature on policy making in the context of COVID-19 but our application also relates to the long tradition in OR to use CTMCs for modeling operational systems. We then calibrate our model using parameter estimates from the literature to arrive at the optimal (dynamic) intervention strategy. From this base-case scenario we then proceed to study the robustness of our policy conclusions, thereby providing a solid theoretical basis for policy evaluation by public health officials.

As pointed out by Tsekrekos & Yannacopoulos (2016) the OR community is well equipped to deal with problems where the objective is to determine the optimal timing decisions to commence and/or terminate a process or operations, when the underlying problem is subject to uncertainty. An established wide range of studies has shown how optimal stopping problems can be applied.⁴ For a review and discussion of the real option valuation literature in OR we refer to Trigeorgis & Tsekrekos (2018).

We consider a model where a social planner can impose a sequence of lockdowns. Our model views the timing issues through a real options lens: for each lockdown the social planner has two, nested, options, which we value under the assumption that the social planner wants to minimize the expected present value of total costs related to the virus and the lockdown. Briefly stated, our main conclusions are that:

1. lockdown should be entered not long after the first infections occur;
2. if the epidemic is discovered too late, i.e., when there are too many infected individuals already, then it is optimal never to introduce a sequence of measures;
3. once the decision to introduce measures has been taken, these should stay in place until the number of susceptible or infected members of the population is close to zero;

³ The acronym SIRD refers to the four groups members of the host population can be assigned to at each point in time: susceptible, infectious, recovered, and deceased.

⁴ This includes applications to, e.g., energy systems (Böckman, Fleten, Juliussen, Langhammer, & Revdal, 2008; Ernstsén & Boomsma, 2018; Muñoz, Contreras, Caamaño, & Correia, 2011; Steffen & Weber, 2016; Thompson, Davison, & Rasmussen, 2004; Tseng & Barz, 2002), infrastructure (Thijssen, 2022; Wang, Liu, Zhang, & Li, 2019), climate change adaptation (Truong & Trück, 2016), natural resources and commodities (Devalkar, Anupindi, & Sinha, 2011; Felix & Weber, 2012; Firoozi & Merrifield, 2003; Støre, Fleten, Hagspiel, & Nunes, 2018), forest plantations (Kallio, Kuula, & Oinonen, 2012), adaptive clinical trials (Thijssen & Bregantini, 2017), and IT and technology (d'Halluin, Forsyth, & Vetzal, 2007; Pendharkar, 2010).

4. it is never optimal to introduce a tier system to phase in measures but it is optimal to use a tier system when exiting a lockdown;
5. a more infectious variant may not only reduce the probability of lockdown being entered so that, before it is imposed, lockdown is less desirable as the virus spreads, but could also lead to a reduction in lockdown duration; and
6. lockdown is only optimal when combined with a strict policy at the country's borders that aims to minimize the risk of infections being brought in by travelers.

These results are, qualitatively, robust to several variations of our base-case model.

Some of the results are rather counter-intuitive when compared to the standard real options literature. The third result confirms the usual intuition that when switching between regimes involves sunk costs, then it is optimal to delay taking the decision. So, once lockdown is entered, one should wait until one is certain enough that the benefits of lifting lockdown are sufficiently higher than the associated costs (also see, e.g., Keogh-Brown, Wren-Lewis, Edmunds, Beutels, & Smith, 2010; Smith, Keogh-Brown, Barnett, & Tait, 2009). The first result is less intuitive. One might expect that the decision to enter lockdown should only be taken when one is sufficiently sure that its benefits outweigh its costs and, therefore, that the decision is taken when “enough” members of the population have been infected. However, our results belie this intuition in the sense that already after few infections it is optimal to go into lockdown. The reason is that the benefits from switching arise from the difference in the rate of infection and the expected evolution of the number of infected, which is greatest at the start of the epidemic (also see, e.g., Farboodi, Jarosch, & Shimer, 2021; Shin, 2020). This is supported by our finding that when the rate at which infections occur from ‘outside’ is higher, the decision maker wants a few more infections to occur before going into lockdown is optimal. When there is not enough time left for benefits to be realized, as is the case when a lot of members of the population have already been infected and/or recovered, then it is never optimal to incur the costs of lockdown, which explains the second result. Even when considering constraints on health care capacity, it is unlikely that lockdown is optimal when many individuals have had been infection. Our result is supported by Nova Scotia's (Canada) success with its strategy involving early and harsh measures to keep infections low in the COVID-19 epidemic (The Globe and Mail, 18 January 2021).⁵ The same intuition feeds into our fifth result. An increase in the infection rate has several implications. When the virus spreads more quickly, the expected time between infections reduces. As a result, the cost of staying in lockdown may no longer outweigh its benefits, especially when a higher proportion of the population has had the virus. This increases the incentive to come out of lockdown earlier when the infection rate is higher.

COVID literature The OR literature has seen various contributions in the context of the recent COVID epidemic. Examples of empirical studies include the use of predictive analytics tools for forecasting and planning (Nikolopoulos, Punia, Schäfers, Tsinopoulos, & Vasilakis, 2021), studying the effectiveness of social media as a humanitarian response (Kumar, Xu, Ghildayal, Chandra, & Yang, 2021b), establishing the socio-economic impact of the epidemic (Amaratunga et al., 2021), how the Chinese government can stimulate consumption (Liu, Shen, Li, & Chen, 2021), studying the role of social learning on closure decisions by firms (de Vaan, Mumtaz, Nagaraj, & Srivastava, 2021), providing evidence that countries that did not experience SARS in 2003 delayed action (Ru, Yang, & Zou, 2021), forecasting mortality rates in the USA (Taylor & Taylor,

⁵ <https://www.theglobeandmail.com/opinion/editorials/article-go-early-go-hard-stay-vigilant-how-atlantic-canada-beat-the-virus/>

2021), addressing the capacity planning decisions of a hemodialysis clinic in Istanbul (Bozkir et al., 2021), the use of an integrated epidemics-testing allocation model to minimize infections (Abdin et al., 2021), and evaluating and optimizing social distancing policies (Chen, Pun, & Wong, 2021).

Theoretical work related to health care and infections includes studies that, for example, develop a model to estimate and study local coronavirus outbreaks (Chang & Kaplan, 2021) or to study the spread of infections in an agent-based model (Ghaderi, 2022), use matheuristic algorithms for resource planning problems in a home health care (Nikzad, Bashiri, & Abbasi, 2021), find facility layouts that minimize risk of infections (Fischetti, Fischetti, & Stoustrup, 2021), apply an efficient large neighbourhood search algorithm to a ‘contagious disease testing’ problem (Wolfinger, Gansterer, Doerner, & Popper, 2021), and address vaccine and testing kit allocations in a multi-agent sequential decision problem (Thul & Powell, 2021). Economic impact is included in work that, for example, uses game theoretical mechanisms to study supply chain networks with labor constraints (Nagurney, 2021), studies pricing during disruptions (Feng, Rong, Shen, & Snyder, 2020), finds out how minimize disruptions to customer service level due to infrastructural deficiencies (Sinha, Kumar, & Chandra, 2021), and optimizes the implementation of measures with a reaction-diffusion process to minimize the economic burden (Rezapour, Baghaian, Naderi, & Sarmiento, 2021). The social planner is included in the theoretical work that, for example, looks at communication strategies by governments to induce compliance with measures (de Véricourt, Gurkan, & Wang, 2021), establishes how test accuracy and availability impact demand and thereby the social outcome by a social planner (Drakopoulos & Randhawa, 2021), and reduces social planner’s cost through optimal policies of safety stock and capital reserve (Zhang, Shi, Huang, Hua, & Teunter, 2021).

Other contributions include the study by Silal (2021), who discusses opportunities for OR to contribute to infectious disease management and to improve health outcomes (also see Choi, 2021). Tippong, Petrovic, & Akbari (2021) review OR applications in emergency medical response coordination in disaster management. Our contribution to the OR literature is to look at the optimal timing decisions of (non-pharmaceutical) interventions.

Other recent theoretical work on optimal policy within an epidemiological framework uses *optimal control theory*, sometimes incorporating economic trade-offs (see, e.g., Caulkins et al., 2021; Djidjou-Demasse, Michalakis, Choisy, Sofonea, & Alizon, 2020; Garriga, Manuelli, & Sanghi, 2020; Kantner & Koprućki, 2020; Piguillem & Shi, 2020; Toxvaerd, 2020). However, only a handful of papers incorporate optimal timing, including Alvarez, Argente, & Lippi (2020); Farboodi et al. (2021); Patterson-Lomba (2020); Shin (2020); Zhang & Enns (2020), and Kruse & Strack (2020). Their findings are generally mixed. For instance, Patterson-Lomba finds that the timing strongly depends on the natural reproduction number (R_0) and Kruse and Strack find that the social distancing can be delayed, whereas some others find that it is optimal to intervene when the first cases of infectives have been confirmed. The study by Federico & Ferrari (2020) is closest to our set-up, because they also incorporate uncertainty in an SIR-like model. However, in their model only the infection rate is stochastically evolving over time whereas in our model all transitions are stochastic. In addition, Federico and Ferrari do not consider sunk costs of switching, which leads to measures being continuously adjusted in the optimal policy. Therefore they find that immediate lockdown is not optimal. Kumar, Choi, Wamba, Gupta, & Tan (2021a) show empirically that rigid measures are less effective to combat infections than moderate long-lasting preventive measures in an SEIR model. They argue that harsh measures lead to more economic damage and only delay the peak. We allow the social planner to switch between rigid and moderate measures and find, indeed, that moder-

ate measures may be better during the end-phase. However, especially when incorporating healthcare constraints, delaying the peak might be better despite the economic damage. Other relevant and related studies include Jia & Chen (2021) who develop an SEIAR model with human uncertainty factors, Bliman, Duprez, Privat, & Vauchelet (2021) who look at optimal immunity control with the aim to minimize the epidemic final size in an SIR model, and Doyle (2021) who predicts the spread of infections in a SEIRD model with heterogeneous agents.

Although some studies have highlighted that any optimal policy is highly sensitive to parameterizations (e.g., Manski & Molinari, 2021, also see Avery, Bossert, Clark, Ellison, & Ellison, 2020 for a discussion), we argue that this is only partly true of our main results. When varying the parameters associated with the process, our main conclusions hold irrespective of the specific parameter values. Although our results are qualitatively the same for a wide range of parameter values, we do not want to underestimate the importance of accuracy when it comes to estimation of the infection rate, mortality rate, and immunity.

The remainder of this paper is organized as follows. Section 2 introduces the mathematical framework. Section 3 solves the social planner’s problem for our benchmark model, which is analyzed and extended in Section 4. Section 5 concludes and an online appendix contains the proofs of propositions as well as further robustness checks.

2. Model description

Our model considers the problem of a social planner in the context of an infectious disease outbreak. The social planner needs to decide when to undertake action and temporarily intervene in an environment where infections of members of the host population result in a (collateral) social cost. Temporal intervention could range from self-isolation to a complete lock-down. In this paper, we consider two models: a simplified benchmark model and a more realistic extended model. This allows us to identify the impact of each model characteristic. For our benchmark model, we assume that the form of intervention is exogenously given and that the costs and benefits (in terms of reduced infection/enhanced recovery) are known. Once temporal intervention has started, the social planner has to decide when to stop it. After that second decision, the social planner has no further decisions to take. In our extended model, we allow the social planner to impose a sequence of lockdowns and could opt to phase in and/or phase out measures.

For each lockdown, the social planner is assumed to hold two nested options, which we value under the assumption that she wants to minimize the expected present value of total social costs related to the virus and the lockdown over the planning horizon, which we assume to be infinite here. These costs are assumed to include (health and non-health) costs related to the number of infected people, the (health and non-health) cost of intervention, and the costs of lost lives and lost QALYs⁶ due to the disease. For simplicity, throughout this paper, we refer to intervention as “lockdown”, although, of course, interventions could entail different types of non-pharmaceutical measures.

The time parameter is continuous and is measured in days. Let $N \in \mathbb{N}$ be the size of the host population. Our model describes a continuous-time Markov chain (CTMC) in which at any time $t \geq 0$, the N individuals in our population are split into four groups, as in a typical SIRD model: I_t members are *infected*, R_t members are *recovered*, D_t members are *deceased*, and $S_t = S(I_t, R_t, D_t) := N - I_t - R_t - D_t$ members are *susceptible*. Initially, all individuals

⁶ Quality Adjusted Life Years is a commonly-used unit of measurement for the benefits of health technologies.




Fig. 1. SIRD model with transition rates.

are susceptible, after which they might move to be infected, after which they will either recover or die in finite time (a.s.).⁷ Our CTMC has state dependent transition rates, in which the interarrival times of *events* (i.e., infections, recoveries, and deaths) are exponentially distributed. In our benchmark model we assume that there are only three groups, setting $D_t = 0$ for all t . In that case, R_t represents both deceased and recovered individuals. Since for the benchmark case we do not have group “D”, the model can be referred to as an SIR-type of model.

When the social planner decides to intervene, a lockdown starts, which is assumed to decrease the infection rate and/or increase the recovery rate. We, therefore, initially, define two policy modes $m = 0, 1$. At time $t = 0$, the system is in $m = 0$, where no action is undertaken, and the social planner has an option to move to mode $m = 1$ where a lockdown is in place. In our extended model we assume that there is an intermediate mode where only mild measures are introduced. The social planner, then, has the option to either switch to the intermediate mode first before imposing lockdown but could also choose to skip this mode. Her objective is to minimize expected costs, which consist of *sunk costs*, associated with switching between modes, and (state-dependent) *running costs*. We assume that all costs are monetized.

Let $\mathbb{N}_N = \{0, 1, \dots, N\}$. The state space of our CTMC is given by

$$E = \{(i, r, d) \in \mathbb{N}_N^3 \mid i + r + d \leq N\}.$$

Here, i, r , and d denote the number of infected, recovered, and deceased members of the population, respectively. For $(i, r, d) \in E$, the number of susceptible members is given by $s = N - i - r - d \geq 0$. That is, the population is closed.⁸ We denote by $\text{int}(E)$ all points $(i, r, d) \in E$ such that $i + r + d < N$ and $i, r, d > 0$.

In our CTMC the number of states the Markov chain can transition to is at most three: either a susceptible member becomes infected so that the Markov chain moves from (i, r, d) to $(i + 1, r, d)$, or an infected member recovers so that the Markov chain moves from (i, r, d) to $(i - 1, r + 1, d)$, or an infected member dies, which implies a move from (i, r, d) to $(i - 1, r, d + 1)$. The first can only happen if $i + r + d < N$ and the second and third can only happen if $i > 0$. This implies that any state $(0, r, d)$ such that $r + d = N$ is always an absorbing state. Fig. 1 summarizes all possible transitions in our SIRD-like model.⁹

We shall proceed with formulating the full model description for our SIR-like benchmark model for which there are only two

modes and for which lockdown can only be imposed once. In our extended, SIRD-like, model, we reformulate where necessary. In addition, for simplicity, we first focus on the simplified case where $D_t = 0$ for all t . In that case we can think of the state space as

$$E = \{(i, r) \in \mathbb{N}_N^2 \mid i + r \leq N\}.$$

The interarrival times between state transitions are assumed to be exponentially distributed with mode- and state-dependent infection and recovery rates $\lambda_m(i, r)$ and $\mu_m(i, r)$, respectively. In order to stay close to standard SIR models, we assume for each mode $m = 0, 1$, that

$$\lambda_m(i, r) = \beta_m i \frac{S(i, r)}{N}, \quad \text{and} \\ \mu_m(i, r) = \gamma_m i,$$

respectively, for some $\beta_m, \gamma_m \in \mathbb{R}_{++}$. Note that the epidemic is over when the number of infectives is equal to 0, resulting in the process being absorbed. For our specification, the set of absorbing states of the chain is, therefore, given by

$$\Lambda = \{(i, r) \in E \mid i = 0\}.$$

In our extended model we consider the possibility of infections occurring as long as $S_t > 0$, i.e., even when $I_t = 0$. This could happen, for instance, when infections can take place from neighboring countries. In that case, the set of absorbing states is redefined as all states such that $i = 0$ and $s = 0$.

Let $X_t := (I_t, R_t) \in E$ represent the state of the system at any time $t \geq 0$, so that $X := (X_t)_{t \geq 0} = (I_t, R_t)_{t \geq 0}$ is our stochastic process. Note that if, at the outset, the process is in a state in Λ , it stays there forever and therefore, lockdown will never be imposed. For that reason, we will only consider scenarios where $X_0 \notin \Lambda$.

For $m = 0, 1$, we define the Markov generator $\mathcal{Q}^m = (q_{xy}^m)_{x, y \in E}$ with, for all $x, y \in E, y \neq x$,

$$q_{xy}^m = \begin{cases} \lambda_m(x_1, x_2) & \text{if } y - x = (1, 0), \\ \mu_m(x_1, x_2) & \text{if } y - x = (-1, 1), \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

and $q_{xx}^m = -\sum_{y \neq x} q_{xy}^m$. Note that when $y - x = (1, 0)$ an infection occurs and when $y - x = (-1, 1)$ an infected member recovers. We then define the (infinitesimal) transition function for $m = 0, 1$ as

$$p_{xy}^m(h) = \begin{cases} hq_{xy}^m + o(h) & \text{if } y \neq x \\ 1 + hq_{xx}^m + o(h) & \text{if } y = x \end{cases}, \quad h > 0,$$

where $o(h)$ collects all terms for which $\lim_{h \downarrow 0} \frac{o(h)}{h} = 0$.

Let \mathcal{E} be the power set of E . From the Kolmogorov extension theorem it follows that there exists a measurable space (Ω, \mathcal{F}) and, for each $x \in E$ and $m \in \{0, 1\}$, a unique probability measure

⁷ Note that recovered individuals do not become susceptible again (in finite time). Therefore, we can refer to the group of individuals that are no longer infected as “removed”, in line with the literature.

⁸ We refer to point 2., on page 11, for a discussion on this assumption.

⁹ The figure includes all transition rates, which we will introduce later.

P_x such that for all $0 \leq t_1 < \dots < t_k, k \in \mathbb{N}$, and all $F_1, F_2, \dots, F_k \in \mathcal{E}$, it holds that

$$P_x^m(F_1, \dots, F_k) = \sum_{y_1 \in F_1} \sum_{y_2 \in F_2} \dots \sum_{y_k \in F_k} p_{xy_1}^m(t_1) p_{y_1 y_2}^m(t_2 - t_1) \dots p_{y_{k-1} y_k}^m(t_k - t_{k-1}).$$

With each mode $m \in \{0, 1\}$, state $x \in E$, and probability measure P_x^m , we associate the expectation operator E_x^m . Throughout we will use the natural filtration $\mathbf{F} := (\mathcal{F}_t^X)_{t \geq 0}$. The set of \mathbf{F} -stopping times is denoted by \mathcal{M} .

Switching from mode m to mode $m', m, m' \in \{0, 1\}, m \neq m'$, is associated with an immediate and sunk cost $K_{mm'} \geq 0$, which represent all costs that are incurred once and are associated with entering or exiting lockdown. Examples of costs of entering include a change in bureaucracy, hiring people to work on legislation, local governments' work on policies, shops to make investments to accommodate new rules, home workers needing office chairs at home, etc. Exiting lockdown can be associated with costs around, e.g., hiring new staff, legislative changes, etc.¹⁰

For each mode $m = 0, 1$, we define a bounded function $c_m : E \rightarrow \mathbb{R}_+$, which represents the running health and non-health costs related to the number of infected members and, for mode $m = 1$, additional economic and indirect health costs (e.g., due to missed diagnoses of other diseases) as a result of lockdown. Letting $\chi_{\{v\}}$ denote the characteristic function that is equal to 1 if v is true and 0 otherwise, we take

$$c_m(i, r) = \delta_E(i, r)i + \delta_L \chi_{\{m=1\}}, \tag{2}$$

for some epidemic related per-patient cost function $\delta_E : E \rightarrow \mathbb{R}_+$, with $\delta_E(0, r) = 0$, and fixed lockdown cost flow $\delta_L > 0$. Although in our analyses the cost functions are chosen to be linear, our methodology can be applied to, e.g., convex specifications of the cost functions.

The social planner is assumed to discount costs at a constant rate $\rho > 0$. The decision problem is modeled as a nested optimal stopping problem:

$$V(i, r) = \inf_{\tau_1 \in \mathcal{M}} E_{(i,r)}^{m=0} \left\{ \int_0^{\tau_1} e^{-\rho t} c_0(I_t, R_t) dt + e^{-\rho \tau_1} K_{01} \right. \\ \left. + \inf_{\tau_2 \in \mathcal{M} | \tau_2 \geq \tau_1} E_{(I_{\tau_1}, R_{\tau_1})}^{m=1} \left[\int_{\tau_1}^{\tau_2} e^{-\rho t} c_1(I_t, R_t) dt + e^{-\rho \tau_2} K_{10} \right] \right. \\ \left. + E_{(I_{\tau_2}, R_{\tau_2})}^{m=0} \left(\int_{\tau_2}^{\infty} e^{-\rho t} c_0(I_t, R_t) dt \right) \right\}. \tag{3}$$

The first, second, and third terms represent the present values of the costs incurred before lockdown ($0 \leq t < \tau_1$), during lockdown ($\tau_1 \leq t < \tau_2$), and after lockdown ($t \geq \tau_2$), respectively.

Finally, we denote the *basic reproduction number* by \mathcal{R}_0^m , for modes $m = 0, 1$. The basic reproduction number represents the average number of infectives emerging from the introduction of a single infectious member into a completely susceptible population. In our model, this is given by

$$\mathcal{R}_0^m = \frac{\beta_m}{\gamma_m}.$$

Discussion of modeling assumptions

After having introduced our basic model, we now briefly discuss some of our assumptions before analysing the model.

1. The basic reproduction number, \mathcal{R}_0^m , is observed without error. This assumption is reasonable when prevalence of the disease can be accurately observed even at low levels (Mutesa, Ndishimye, Butera et al., 2021).

¹⁰ It should be noted though that these costs need not necessarily be positive. Also see our robustness analysis in Appendix C.

2. The system is closed, i.e., we assume no births and no non-COVID deaths. Births can be interpreted in two ways: either they could represent migration into the population or they could represent newborns. For the former, the impact of migration can be studied separately, as we do in Section 4.5. The latter is not expected to drive any of our results and is, therefore, ignored.
3. We restrict the social planner's policy choice to two or three modes. While, theoretically, a social planner could consider an infinite set of modes, many governments have opted for a relatively simple traffic-light system in their responses to the COVID-19 pandemic. Hence, our model is reasonably close to observed practice. In fact, we show in Section 4.4 that when measures are added or taken away from the intermediate mode only the speed with which measures are phased out is impacted, but not the sequencing of the optimal policy.
4. The social planner can only impose new restrictions once mode $m = 0$ is reached. While this assumptions looks restrictive, we assume throughout that switching from the intermediate mode to mode $m = 0$ is costless. Therefore, in principle, mode $m = 0$ can be skipped and stricter measures be introduced at any time.
5. The cost of infection is constant over the population. Throughout we use the weighted average of the cost per patient when in need of (hospital) treatment and the cost when no treatment is necessary. In the latter case, costs could be related to, e.g., self-isolation. Although this assumption might not be as realistic when I_t is small, since the population in our model aims to represent a country's population (in our case the UK), small values of I_t may still represent a large number of individuals.¹¹
6. The cost of being in lockdown is constant over time. This may not always be a realistic assumption, especially when measures are in place for a long time. Therefore, we relax this assumption in Appendix C.3 and show that our results do not change qualitatively.
7. Recovered members of the population are assumed to acquire immunity and that infection, recovery, and death rates are known to the social planner. These are standard assumptions in SIR(D) models, which we further discuss in Section 5 where we point to avenues for future research to weaken these assumptions. In our model, the recovery rate incorporates the weighted sum of the average recovery times of patients with and without symptoms, including those receiving hospital treatment (see Section 3.2). The incubation time can be incorporated in a similar way.

2.1. Properties of the stochastic process

We first look at how X behaves as a function of time in our benchmark model. Therefore let us first study the change in $X_t = (I_t, R_t)$. One can easily see that this can be written as

$$\frac{1}{dt} E_x^m \left[(dI_t, dR_t) \mid X_t \right] = \frac{1}{dt} E_x^m \left[dX_t \mid X_t \right] \\ = \frac{1}{dt} \sum_{y \in E \setminus \{x\}} q_{xy}^m (y - x) dt \\ = (\lambda_m(I_t, R_t) - \mu_m(I_t, R_t), \mu_m(I_t, R_t)),$$

for all $x \in \text{int}(E), m = 0, 1$, where the last step follows from (1). This gives the intuitive result that the number of infectives is expected to go up if, and only if, $\lambda_m(I_t, R_t) > \mu_m(I_t, R_t)$ and the num-

¹¹ An alternative way of modeling would be to consider δ_E as two branches where a fraction of I_t receives hospital treatment. Since we already take into account that different members require different treatment (if at all), we do not expect that this more complicated version would lead to qualitatively different results.

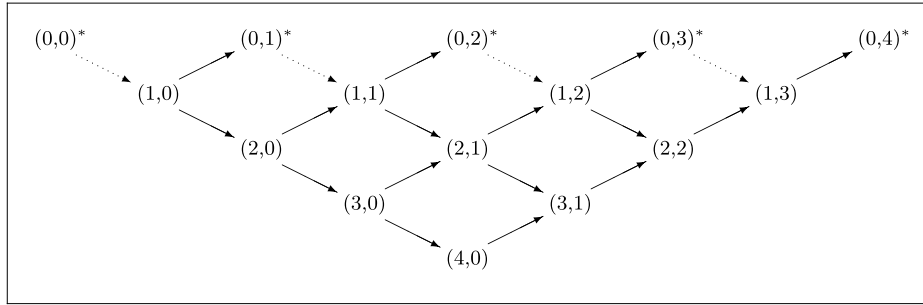


Fig. 2. Example of a lattice for $N = 4$, describing all potential paths. All absorbing states in the benchmark model are labeled with an asterisk. For our full model we add the additional transitions with the dotted arrows when there is an exogenous source of infections.

ber of recovered members is always expected to increase. If control measures in mode 1 lead to a decrease in the transmission rate, i.e., if $\lambda_2(\cdot) < \lambda_1(\cdot)$, then there exists a set of states in E for which the number of infections is expected to go down in mode 1, whereas it is expected to go up for the same states in mode 0. For our SIR-like model, the expected change in X_t equals

$$\frac{1}{dt} E_x^m [dI_t | X_t] = \beta_m I_t \frac{S(I_t, R_t)}{N} - \gamma_m I_t, \quad \text{and}$$

$$\frac{1}{dt} E_x^m [dR_t | X_t] = \gamma_m I_t,$$

so that in expectation we obtain the dynamics of the traditional deterministic SIR model.

One of the convenient features of the SIR model is that the state space, E , can equivalently be represented by a simple lattice structure; that is, states can be ordered to form a very simple tree describing all potential paths. Fig. 2 illustrates this lattice structure for the case $N = 4$. These lattice structures are useful in determining the value functions in all modes and determining the stopping sets of our optimal stopping problems.

The volatility of the process is given by

$$\begin{aligned} \frac{1}{dt} E_x^m [(dI_t^2, dR_t^2)] &= \frac{1}{dt} \sum_{y \in E} p_{xy}^m(dt) (y - x) \diamond (y - x) \\ &= \sum_{y \in E \setminus \{x\}} q_{xy}^m (y - x) \diamond (y - x) \\ &= \lambda_m(I_t, R_t) ((1)^2, 0^2) + \mu_m(I_t, R_t) ((-1)^2, (1)^2) \\ &= (\lambda_m(I_t, R_t) + \mu_m(I_t, R_t), \mu_m(I_t, R_t)), \end{aligned}$$

where \diamond denotes the component-wise multiplication of vectors and where we use that $(E_x^m dI_t)^2 = (E_x^m dR_t)^2 = 0$ for all $m = 0, 1$.

Note that when the function $\lambda_m(\cdot)$ and/or $\mu_m(\cdot)$ changes, there is a direct effect on the speeds with which the process moves through the tree. When the volatility increases, costs are incurred over a shorter period of time (in expectation). Therefore, an increase in, e.g., β has two effects. Individuals are infected at a higher rate, but at the same time, the volatility of the process is also increased. The basic reproduction number is given by $\mathcal{R}_0^m = \frac{\beta_m}{\gamma_m}$, for modes $m = 0, 1$. Therefore one can alternatively write

$$\frac{1}{dt} E_x^m [(dI_t^2, dR_t^2)] = (\beta_m I_t \frac{S(I_t, R_t)}{N} + \gamma_m I_t, \gamma_m I_t) = \gamma_m I_t (\mathcal{R}_0^m \frac{S(I_t, R_t)}{N} + 1, 1).$$

as more commonly done for SIR models.

3. Model solution

We solve the optimal stopping problem (3) by working backwards. So, we start with the situation where the social planner has come out of lockdown and there are, thus, no options left. We illustrate our findings with our benchmark model. In Section 4 the full model is analyzed. Note that then, the social planner has a new option to go into lockdown.

3.1. Value after lockdown

For our benchmark model, after lockdown has been entered and exited, the social planner has no more decisions to take, i.e., we remain in mode $m = 0$ perpetually. Recall that $X_t = (I_t, R_t)$ denotes the state of the process at any time t . For a given initial state $x = (i, r) \in E$, we denote the total expected discounted costs after lockdown (under P_x^0) by $G(x)$. That is,

$$G(x) = E_x^{m=0} \left[\int_0^\infty c_0(X_t) e^{-\rho t} dt \right]. \tag{4}$$

The following proposition shows that $G(x)$ can be defined recursively, that is, it can be written as a function of all states the chain can transition to when leaving x .

Proposition 1. *Let $x \in E$. Then, $G(x)$ is recursively defined as*

$$G(x) = \frac{c_0(x)}{\rho - q_{xx}^{m=0}} + \sum_{y \in E \setminus \{x\}} \frac{q_{xy}^{m=0}}{\rho - q_{xx}^{m=0}} G(y). \tag{5}$$

Moreover, for $x \in \Lambda$, it holds that

$$G(x) = \frac{c_0(x)}{\rho}. \tag{6}$$

Proof. See Appendix A.

The expression in (5) consists of two terms. The first term relates to the costs incurred while being in state x . After a transition to some state $y \neq x$, the costs after the transition are represented by $G(y)$. The present value of these costs after the event of a transition, for all possible subsequent states, is then given by the second term multiplied by q_{xy}^m from the Markov generator which relates to the transition probability. Thus, the second term is the expected present value of costs after the next transition.

Proposition 1 provides the solution to the general problem in (4). Note that since in our formulation, see e.g. (2), $c_0(x) = 0$ for all $x \in \Lambda$, one could simplify (6) to $G(x) = 0$.

We observe that it is not guaranteed that the recursive formulation has a solution. However, due to the nature of the lattice that corresponds to our problem, a sufficient condition would be that if the process starts in a state that is not in Λ , i.e., the process is not absorbed, then the process will reach a state in Λ a.s. at some future point in time.

Since the chain is absorbing, it is easy to see that $G(x)$ can be determined for all x by working recursively through the tree. First, for $x \in \Lambda$, $G(x) = 0$. For $x \in E \setminus \Lambda$, G in (5) can be determined using a lattice method, finding $G(x)$ for all $x = (i, r) \in E$ such that $i + r = N, N - 1, N - 2, \dots, 1$. Then, for any x in the interior of the state space, i.e. for any $x \in \text{int}(E)$, one can check that $G(x)$ is equal to

$$G(x) = \frac{c_0(x)}{\rho + \mu_0(x) + \lambda_0(x)} + \frac{\mu_0(x)}{\rho + \mu_0(x) + \lambda_0(x)} G(x + (-1, 1)) + \frac{\lambda_0(x)}{\rho + \mu_0(x) + \lambda_0(x)} G(x + (1, 0)).$$

3.2. Value in lockdown

In mode 1, the social planner has lockdown in place. This ends when the social planner decides to switch back to mode $m = 0$ at sunk cost K_{10} . After this switch, the value function is given by G , cf., (4).

Therefore, in mode 1, the social planner is faced with the optimal stopping problem

$$F(x) = \inf_{\tau \in \mathcal{M}} \mathbb{E}_x^{m=1} \left[e^{-\rho\tau} (G(X_\tau) + K_{10}) + \int_0^\tau c_1(X_t) e^{-\rho t} dt \right], \tag{7}$$

where G is defined in (4). The integral captures the present value of costs while being in mode $m = 1$, i.e., until lockdown is exited at time τ . The present value of costs incurred after τ are given by the first term.

Proposition 2. *There exists a unique function $\varphi : E \rightarrow \mathbb{R}_+$ such that*

$$\varphi(x) = \min \left\{ G(x) + K_{10} \frac{c_1(x)}{\rho - q_{xx}^{m=1}} + \sum_{y \neq x} \frac{q_{xy}^{m=1}}{\rho - q_{xx}^{m=1}} \varphi(y) \right\}. \tag{8}$$

The function φ solves the optimal stopping problem in (7), i.e., $F = \varphi$. The optimal stopping time is the first exit time of the set C_1 , where

$$C_1 = \{x \in E : \varphi(x) < G(x) + K_{10}\}.$$

Proof. See Appendix A.

Proposition 2 shows that in order to find the value, we need to distinguish states where it is decided to switch to mode 0 and where it is decided to delay exiting lockdown. Thus, problem (7) has a solution that allows us to split the state space E into a continuation set $C_1 \in \mathcal{E}$ and a stopping set $D_1 = E \setminus C_1 \in \mathcal{E}$ such that the social planner decides to stay in mode 1 for all $x \in C_1$ and to switch to mode 0 for $x \in D_1$.

The first term on the right-hand side in (8) represents the value if the social planner decides to switch to mode $m = 0$ immediately, i.e., when x is in the stopping set. The second term is the value when x is in the continuation set so that the decision to come out of lockdown is “delayed”: it is the sum of the costs while being in state x and the expected present value of costs after transition.

The stopping set, for mode $m = 1$, is then given by

$$D_1 = E \setminus C_1 = \{x \in E : \varphi(x) = G(x) + K_{10}\}.$$

Since for all $x \in \Lambda$ it holds that $q_{xy}^m = 0$ for all $y \in E$, F in (8) can directly be determined. For $x \in E \setminus \Lambda$, $F(x)$ is recursively determined as a function of the states the chain can transition to when leaving x . Equivalent to what we found for G , F can be found using an induction method where (8) is solved for all $x = (i, r) \in E$, such that $i + r = N, N - 1, N - 2, \dots, 0$.

For further reference, the first exit time of a set $C \in \mathcal{E}$ is defined as the stopping time

$$\tau_C := \inf\{t \geq 0 \mid X_t \notin C\}.$$

Notice that, given the set $C_1 \in \mathcal{E}$, it follows that F can be written as

$$F(x) = \begin{cases} \frac{c_1(x)}{\rho - q_{xx}^{m=1}} + \sum_{y \in E \setminus \{x\}} \frac{q_{xy}^{m=1}}{\rho - q_{xx}^{m=1}} F(y) & \text{if } x \in C_1, \\ G(x) + K_{10} & \text{if } x \in E \setminus C_1. \end{cases}$$

Illustration

Table 1

Model parameters and their baseline values. Our choice of N can be interpreted as the size of the grid that we impose on the population. For example, if $N = 500$, then each change in, say, the number of infected from i to $i + 1$ represents an increase in the infected population equal to 0.2% of the total population.

Param.	Description	Value
β_0	Infection rate in mode 0 [days ⁻¹]	0.3
β_1	Infection rate in mode 1 [days ⁻¹]	0.15
γ_0, γ_1	Recovery rate [days ⁻¹]	0.1
N	Size of host population	500
ρ	Discount rate	0.1/365
K_{01}	Sunk costs of entering lockdown	£2 mln
K_{10}	Sunk costs of exiting lockdown	0
δ_E	cost per patient	£4000 per day
δ_L	per-period cost of lockdown	£18,840 per day

Table 1 summarizes our baseline parameterization. The resulting corresponding basic reproduction numbers are $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$ for modes 0 and 1, respectively. Starting with mode 0, the empirical literature is not unanimous in its estimates for the basic reproduction rate \mathcal{R}_0^0 . However, as pointed out by, e.g., Liu, Gayle, Wilder-Smith, & Rocklöv (2020), the median and average estimates are roughly 3 for empirical studies on mostly China. Toda (2020) estimates \mathcal{R}_0^0 to be equal to 2.7 for the UK and 3.7 for the USA. An estimate close to 3 is also in line with empirical studies by, e.g., Roques, Klein, Papaix, Sar, & Soubeyrand (2020); Salje et al. (2020); Zhao et al. (2020), and You et al. (2020). The value of γ_m is estimated to be around $\frac{1}{10}$ by, e.g., Zhou et al. (2020), i.e. it takes on average $\frac{1}{\gamma_m} = 10$ days for people to recover. Setting $\gamma_m = \frac{1}{10}$ is also done by, e.g., Patterson-Lomba (2020); Roques et al. (2020), and Toda (2020). He et al. (2020) find 9.5 days: 2.5 for being infected without symptoms and 7 days with symptoms. Moreover, it was estimated by Salje et al. (2020) that it takes 2.3 days for patients to recover with mild symptoms and, estimated by Di Domenico, Pullano, Sabbatini, Boëlle, & Colizza (2020), it takes 17 days for people in ICU.

For mode 1, we follow, e.g., Garriga et al. (2020); Patterson-Lomba (2020), and Kruse & Strack (2020), in assuming that \mathcal{R}_0^1 is roughly half of \mathcal{R}_0^0 (for some empirical studies see, e.g., You et al., 2020), while assuming that the rate of recovery is not affected by the lockdown.

We assume that all costs are measured in thousands GBP. For the epidemic related cost parameter δ_E we use the study by Bethune & Korinek (2020) and assume that the virus imposes a cost of \$50k for each infective, so that, using the exchange rate $\$1 = \text{£}0.8$, $\delta_E = (0.8)(50)\mu_m = 4$. The lockdown related costs can be estimated to be $\text{£}2.4b$ per day in Britain,¹² which then translates to $\delta_L = \frac{2.4b}{1000} \times \frac{N}{63.7m} = \frac{1200}{63.7} \approx 18.84$ for our population of size N , assuming the British population to be 63.7m.

We solve (8) employing a lattice method as described above. When starting with states in Λ , it can be easily determined whether staying in lockdown perpetually is optimal or whether coming out of lockdown is optimal. When assuming D_1 and C_1 to be empty at the outset, states in Λ can either be allocated to D_1 or C_1 . Next, states that transition to states in Λ can be studied: the value as given by (8) can be determined and these states can be assigned to the appropriate set. Recursively, one can continue this process until all states have been treated.

Panel (a) of Fig. 3 illustrates the stopping set for this parameterization. Throughout this paper, we use orange to represent states where it is optimal to come out of (full) lockdown. Generally, the stopping set consists of two types of states: states that are in Λ , i.e. all absorbing states, or are close to Λ , and states (i, r) such that

¹² The Times, April 6, 2020, <https://www.thetimes.co.uk/edition/business/consumer-confidence-at-its-lowest-since-the-financial-crisis-p6phf7x3k>.

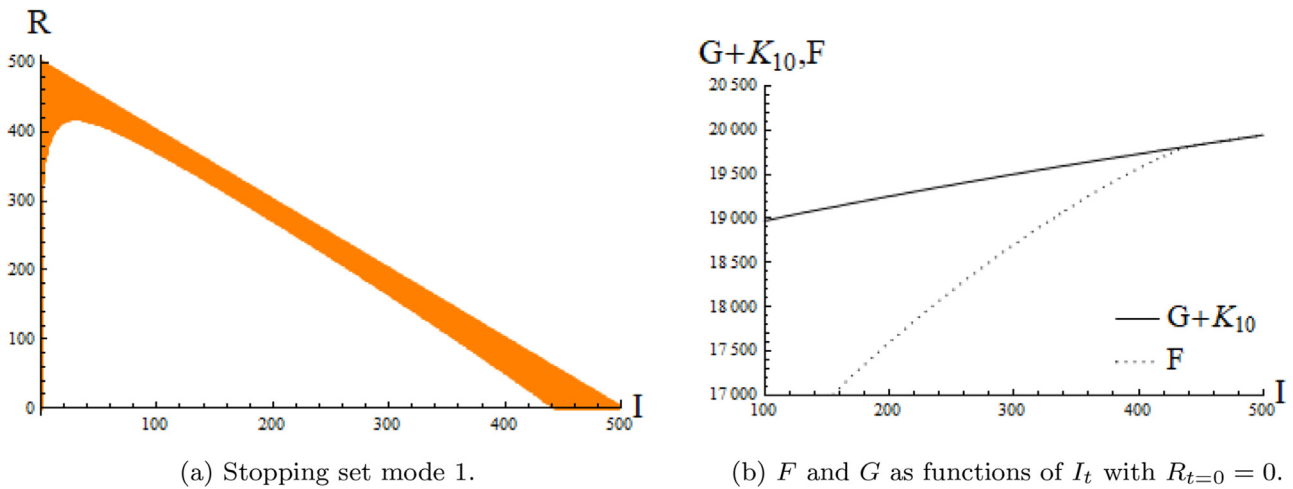


Fig. 3. Stopping set (Panel (a)) and value function (Panel (b)) for the optimal stopping problem for switching to mode 0. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, and $\delta_L = \frac{1200}{63.7}$.

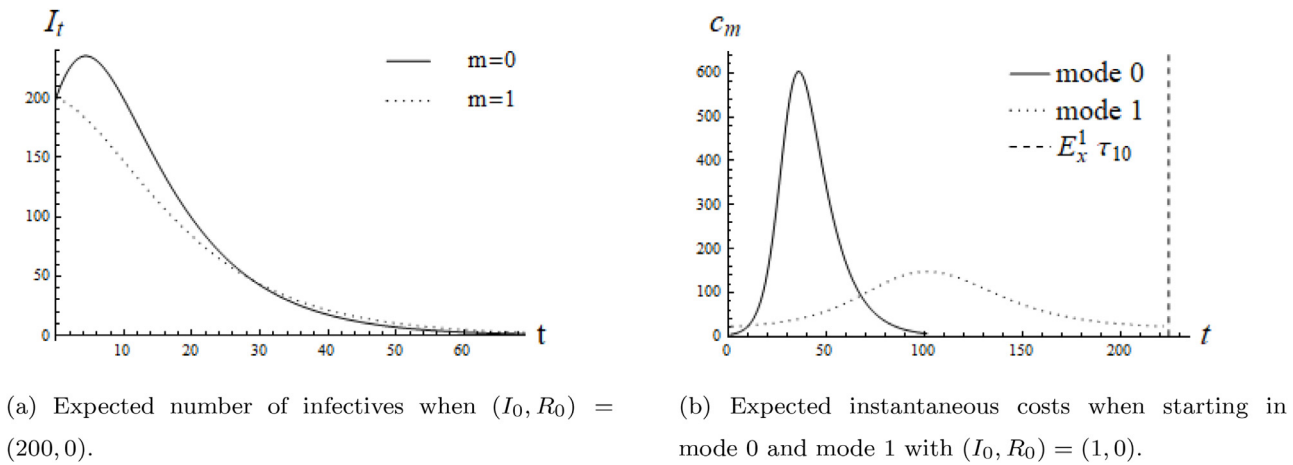



Fig. 4. Expected development of I_t for both modes (Panel (a)) when $(I_0, R_0) = (200, 0)$ and expected instantaneous costs c_m (Panel (b)) for the optimal stopping problem for switching to mode 0 as functions of time. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, and $\delta_L = \frac{1200}{63.7}$.

$i + r$ is close to N . In the former case, the process is close to absorption or is absorbed, which means that lockdown-related costs outweigh any benefits of suppressing the epidemic. Therefore, it is optimal to lift lockdown. In the latter case, the number of susceptibles is relatively small. Since staying in mode 1 only keeps the infection rate low, staying in mode 1 is costly and, therefore, it is optimal to switch and lift lockdown.


Panel (b) of Fig. 3 shows a cross-section of the value function F , as well as the cost function G , for the case where $R_t = 0$ for $t = 0$. Other cross-sections of the functions F and G , qualitatively, look the same. Panel (b) illustrates that the value function $F(X)$ in the continuation set exhibits lower total costs than $G(X) + K_{10}$, i.e. if it were to stop. To understand this, consider the total expected accumulated discounted costs if the social planner decides ex-ante to always stay in mode 1. In that case, the lockdown related costs are incurred even when the number of infectives is low, but exhibits lower expected costs per unit of time when the number of infectives is high relative to being in mode 0. The curve F takes advantage of lower expected cost per unit of time when the number of infectives is high, with the prospect of not incurring lockdown related cost when the number of infectives is low. Comparing $G(X) + K_{10}$ to $F(X)$ illustrates that the option to switch has led to a lower present value of expected total costs in the continuation set. In the stopping set, $F(X) = G(X) + K_{10}$.

Fig. 4 illustrates the development of the disease as a function of time. We illustrate two cases, one where there are already a high number of infectives and one at the start of the epidemic. Panel (a) relates to the situation depicted in Fig. 3(b), and shows the expected number of infectives when starting in state $(i, r) = (200, 0)$ for modes 0 and 1, respectively. The total number of infectives and the resulting (health and economic) related costs are lower in mode 1 but, in addition, lockdown costs are incurred. Irrespectively, there is a steep incline of infectives in mode 0 at $t = 0$ (a “second wave”), which results in it being optimal for the social planner to delay the moment of exiting lockdown. Therefore the optimal total costs are lower when remaining in mode 1 for a (short) period of time and then switching to mode 0.

Panel (b) depicts the incurred cost per unit of time after lockdown (black) and when starting while in lockdown (dotted), when considering the expected path of (I_t, R_t) , i.e. $c_m(E_{X_0}^m(X_t))$, $m = 0, 1$, with $X_0 = (I_0, R_0) = (1, 0)$. Note that the dashed line represents the expected point in time when switching to mode 0 is optimal. For the dotted curve, i.e. when starting in mode 1, notice that at $t = 0$ the curve exceeds the solid curve representing costs incurred when starting in mode 0. Even though in both cases $(I_t, R_t) = (1, 0)$, in mode 1 additional lockdown-related costs are incurred which are independent of the number of infectives. Clearly, the costs peak earlier for mode 0, reflected by the solid curve, and at the same



(a) Stopping set for switching to mode 1 (blue) and to mode 0 (orange). Gray states are states in the continuation set for which the blue area may be reached.



(b) V , F , and G as functions of I_t with $R_{t=0} = 0$.

Fig. 5. Stopping set (Panel (a)) and value functions (Panel (b)) for the optimal stopping problem for switching to mode 1. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1, gray: states from which reaching the blue area is possible. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, and $\delta_L = \frac{1200}{63.7}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

time the expected costs for all $t < 68$ are higher than the costs incurred when in mode 1. This also illustrates that lockdown is lifted only when the number of infectives is well past its peak to prevent a (severe) second wave.

3.3. Value before lockdown

In mode 0, before lockdown has started, the social planner is faced with the optimal stopping problem

$$V(x) = \inf_{\tau \in \mathcal{M}} E_x^{m=0} \left[e^{-\rho\tau} (F(X_\tau) + K_{01}) + \int_0^\tau c_0(X_t) e^{-\rho t} dt \right], \quad (9)$$

where F is defined in (7). The integral captures the present value of costs while being in mode $m = 0$, i.e., until lockdown is imposed at time τ . The present value of costs incurred after τ are given by the first term.

The following proposition shows that also for the stopping problem in (9) the solution splits the state space into a stopping set and a continuation set.

Proposition 3. *There exists a unique function $\psi : E \rightarrow \mathbb{R}_+$ such that*

$$\psi(x) = \min \left\{ F(x) + K_{01} \frac{c_0(x)}{\rho - q_{xx}^{m=0}} + \sum_{y \neq x} \frac{q_{xy}^{m=0}}{\rho - q_{xx}^{m=0}} \psi(y) \right\}. \quad (10)$$

The function ψ solves the optimal stopping problem (9), i.e., $V = \psi$. The optimal stopping time is the first exit time of the set C_0 , where

$$C_0 = \{x \in E : \psi(x) < F(x) + K_{01}\}.$$

Proof. See Appendix A.

Similar to Proposition 2, Eq. (10) contains two terms: the first term giving the value when going into lockdown immediately and the second term representing the present value of costs when x is in the continuation set so that lockdown is delayed.

The stopping set, for mode $m = 0$, is then given by

$$D_0 = E \setminus C_0 = \{x \in E : \psi(x) = F(x) + K_{01}\}.$$

It is easy to see that (10) can be determined using an induction method as described for G and F .

Illustration We illustrate Proposition 3 for the parameterization of Section 3.2. Fig. 5 illustrates the stopping sets for both stopping


problems: the blue and orange areas represent all states for which it is optimal to switch from mode 0 to mode 1 and vice versa, respectively. The gray area is the set of all states for which there is a positive probability that the blue area can be reached. So, when the process is in a state in the white area, then it will never be optimal to enter lockdown. Going back to Fig. 4(b), the difference between the two curves illustrates that instantaneous costs, when not in lockdown, increase more rapidly for small values of t . As such, an expected positive gain from switching to mode 1 vanishes if the social planner waits too long, demonstrating that switching in early stages of the process (when I_t and R_t are both “small”) is optimal and that switching is not optimal in later stages of the process.

The value functions V and F as well as the total costs function G are depicted in Fig. 5(b). The dash-dotted curve represents the value in mode 0, V , at time $t = 0$. Note that the solid curve representing the graph of the function G , i.e., the value in mode 0 after the second switch, also represents the graph of the value function V if it is optimal to stay in mode 0 from time $t = 0$ onward, i.e., when lockdown is never implemented.¹³


We can now observe the following for the two options. The continuous pasting principle is visible for the option to switch to mode 0 from mode 1: the dotted curve and the solid curve have approximately the same ascend around the boundary of the stopping set. For the first switching option, instantaneous switching is optimal when I_0 is small, as Panel (a) illustrated. Therefore, $V(I_0, 0)$ coincides with $F(I_0, 0) + K_{01}$ for sufficiently small I_0 , the latter represented by the upper dotted curve in Panel (b). Then, instantaneous switching is optimal as long as the value to permanently stay in mode 0 exceeds $F + K_{01}$. Indeed, V coincides with G for high values of I_0 where switching to mode 1 is not optimal.

Nevertheless, for I_0 (very) close to zero, there is a positive value of waiting: it is optimal for the social planner to delay switching until D_0 is reached. For these values there is a positive probability that at the next transition, or in a small number of transitions, the process hits an absorbing state and in that case, the disease will not be able to spread further so that lockdown is redundant. Since

¹³ Note that, in contrast to Fig. 3, here, the solid line represents the costs of never going into lockdown. Thus, when comparing to Fig. 3, K_{10} is not incurred for the scenario in Fig. 5.



(a) $F + K_{01}$ as a function of I_t with $R_{t=0} = 265$ and the expected total cost if the social planner decided to stay in mode 0 permanently (dashed).



(b) F , V , and G as functions of I_t with $R_{t=0} = 265$.

Fig. 6. Value functions and total cost functions with $R_0 = 265$ for the optimal stopping problem for switching to mode 1 as functions of time. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, and $\delta_L = \frac{1200}{63.7}$.

lockdown is associated with a fixed switching cost, the expected total cost involved with entering lockdown is not outweighed by the expected gain. This is illustrated by Fig. 6. Panel (a) shows for the same parameterization as before that, for $I_0 < 3$ and $R_0 = 265$, waiting is optimal. The graphs of the resulting functions V , F , and G are shown in Panel (b).

4. Optimal policy

In this section we first study the social planner's optimal stopping strategy in our SIR-like base-case model further. Then, we expand our benchmark model in the following ways.

1. We first extend the number of classes individuals can be in. In addition to agents being infected (I), recovered (R), and susceptible (S), we allow for events where individuals are deceased (D). This model is the equivalent of a standard SIRD model.
2. Next, we relax the assumption that lockdown can only be imposed once.
3. Additional modes are introduced, allowing the social planner to phase lockdown in and/or out.
4. The policy regarding open borders is studied. Infections can occur even when $I_t = 0$.
5. Restrictions imposed by health care capacity are taken into account.

Finally, we execute an extension where we study the stopping sets in a scenario where the lockdown related costs are increasing over time. The analysis is concluded by a short summary of the comparative statics; the full analysis can be found in Appendix C.

4.1. SIR

Recall that the transition rates are given by

$$\begin{aligned} \lambda_m(I_t, R_t) &= \beta_m I_t \frac{S(I_t, R_t)}{N} \\ &= \mathcal{R}_0^m \gamma_m I_t \frac{N - I_t - R_t}{N}, \end{aligned}$$

$$\mu_m(I_t, R_t) = \gamma_m I_t,$$

for modes $m = 0, 1$, where $\mathcal{R}_0^m = \frac{\beta_m}{\gamma_m}$. In this section, we parameterize the dynamics by \mathcal{R}_0^m and μ_m , for modes $m = 0, 1$. Recall that

the process has the following properties

$$\frac{1}{dt} \mathbb{E}_x^m \left[(dI_t, dR_t) \mid X_t \right] = \gamma_m I_t \left(\mathcal{R}_0^m \frac{S(I_t, R_t)}{N} - 1, 1 \right), \quad (11)$$

$$\frac{1}{dt} \mathbb{E}_x^m \left[(dI_t^2, dR_t^2) \mid X_t \right] = \gamma_m I_t \left(\mathcal{R}_0^m \frac{S(I_t, R_t)}{N} + 1, 1 \right). \quad (12)$$

For the illustrations we use the same parameterization as before.

Fig. 7 depicts the expected evolution of the number of infected, susceptible, and removed individuals as a function of time t . For panels (a) and (b), the brown region represents the total number of active cases of infectives at each moment in time. The blue area depicts the number of susceptible individuals and the red area, then, are all remaining individuals that are labeled as removed/recovered. At time $t = 0$, 490 individuals are susceptible and at time $t = 80$, in mode 0, a very large share of the host population has been infected and subsequently recovered. This is different for mode 1 where a significant share of the host population has not been infected. This would apply if at time $t = 0$, the social planner had decided to immediately switch to mode 1. Notice that the curve of infectives in mode 1 is flatter than the curve in mode 0. Panel (c) juxtaposes the expected number of infectives at each point in time for modes 0 and 1.

Fig. 8 illustrates the stopping sets. In both panels the blue area is the set of all states where it is optimal to switch from mode 0 to mode 1. The orange area represents all states where it is optimal to switch back to mode 0 from mode 1. The panels also feature stream plots for the system

$$\dot{I} = \mathcal{R}_0^m \gamma_m I (N - I - R) - \gamma_m I,$$

$$\dot{R} = \gamma_m I.$$

This system represents the expected evolution of I_t and R_t for modes $m = 0, 1$ for any point in the grid that represents the state at time $t = 0$. The black affine curve represents all points (i, r) such that the expected change in the number of infectives is 0 in the next instance of time, i.e., $\dot{I} = 0$. Panels (a) and (b) show the stream plot for modes 0 and 1, respectively. The plot shows that, in expectation, the maximum number of infectives at any time t is lower for mode 1, i.e. $\sup_t \mathbb{E}_x^{m=1} I_t \leq \sup_t \mathbb{E}_x^{m=0} I_t$. From the blue area in the figure, two implications for the optimal decision to impose lockdown jump out. First, in mode 0, it is optimal to undertake action as soon as the first infectives have been confirmed. The stopping




Fig. 7. Expected number of infected (brown), susceptible (blue), and recovered/removed (red) individuals for mode 0 (Panel (a)) and mode 1 (Panel (b)). Panel (c) shows the number of infectives in expectation together. Here, $(I_0, R_0) = (10, 0)$, $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, and $N = 500$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

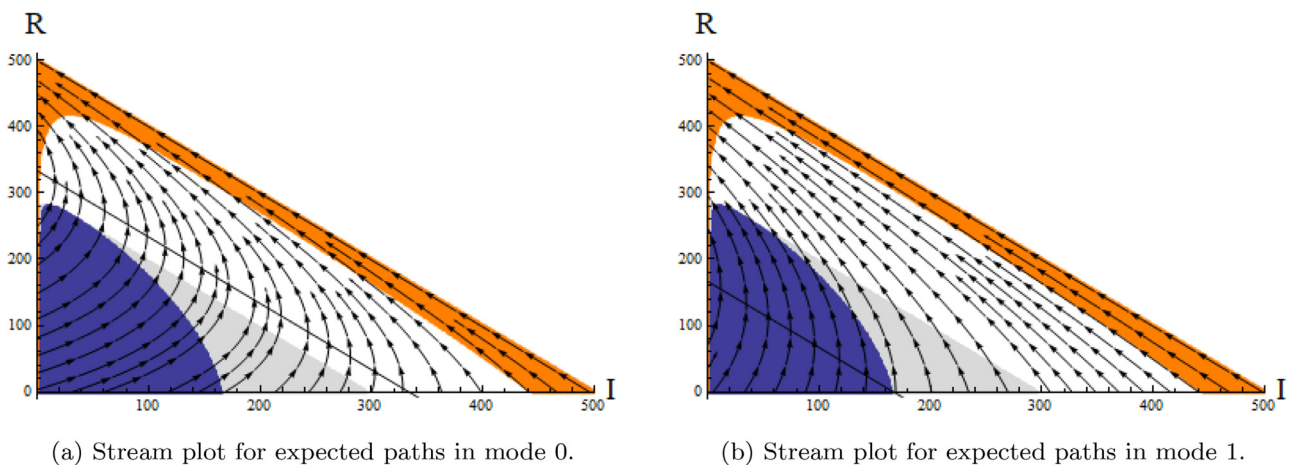


Fig. 8. Stopping sets for switching between mode 0 and 1 (blue) and for switching between mode 1 back to mode 0 (orange). Figures include a vector/stream plot which indicate the direction of I_t and R_t in expectation. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1, gray: states from which reaching the blue area is possible. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, and $\delta_L = \frac{1200}{63.7}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

set in mode 0, the blue area, covers the combinations of infectives and recovered individuals close to the origin of the plot, which indicates that switching is optimal during the early stages of a potential epidemic. Second, if the social planner waits too long or when the disease is discovered late, the process may start outside the blue and gray areas, implying that it will a.s. never be optimal to switch to mode 1. This could happen, for instance, due to a lack of testing in the early stages. If the epidemic gets discovered when the state is in the gray area, then it might be optimal to impose a lockdown in the future, but not immediately. So, the social planner should act swiftly if the disease is detected early on. If the process leaves the blue area, before switching has taken place, the expected gain from switching is not enough to compensate for the total expected lockdown costs and the sunk cost of entering lockdown.

4.2. SIRD

In our SIR-like model it is assumed that all infected individuals will recover. In order to study the impact of infection fatality rates on the optimal switching times, the model can be extended, as described in Section 2, to include an additional class of individuals: deceased (D). For our continuous-time Markov chain, $X_t = (I_t, R_t, D_t)$ represents the state of the system at each time $t \geq 0$ such that $X_t \in E$ with $E = \{(i, r, d) \in \mathbb{N}_N^3 \mid i + r + d \leq N\}$, where D_t is the number of deceased individuals. Denote by $\theta_m(I_t, R_t, D_t) =$

$\frac{\phi_m}{1-\phi_m} \gamma_m I_t$ the mortality rate, where ϕ_m denotes the death rate, i.e., the fraction of member that dies due to an infection (also see, e.g., Bastos & Cajueiro, 2020). Then, the Markov generator is defined as $Q^m = (q_{xy}^m)_{x,y \in E}$ with, for $m = 0, 1$,

$$q_{xy}^m = \begin{cases} \lambda_m(x_1, x_2, x_3) & \text{if } y - x = (1, 0, 0), \\ \mu_m(x_1, x_2, x_3) & \text{if } y - x = (-1, 1, 0), \\ \theta_m(x_1, x_2, x_3) & \text{if } y - x = (-1, 0, 1), \\ 0 & \text{otherwise.} \end{cases}$$

We will use an estimate of 2% for the death rate.¹⁴ The estimated average number of years lost per deceased individual is estimated to be 16 (Pifarré i Arolas et al., 2021) and for the quality-adjusted life-year (QALY), we assume that the cost per QALY is equal to £6.8k.¹⁵ Thus, per time unit the cost per individual in D is $\delta_D = (16)(6.8)\rho$ in each mode. Note that this model can be solved

¹⁴ We base ourselves on the data presented on <https://coronavirus.data.gov.uk/> where one can check that the death rate is highly volatile and varies between roughly 0.4% and 3.5%.

¹⁵ To estimate the cost per QALY, we refer to a report presented by the UK's Department of Health and Social Care (DHSC) and Office for National Statistics (ONS), published on 9 September 2021, "Direct and Indirect health impacts of COVID-19 in England". The cost per QALY differs per region and we have chosen to use Yorkshire's estimate, which is below London's but above estimates for some other regions (e.g., North East). We will argue that a deviation from this estimate does not qualitatively change our results.

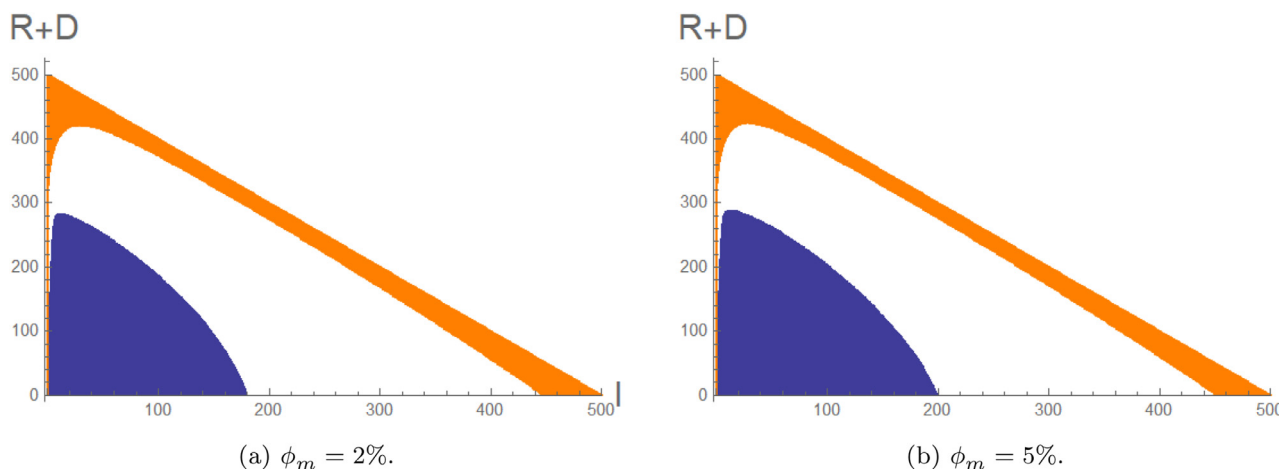


Fig. 9. Stopping sets for SIRD model. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_m = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, $\delta_I = \frac{1200}{63.7}$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

recursively as well, using the same lattice method as described in Section 3.2.

Results What we find is that if in our SIR-like model a state (i, r) is in a stopping set, then any state $(i, r', d') \in E$ such that $r' + d' = r$ is in the equivalent stopping set for the SIRD model. Intuitively, since the cost associated with any deceased individual is sunk, further actions taken by the social planner cannot change outcomes that happened in the past. Rather, the social planner will need to take into account the cost of future deaths. Note that this result does not depend on our parameterization. In the SIRD-like model the total costs are higher compared to the SIR-like model because, in addition to the lockdown related costs and the epidemic related per-patient costs, we add that, on average 2% of the transition out of I are associated with a positive shock of $(16)(6.8)k$ to the present value of total costs. Since the costs are higher, we find that the blue area has slightly increased and the orange area has slightly shrunk, as illustrated by Fig. 9a. Fig. 9b then shows the impact of the mortality rate. Setting the rate as high as 5% allows us to more clearly observe the impact. In line with what we saw for Panel (a), a higher mortality rate leads to an expansion of the blue set, especially for states where $R_t + D_t$ is small. In addition, we observe a slight change in the orange area in a way that the social planner delays coming out of lockdown, predominantly for states where $R_t + D_t$ is small.

An alternative way to model the cost associated with D is by incurring a one-off cost at each transition, i.e., when an individual moves from I to D , a cost of $(16)(6.8)k$ is incurred. This way of modeling would result in the same (qualitative) result.

For what follows, we continue by illustrating the stopping sets in a way as done in Fig. 9, i.e., by adding up all individuals in R and D .

In a similar way other models can also be studied using our set-up, by extending the number of classes of individuals. This includes for instance, the SEIR model (see, e.g., Atkeson, 2020; Wang et al., 2020) where susceptible individuals become exposed before infected.

4.3. A sequence of lockdowns

Since costs associated with lockdown can be high, it might be optimal for the social planner to impose a sequence of shorter lockdowns, rather than imposing one long one. This means the social planner has the option to alternate between modes. In terms

of our value function, for any lockdown $k = 1, \dots, K$,¹⁶ we can denote V_k , F_k , and G_k as the value functions equivalent to (9), (7), and (4), respectively. Then, G_k coincides with V_{k+1} and thus has embedded the option to go back into lockdown. This applies to any $k < K$, i.e., for $k = K$, G_k is the value as defined in (4). We assume that the switching costs associated with the first time lockdown is imposed is not smaller than the costs assumed for the benchmark model. These costs include, e.g., investments undertaken by companies to provide a safe environment such as the provision of sneeze guards or, e.g., the purchase of equipment and software to be able to work from home. However, for subsequent lockdowns these investment do not need to be undertaken again and as such the switching costs for later lockdowns are assumed to be a fraction of the K_{01} .

Results We start by assuming that the switching costs are equal 1% of K_{01} . Fig. 10 shows the stopping set for our parameterization. Panel (a) shows that the stopping set for the first lockdown is not qualitatively different from what we have found before. The result that when the epidemic is discovered too late (i.e., when the number of infected individuals is sufficiently high), intervention is no longer optimal is also present here. Since the first lockdown is concerned with substantial switching costs, the intuition as found for the SIR-like model applies here as well: the time period during which there are gains from interventions is too short to weigh against the costs. In addition, we find again that it is optimal to go into lockdown immediately when the first infections occur.

Panels (b) and (c) illustrate that, when switching costs are very low, subsequent lockdowns are always immediately imposed unless we have reached a state that lies in the stopping set for coming out of lockdown. In other words, lockdown ends when the number of infections or the number of susceptible individuals is low, as before, but is imposed again when infections are rising. Surprisingly, the orange area in Panel (a) is not substantially larger than the orange area in the SIR model. This means that, despite our initial intuition, it is not optimal to come out of lockdown temporarily: the increase in the infection rate and thereby the increase in infections is not outweighed by a recess of control measures. The only noticeable difference is for states where I_t is close to 0. Then, for $R_t + D_t > 290$ we see that the orange area is somewhat thicker which indicates that the option for subsequent lockdowns

¹⁶ For illustration purposes we assume that the social planner can only impose $K \geq 2$ lockdowns. However, this may not be an unrealistic assumption when people are more reluctant to obey restrictions as more lockdowns are imposed.




Fig. 10. Stopping sets when the social planner can alternate between modes K times with switching costs $0.01K_{01}$ for imposing subsequent lockdowns. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_m = 0.1$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, $\delta_L = \frac{1200}{63.7}$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 11. Stopping sets when the social planner can alternate between modes K times with switching costs $0.5K_{01}$ for imposing subsequent lockdowns. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_m = 0.1$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, $\delta_E = 4$, $\delta_L = \frac{1200}{63.7}$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

makes that the social planner comes out of lockdown sooner. Intuitively, these are the states where the increase in infection rate is least “harmful”. The same applies to subsequent lockdowns.

Fig. 11 illustrates the same figures but when switching costs for later lockdowns are 50% of K_{01} and confirms what we found before. Although the blue area has shrunk for lockdowns $k \geq 2$, the social planner does not change its strategy for the first lockdown.

In what follows, we analyze the stopping sets for the first lockdown as long as subsequent lockdowns are not qualitatively different from what we have discussed above.

4.4. Phased lockdown

So far, we have only considered two states: state $m = 0$ where no intervention takes place and state $m = 1$ where a lockdown is imposed. Our set-up allows for extensions where more modes can be considered. Thereto an intermediate mode is introduced. In this mode only mild measures are in place, such as social distancing. These come with a running cost that is lower than those of a full lockdown and higher than in mode $m = 0$. Costs for the intermediate mode can be associated with, e.g., shops and restaurants allowing a lower number of consumers and investments made to be able to adhere to these measures.

We redefine and relabel the modes in the following way. In mode $m = 0$ there is no intervention, in mode $m = 1$ mild measures are in place, and mode $m = 2$ represents a full lockdown. In other words, what was previously called mode 1 is now mode 2. Correspondingly, we can distinguish lockdown related costs δ_L^m for modes $m = 1$ and $m = 2$. Our parameterization for illustration purposes for modes $m = 0$ and $m = 2$ is the same as before, i.e. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^2 = 1.5$, and $\gamma_0 = \gamma_2 = 0.1$. For mode 1, we consider two cases: $\mathcal{R}_0^1 = 2$ and $\mathcal{R}_0^1 = 2.25$, with $\gamma_1 = 0.1$. For mode $m = 2$, lockdown

related costs are equal to $\frac{1200}{63.7}$, as before, and for mode $m = 1$ we consider three cases: $\delta_L^1 = \frac{1}{3}\delta_L^2$, $\frac{1}{2}\delta_L^2$, and $\frac{2}{3}\delta_L^2$. Finally, switching costs are $K_{01} = K_{02} = 2000$ and $K_{12} = K_{21} = K_{10} = K_{20} = 0$. Assuming $K_{12} = 0$ allows us to consider scenarios where the social planner decides to never stay in mode 2 with a positive period of time and to therefore decide to switch to mode $m = 0$ from $m = 1$, thereby essentially skipping mode $m = 2$. Later we will study what happens if $K_{12} > 0$. We will again use blue to denote states where it is optimal to leave mode 0 and, for consistency, we use orange for states where it is optimal to leave the mode with the highest R_0 , in this case mode 2. We introduce brown as the colour related to coming out of our, new, intermediate mode.

Fig. 12 illustrates the stopping sets for $\mathcal{R}_0^1 = 2$ and $\delta_L^1 = \frac{1}{2}\delta_L^2$. Panel (a) illustrates the sets for the scenario with only 2 modes as in the previous section. Panel (b) illustrates the stopping sets for the scenario with 3 modes. The brown area contains all states of the stopping set where it is optimal to switch from the intermediate mode to mode 0, and the orange area contains states that are part of the stopping set for switching from a full lockdown to social distancing. Although not visible, states of the brown set are also states of the stopping set for switching from full lockdown to social distancing: if (I_0, R_0) is in the brown area, then it is optimal to switch from mode 2 immediately to mode 0. The orange set in Panel (a) contains fewer states than the orange (plus brown) region in Panel (b), but more than the brown region itself.

The blue area consists of all points where it is optimal to switch from mode 0 to mode 1. Panel (c) distinguishes the sets for switching to mode 1 (blue) from mode 0 and to mode 2 (blue and light blue) from mode 1. This panel illustrates that all states for which it is optimal to introduce social distancing form a subset of area where switching to a full lockdown is optimal. Therefore, the only optimal policy is to implement a full lockdown. Numerical analysis




Fig. 12. Stopping sets with three modes with $\mathcal{R}_0^1 = 2$ and $\delta_L^1 = \frac{1}{2}\delta_L^2$. Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 1 and 1 to 2 (light). $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^2 = 1.5$, $\gamma_m = 0.1$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 13. Stopping sets with three modes with different \mathcal{R}_0^1 and δ_L^1 . Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^2 = 1.5$, $\gamma_m = 0.1$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 14. Effect of K_{12} on stopping sets. Blue: switching from mode 0 to 1, light blue: from mode 1 to 2 (light). $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^2 = 2$, $\mathcal{R}_0^1 = 1.5$, $\gamma_m = 0.1$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

shows that this result holds for a wide range of a parameter values. Further, notice that there are two light blue areas in Panel (c). The set close to the line $i + r = N$ denotes all points such that, when currently in mode 1, it is optimal to switch to mode 0 without spending a positive amount of time in mode 2, since these states are also part of the orange and brown area.

Fig. 13 illustrates what happens when changing δ_L^1 and \mathcal{R}_0^1 . In all panels, the blue area is not qualitatively different, which does not come as a surprise since no time is spent in mode 1 after leaving mode 0. Panels (a) and (b) show the stopping sets for $\mathcal{R}_0^1 = 2$ with $\delta_L^1 = \frac{2}{3}\delta_L^2 > \frac{1}{2}\delta_L^2$ and $\delta_L^1 = \frac{1}{3}\delta_L^2 < \frac{1}{2}\delta_L^2$, respectively. When staying in mode 1 is more expensive, it is no longer optimal to switch to mode 1 before switching to mode 0, when currently in mode 2. The opposite applies when the lockdown related cost is smaller so that the region where the social planner switches from mode 1 to mode 0 shrinks, and the region where the social planner switches from mode 2 to mode 1 expands. For the cases where $\mathcal{R}_0^1 = 2.25$, it is also less attractive to spend a positive amount of time in mode 1, as illustrated by Panel (c) and Panel (d).

Finally, Fig. 14 illustrates the effect of K_{12} , assuming that the social planner always switches to mode 2 before lifting control mea-

sures. As switching to mode 2 becomes more expensive, the set of states where it is optimal to switch shrinks, in line with the intuition for our main model. However, the light blue area has not shrunk considerably, which means that once switched to mode 1, from mode 0, it remains optimal to immediately switch to mode 2 for almost all states.

4.5. Open border policies

This section studies the effect of a strict border policy on lockdown timing decisions. Thus far it has been assumed that no new infections can occur when states are reached where $I_t = 0$. We continue to make this assumption when in mode $m = 2$, but not when in mode $m = 0$. For mode $m = 1$ the social planner can choose whether or not to restrict entry at the border. Note that no infections occurring when $I_t = 0$ is not only equivalent to a closed border policy but could also include, e.g., the requirement of individuals entering the country to quarantine in a hotel. In mode $m = 1$, the social planner could opt to ask individuals to self-isolate so that the infection rate is somewhere between the infections rates of modes 0 and 2.

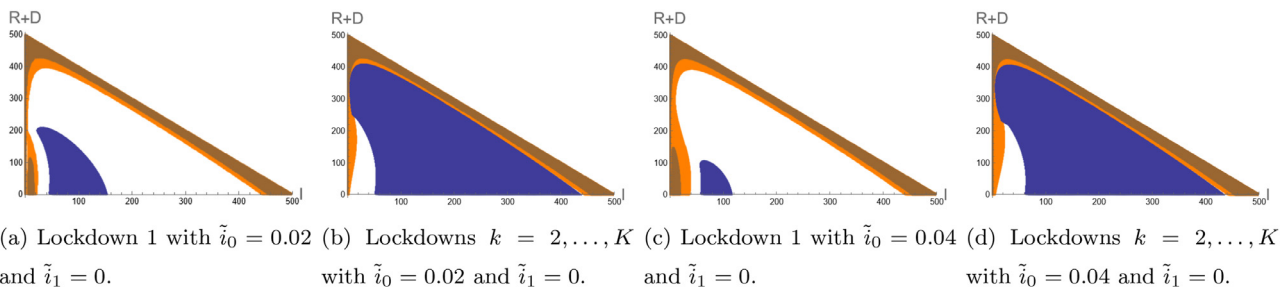


Fig. 15. Stopping sets with different border policies. Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 2. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_1^0 = 2$, $\mathcal{R}_2^0 = 1.5$, $\gamma_m = 0.1$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^0 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^0/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

We model this extension using the infection rate

$$\lambda_m(i, r) = \beta_m(i + \tilde{i}_m) \frac{S(i, r)}{N},$$

where $\tilde{i}_2 = 0$ and $0 \leq \tilde{i}_1 \leq \tilde{i}_0$ are constants.¹⁷

Results Fig. 15 illustrates how the stopping sets have changed when $\tilde{i}_1 = \tilde{i}_2 = 0$. Panels (a) and (b) illustrate the case where $\tilde{i}_0 = 0.02$, while Panels (c) and (d) depict the case $\tilde{i}_0 = 0.04$. We first note that the blue area for the first lockdown has significantly shrunk, relative to a situation where $\tilde{i}_0 = 0$. Our numerical exploration reveals that for any lockdown to ever be potentially optimal, \tilde{i}_0 should not be larger than 0.05, i.e., for any larger risks it is never optimal to impose restrictions. Intuitively, any increase in the infection rate leads to an increase in costs: it takes longer, on average, to bring the number of infectious individuals down and, in addition, more infections occur. This leads to our conclusion that because open borders heavily reduce the effectiveness from lockdown, any gains from lockdown cannot weigh against the (sunk) cost involved with imposing restrictions.

The blue areas in Panels (b) and (d) are very similar, i.e., the main impact of a change in the infection rate from abroad has an impact on strategy for the first lockdown only. Thirdly, the orange area has slightly expanded for the states where I_t is close to zero. Recall that the orange area is the set of states for which it is optimal to leave the mode with the highest infection rate, in this case mode $m = 2$. Because the increase in the infection rate has the largest impact on these states in particular, the net gain from staying in mode $m = 2$ is smaller when $\tilde{i}_0 > 0$ and the social planner prefers to switch to mode $m = 1$. Note that the brown area, the area where it is optimal to come out $m = 1$, has barely expanded so that it remains optimal to keep measures in place.

It is essential to notice, though, that although there exists a small brown region near the origin, it is unlikely that these states are reached. Because $I_t + R_t + D_t$ cannot decrease over time,¹⁸ whenever the blue area is reached, for the process to reach this brown region subsequently we would require: a much higher than expected number of infections before the blue region is reached and a much higher than expected number of recoveries after lockdown is imposed while in the mean time close to nil infections occur.

Fourthly, we find that there are states where it is optimal, once switched to mode 1, to never switch to mode 2. These are states where the blue area overlaps with the orange area. In these states, the social planner switched immediately from mode 0 to mode 1,

in order to switch back to mode 0 in the future. For these states the epidemic has evolved too much for a full lockdown to be beneficial.

Finally, Fig. 16 illustrates that a positive value for \tilde{i}_1 does not qualitatively change the social planner’s decision and that an increase has no significant impact on whether or not lockdown is optimal for positive \tilde{i}_0 . This is not surprising in light of earlier figures where we showed that “phasing in” lockdowns by entering mode 1 before moving to mode 2 are not optimal.

4.6. Health-care imposed capacity constraints

Next we study the effect of constraints on healthcare capacity, by generalizing δ_E in the following way:

$$\delta_E(I_t, R_t) = \begin{cases} 4 & \text{if } I_t < 150, \\ 8 & \text{if } I_t \geq 150. \end{cases} \quad (13)$$

Here, 150 is the assumed threshold that represents a limitation of the health care system. It represents the fact that (health-related) costs go up once the health system has reached capacity. We need to keep in mind that only a fraction of infectious individuals need hospital treatment and therefore the number 150 should be interpreted as the equivalent number of total infections such that the fraction of infectious individuals that requires hospital treatment is equal to the imposed “critical care capacity”.¹⁹ Fig. 17 shows the results of adapting the baseline parameterization. This figure illustrates that the stopping set in mode 0 now includes states where either the number of infectives is above 150 or where the risk is sufficiently large to reach states where the number of infectives is above 150. As illustrated in Panel (a), states in the extended region with $I_t < 150$ are those where it is expected that the process will reach states where $I_t \geq 150$ when remaining in mode 0. This explains the shape of the border of the stopping set where there seems to be an “inaction” region.

However, note that when taking into account the stream plots in Fig. 8, we can see that the expected path, when starting around (0,0), does not go through any of the ‘newly added’ states. Thus, the extended area includes states where the decision maker wishes to bring down the number of infections rapidly when, against the odds, a spike of infections has occurred.

Also notice that the orange region has shrunk, especially for all states where $I_t > 150$: the social planner delays the moment to switch back to mode 0, so that the expected amount of time spent in states for which $I_t > 150$ is minimized.

An alternative way to model hospital constraints is as follows. Governments may wish to avoid to have a number of infectious individuals that is near to the constraint. Therefore, δ_E could be mod-

¹⁷ In this extension, in order to be able to focus on the impact of \tilde{i}_m , additional effects that stem from costs associated with restrictions at the border are not taken into account.

¹⁸ Note that this effectively means that, when the process is in a given state, the process can only either move to the east (when a new infection occurs) or north-west (when a recovery or death occurs).

¹⁹ In November 2021, in the UK, around 2% of people who tested positive were admitted into hospital. Source: <https://coronavirus.data.gov.uk/>

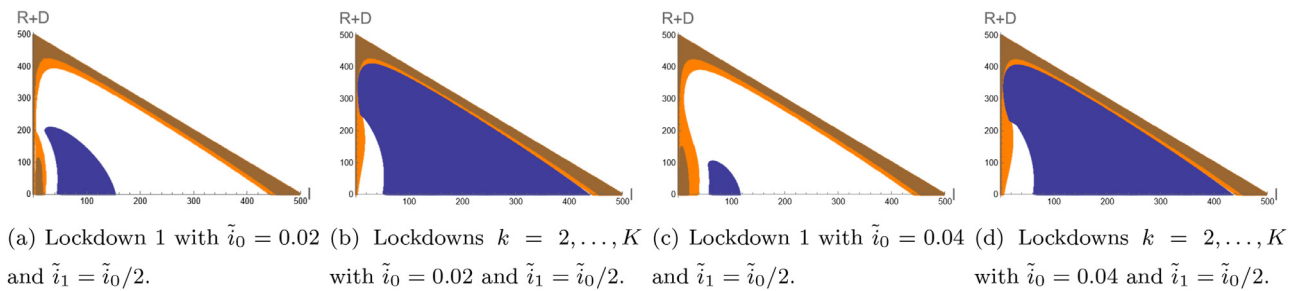


Fig. 16. Stopping sets with different border policies. Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 2. $\mathcal{R}_0^0 = 3, \mathcal{R}_0^1 = 2, \mathcal{R}_0^2 = 1.5, \gamma_m = 0.1, \phi_m = 2\%, N = 500, \rho = 0.1/365, K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4, \delta_L^2 = \frac{1200}{63.7}, \delta_L^1 = \delta_L^2/2,$ and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 17. There is a limit on hospitals admissions that leads to an equivalent of a limit of 150 on infections (see (13)). Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 2. $\mathcal{R}_0^0 = 3, \mathcal{R}_0^1 = 2, \mathcal{R}_0^2 = 1.5, \gamma_m = 0.1, \tilde{i}_0 = 0.02, \tilde{i}_1 = \tilde{i}_2 = 0, \phi_m = 2\%, N = 500, \rho = 0.1/365, K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_L^2 = \frac{1200}{63.7}, \delta_L^1 = \delta_L^2/2,$ and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

eled as an increasing function of I_t when I_t gets close to the constraint imposed by the critical care capacity. Appendix B.1 shows that, in that case, the figures look qualitatively the same.

4.7. Extension: time-dependent cost structures

In this section, we relax our assumption that lockdown related costs are constant over time. One could argue that the length of lockdown has an impact on the associated lockdown related costs and the willingness of individuals to comply. Especially for business, a lockdown lasting for a longer period of time will have a long lasting impact and comes with more economic damage per unit of time than a lockdown imposed for a shorter period of time. Here, we look at our SIRD-like model as presented in Section 4.2 with $\tilde{i}_0 > 0$, but without health-care constraints. We consider the following cost structures,

$$c_m(i, r) = \delta_E(i, r)i + \delta_L(1 + 0.07\sqrt{t})(1 - \chi_{\{m=0\}}),$$

$$c_m(i, r) = \delta_E(i, r)i + \delta_L(1 + 0.01t)(1 - \chi_{\{m=0\}}), \text{ and}$$

$$c_m(i, r) = \delta_E(i, r)i + \delta_L(1 + 0.00015t^2)(1 - \chi_{\{m=0\}}).$$

In this section we analyse the second case with a linear relationship; the convex and concave cases can be found in Appendix B.2, where we conclude that the results are qualitatively the same. We start by assuming that the infection rate \tilde{i}_0 is low, to clearly identify the direct impact of the cost structure. We assume that $\tilde{i}_0 = 0.002$ and $\tilde{i}_1 = 0$. To see the impact of an increase in costs, consider that when lockdown lasts, e.g., 80 days, then $0.01(80) = 0.8$, so that the daily costs increase by £15,072 for the linear specification. This is equivalent to, by the end of lockdown, £2b per day for a population of 67.22m in the UK.

Fig. 18 illustrates the stopping sets for the linear case. The orange area in Fig. 18a illustrates the stopping set for coming out of lockdown when $t = \tau_{C_0}$, i.e., the moment when the social planner decides to impose restrictions. Fig. 18b and c illustrate that the orange area grows larger as time goes by, i.e., the social planner comes out of lockdown earlier when costs are higher. However, despite the daily equivalent increase of £760m for (b) and £1.5b for (c), the increase does not significantly change the shape of the orange area. The blue area corresponds, as before, to the decision of going into lockdown, which does not depend on time and it therefore remains constant. Interestingly, there has appeared an “inaction” for states where both I_t and $R_t + D_t$ are small but positive. For these states going into lockdown might mean that lockdown lasts too long meaning that the social planner incurs large costs, instead the social planner waits until I_t has sufficiently grown to limit the length of lockdown. However, when I_t is very close to 0, it is optimal to keep I_t low and impose lockdown. Since for these states, as the stream plots in Fig. 8 in Section 4.1 illustrated, lockdown may be in place for a short period of time, the increase in costs per unit of time does not impact the inclusion of these states in the stopping set. Hence, we see that the social planner comes out of lockdown slightly earlier and may delay lockdown only when it is discovered too late. Fig. 19 illustrates what happens when \tilde{i}_0 is increased. Then the inaction region grows but the impact on the orange sets is negligible.

4.8. Comparative statics

Appendix C covers an extensive comparative statics analysis for our parameters. Here we briefly summarize the results.




Fig. 18. Stopping sets when lockdown related costs are increasing over time with $\tilde{i}_0 = 0.002$. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 2$, $\mathcal{R}_0^2 = 1.5$, $\gamma_m = 0.1$, $\tilde{i}_0 = 0.002$, $\tilde{i}_1 = \tilde{i}_2 = 0$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 19. Stopping sets when lockdown related costs are increasing over time with $\tilde{i}_0 = 0.004$. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 2$, $\mathcal{R}_0^2 = 1.5$, $\gamma_m = 0.1$, $\tilde{i}_0 = 0.004$, $\tilde{i}_1 = \tilde{i}_2 = 0$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

First we look at an increase in the infection rate, which may result from, e.g., a mutation, and find that the blue area shrinks. This happens, because when the virus is more infectious, the expected interarrival times are shorter thereby reducing the length of a potential lockdown but also making lockdown less effective. The social planner therefore also comes out of lockdown earlier. We find the (intuitively clear) opposite effect when measures are more effective so that $\mathcal{R}_0^2 < 1$.

Next we consider how the recovery rate γ_m impacts the social planner’s strategy. There are mixed effects on the stopping sets when increasing γ_m , where one effect generally dominates: the blue area shrinks when the recovery rate goes up. In addition, the social planner expands the orange area.

Finally, Appendix C.3 illustrates that increasing δ_E , δ_L , or K_{01} has a similar impact: the blue area shrinks. The orange area expands for the first two cases.

In all cases we find that the figures look qualitatively the same and we therefore conclude that our results are robust against (moderate) changes in parameter values.

5. Concluding remarks and future research

Motivated by the recent COVID-19 outbreak, we have developed a continuous-time Markov chain to study the optimal timing of interventions in an SIRD-inspired epidemiological model of the evolution of a disease. A social planner has to decide when to enter “lockdown” and, subsequently, when to lift it. Although, traditionally, epidemiology models are assuming a deterministic evolution of the disease, there is still uncertainty about how the disease spreads, especially when the number of infected individuals is

low. In such a scenario the disease could die out before the disease spreads. Nonetheless, surprisingly, we find that it is optimal to enter lockdown in the very early stages of the disease. Moreover, it is found that it is never optimal to introduce a lockdown when the prevalence of the disease is too high, i.e. a lockdown is optimally started only at low prevalence of the disease. In addition, despite high economic cost, it is optimal for the social planner to wait with exiting the lockdown until either the fraction of susceptible members of the population or infected members of the population is close to zero. This holds, even when the lockdown-related costs are increasing over time.

If there is a capacity constraint on the health system, which is modeled by a jump in the per-infective cost once the number of infectives exceeds a given threshold, then the region where never entering lockdown is optimal shrinks. In addition, lockdown should be kept in place for longer.

When we allow for a phased introduction and exiting of lockdown (e.g., when using a 3-tier traffic-light system), it is found that, while a phased exiting strategy is optimal, a phased introduction is not. Our analysis also shows that lockdowns are only optimal when entry into the country is restricted.

Our model has several limitations that present avenues for future research. First, we assume that the number of infected individuals is known. In reality this only holds by approximation, even when testing is done frequently and on a large scale. This provides scope for research on models where decisions and policies are made based on imperfect information. Since we find that immediate action is optimal when the first cases of infected individuals have been detected, we expect that such an extension does not lead to qualitatively different optimal policies. Nevertheless, it

stresses the importance of testing, especially since we find that if the epidemic is detected too late, then intervention is no longer optimal.

Other potential extensions of our work include considering a heterogeneous host population, by extending the SIRD framework to include more classes of individuals (as discussed to some extent in Section 4.2). An interesting, but technically substantially more demanding, extension for future research would be to allow recovered agents to become susceptible again. Although herd immunity might be more cost effective than an infinite loop of lockdown measures, such a model would include the expectation that over time the virus becomes less dangerous and as such we expect that the decision to come out of lockdown is not qualitatively different. In addition, the result that lockdown is not optimal when the virus is discovered too late, might still hold. Having said that, it seems that many governments, at least in many EU countries, currently seem to be willing to impose restrictions every winter for the foreseeable future, which is in line with how we have modeled the problem of the social planner.²⁰

Appendix A. Proofs

Proof of Proposition 1. For any measurable function $f : E \rightarrow \mathbb{R}$, $x \in E$ and $m = 0, 1$ it holds that,

$$\begin{aligned} E_x^m[df(X_t)] &= \sum_{y \in E} (f(y) - f(X_t)) P_x^m(\{X_{t+dt} = y\} | X_t) \\ &= \sum_{y \neq X_t} q_{X_t, y}^m f(y) dt + q_{X_t, X_t}^m f(X_t) dt. \end{aligned}$$

Here we have used the fact that $q_{xx}^m = -\sum_{y \neq x} q_{xy}^m$. Hence, it follows that the Bellman equation

$$\rho G(x) = \lim_{dt \downarrow 0} \frac{1}{dt} E_x^{m=0}[dG(x)] + c_0(x),$$

can be written as

$$\rho G(x) = \sum_{y \neq x} q_{xy}^{m=0} G(y) + q_{xx}^{m=0} G(x) + c_0(x).$$

Rewriting leads to (5). \square

Proof of Proposition 2. We first assume that φ solves (7). Let τ_C denote the first exit time of any set $C \in \mathcal{E}$. Note that (8) can be written as

$$\min\{G(x) + K_{10} - \varphi(x), \mathcal{L}_x^{m=1}\varphi(x) + c_1(x)\} = 0, \quad x \in E,$$

where \mathcal{L}_x^m denotes the generator of $X = (X_t)_{t \geq 0}$ for mode m , $m = 0, 1$. Since it holds that

$$E_x^m \varphi(X_{dt}) = \sum_{y \in E} \varphi(y) P_x^m(\{X_{dt} = y\}) = \sum_{y \neq x} q_{xy}^m \varphi(y) dt + \varphi(x)(1 + q_{xx}^m dt),$$

it follows that

$$\mathcal{L}_x^m \varphi(x) = \lim_{dt \downarrow 0} \frac{E_x^m \varphi(X_{dt}) - \varphi(x)}{dt} - \rho \varphi(x) = \sum_{y \in E \setminus \{x\}} q_{xy}^m \varphi(y) + (q_{xx}^m - \rho) \varphi(x),$$

for all $x \in E$. Then,

$$\mathcal{L}_x^{m=1} \varphi(x) \geq -c_1(x), \quad x \in E, \quad \text{and}$$

$$\mathcal{L}_x^{m=1} \varphi(x) = -c_1(x), \quad x \in C_1.$$

Fix $T > 0$. From Dynkin's formula it follows that

²⁰ A model inspired by SIS frameworks where transitions are modeled through a CTMC as well, similar to our work, is studied by Huberts & Thijssen (2021) who consider the investment problem of a monopolist in a market with network externalities. The underlying state process becomes mean reverting, which is in line with our intuition for the extension that is discussed here.

$$\begin{aligned} E_x^{m=1} [e^{-\rho(\tau_{C_1} \wedge T)} \varphi(X_{\tau_{C_1} \wedge T})] &= \varphi(x) + E_x^{m=1} \left[\int_0^{\tau_{C_1} \wedge T} e^{-\rho t} \mathcal{L}_x^{m=1} \varphi(X_t) dt \right] \\ &= \varphi(x) - E_x^{m=1} \left[\int_0^{\tau_{C_1} \wedge T} e^{-\rho t} c_1(X_t) dt \right]. \end{aligned}$$

Since

$$\begin{aligned} E_x^{m=1} [e^{-\rho(\tau_{C_1} \wedge T)} \varphi(X_{\tau_{C_1} \wedge T})] &= E_x^{m=1} [e^{-\rho(\tau_{C_1})} (G(X_{\tau_{C_1}}) + K_{10}), \tau_{C_1} < T] \\ &\quad + E_x^{m=1} [e^{-\rho(T)} \varphi(X_T), \tau_{C_1} > T] \\ &\rightarrow E_x^{m=1} [e^{-\rho(\tau_{C_1})} (G(X_{\tau_{C_1}}) + K_{10})], \quad \text{as } T \rightarrow \infty, \end{aligned}$$

we find that

$$\varphi(x) = E_x^{m=1} \left[\int_0^{\tau_{C_1}} e^{-\rho t} c_1(X_t) dt + e^{-\rho(\tau_{C_1})} (G(X_{\tau_{C_1}}) + K_{10}) \right].$$

Now take any stopping time $\tau \in \mathcal{M}$ and fix $T \geq 0$. Then

$$\begin{aligned} E_x^{m=1} \left[\int_0^{\tau \wedge T} e^{-\rho t} c_1(X_t) dt + e^{-\rho(\tau \wedge T)} (G(X_{\tau \wedge T}) + K_{10}) \right] \\ \geq -E_x^{m=1} \left[\int_0^{\tau \wedge T} e^{-\rho t} \mathcal{L}_x^{m=1} \varphi(X_t) dt \right] + E_x^{m=1} [e^{-\rho(\tau \wedge T)} (G(X_{\tau \wedge T}) + K_{10})] \\ = \varphi(x) - E_x^{m=1} [e^{-\rho(\tau)} \varphi(X_\tau)] + E_x^{m=1} [e^{-\rho(\tau)} (G(X_\tau) + K_{10})] \\ \geq \varphi(x). \end{aligned}$$

Since T was chosen arbitrarily, the inequality also holds as $T \rightarrow \infty$. Therefore, $\varphi = F$ and τ_{C_1} is the optimal stopping time.

Existence and uniqueness follow from the fact that F is the solution to the fixed point problem

$$F = \Gamma(F),$$

where

$$\Gamma(F) = \min \left\{ G(x) + K_{10}, \frac{c_1(x)}{\rho - q_{xx}^{m=1}} + \sum_{y \neq x} \frac{q_{xy}^{m=1}}{\rho - q_{xx}^{m=1}} F \right\}.$$

Note that the function F can be thought of as a vector in \mathbb{R}^n , where $n = |E|$. It is easy to check that Γ satisfies Blackwell's conditions, so that Γ is a contraction mapping (Stokey & Lucas, 1989, Theorem 3.3). Existence and uniqueness of F then follow from the Banach fixed point theorem (Stokey & Lucas, 1989, Theorem 3.2).

As a final note, since for all $x \in \Lambda$ it holds that $q_{xy}^m = 0$ for all $y \in E$, $\varphi(x)$ is determined unambiguously. For $x \in E \setminus \Lambda$, $\varphi(x)$ is recursively determined as a function of the states the chain can transition to when leaving x . Hence, by induction, we establish that φ is unambiguously determined in (8) for all $x \in E$. \square

Proof of Proposition 3. The proof is analogous to the proof of Proposition 2. \square

Appendix B. Further analysis

In this Appendix, we provide some additional figures for Sections 4.6 and 4.7.

B1. Critical care capacity constraints

Here we investigate how the stopping sets are impacted when the epidemic related costs are increasing in I_t when I_t gets close to the threshold of 150, rather than the cliff-edge from Section 4.6, i.e., when

$$\delta_E(I_t, R_t) = \begin{cases} 4 & \text{if } I_t < 100, \\ 4(1 + (I_t - 100)/50) & \text{if } 100 \leq I_t < 150, \\ 8 & \text{if } I_t \geq 150. \end{cases} \quad (14)$$

Fig. 20 illustrates the corresponding stopping sets and we find that these look qualitatively the same. The most noticeable difference is for states such that $I_t \in (100, 150)$ since for these states the costs are higher compared to (13).

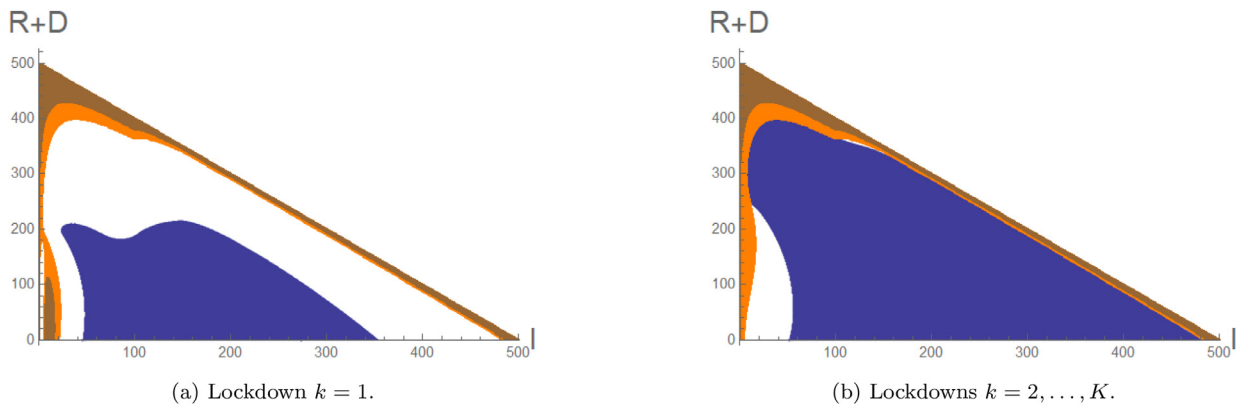


Fig. 20. Stopping sets when the per-patient costs are given by (14). Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 2. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 2$, $\mathcal{R}_0^2 = 1.5$, $\gamma_m = 0.1$, $\tilde{i}_0 = 0.02$, $\tilde{i}_1 = \tilde{i}_2 = 0$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_L^2 = \frac{1200}{83.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

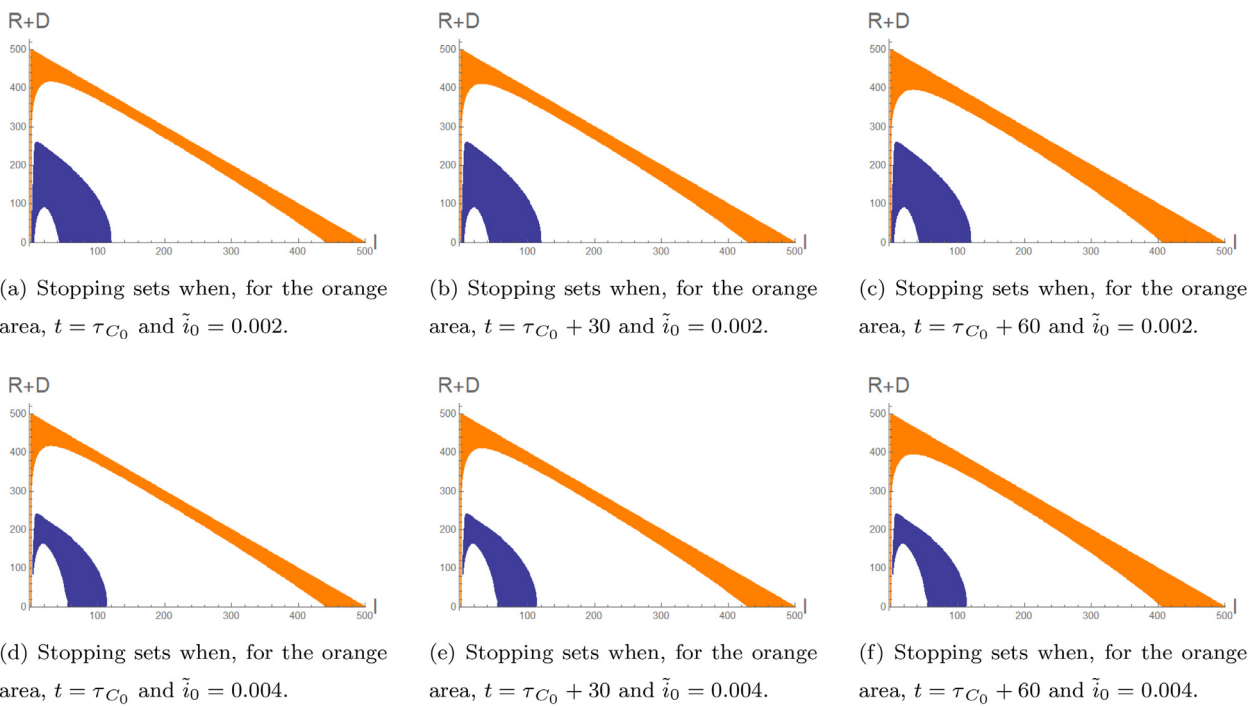


Fig. 21. Stopping sets when the increase in lockdown related costs is $\delta_L 0.00015t^2$.

B2. Time-dependent cost structures

Figs. 21 and 22 illustrate the stopping sets equivalent to those in Figs. 18 and 19. These figures illustrate that when lockdown costs are convex in t or concave in t , the resulting stopping sets are qualitatively the same.

Appendix C. Robustness

C1. Infection rate

Here we investigate how the optimal policy is affected by the discovery of a more infectious mutant virus, as well as the effect of a lower infection rate due to, e.g., effective measures.

A higher infection rate β_m , $m = 0, 1, 2$, has two effects on the social planner’s problem. First, infections occur more often, relative to a case with a low infection rate. This means that states with a high number of infectives are reached with a higher probability, leading to a bigger need for lockdown. Second, a higher infection

rate reduces the expected interarrival times, so that the process moves through more states during the same time interval. Especially for the part of the state space with a high number of infectives, the cost of staying in lockdown may no longer outweigh the benefits of staying in lockdown. It becomes more interesting for the social planner to phase out restrictions for a larger set of states, with the potential consequence that less time is spent in lockdown. It is therefore, *a priori*, not clear what the net result will be on the (expected) switching times.

To illustrate these conflicting incentives, we compare our baseline parameterization to a case where $\mathcal{R}_0^2 = 5$. In line with the literature we again assume that $\mathcal{R}_0^0 = \mathcal{R}_0^2/2$ and we set $\mathcal{R}_0^1 = 3.5$. Although we appreciate that $\mathcal{R}_0^1 = 5$ is high, this choice allows for the net effects on the stopping sets to be more distinguishable.

Fig. 23 illustrates the stopping sets for this parameterization. For the first lockdown, the blue region has shrunk indicating that only when still very few people have been infected switching to mode $m = 2$ is optimal. Discovering the epidemic too late and thereby never imposing lockdown, is more likely when the so-

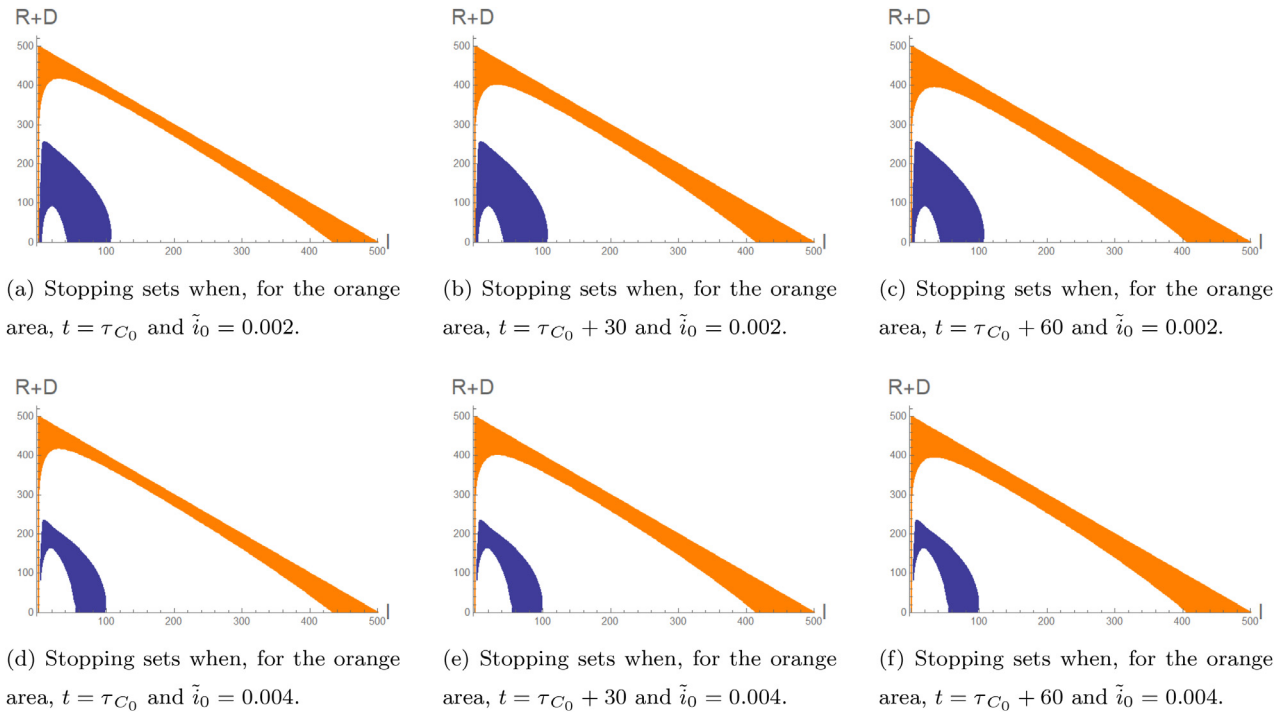


Fig. 22. Stopping sets when the increase in lockdown related costs is $\delta_L 0.07\sqrt{t}$. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 2$, $\mathcal{R}_0^2 = 1.5$, $\gamma_m = 0.1$, $\tilde{i}_0 = 0.002$, $\tilde{i}_1 = \tilde{i}_2 = 0$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

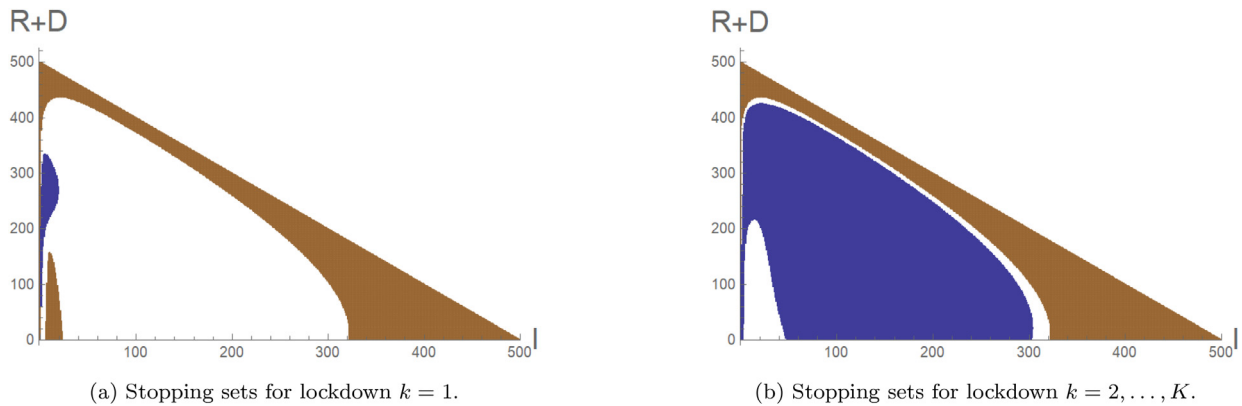


Fig. 23. Stopping sets for lower infection rates. Here, $\mathcal{R}_0^m = 5$ for $m = 0$, $\mathcal{R}_0^1 = 3.5$, and $\mathcal{R}_0^2 = 2.5$. Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 2. $\mathcal{R}_0^0 = 5$, $\mathcal{R}_0^1 = 3.5$, $\mathcal{R}_0^2 = 2.5$, $\gamma_m = 0.1$, $\tilde{i}_0 = 0.002$, $\tilde{i}_1 = \tilde{i}_2 = 0$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 24. Stopping sets for lower infection rates. Here, $\mathcal{R}_0^m = 0.9$ for $m = 2$, $\mathcal{R}_0^1 = 1.2$, and $\mathcal{R}_0^0 = 1.8$. Brown: switching from mode 1 to mode 0, orange: from mode 2 to mode 1, blue: from mode 0 to 2. $\mathcal{R}_0^0 = 1.8$, $\mathcal{R}_0^1 = 1.2$, $\mathcal{R}_0^2 = 0.9$, $\gamma_m = 0.1$, $\tilde{i}_0 = 0.002$, $\tilde{i}_1 = \tilde{i}_2 = 0$, $\phi_m = 2\%$, $N = 500$, $\rho = 0.1/365$, $K_{01} = K_{02} = 2000$ and $K = 0$ otherwise, $\delta_E = 4$, $\delta_L^2 = \frac{1200}{63.7}$, $\delta_L^1 = \delta_L^2/2$, and $\delta_D = \frac{10.88}{365}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 25. Expected number of infected (brown), susceptible (blue), and recovered (red) individuals for mode 0 (Panel (a)) and mode 1 (Panel (b)). Panel (c) juxtaposes the expected number of infectives in both modes. Here, $(I_0, R_0) = (10, 0)$. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.3$, $N = 500$, and $\rho = 0.1/365$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

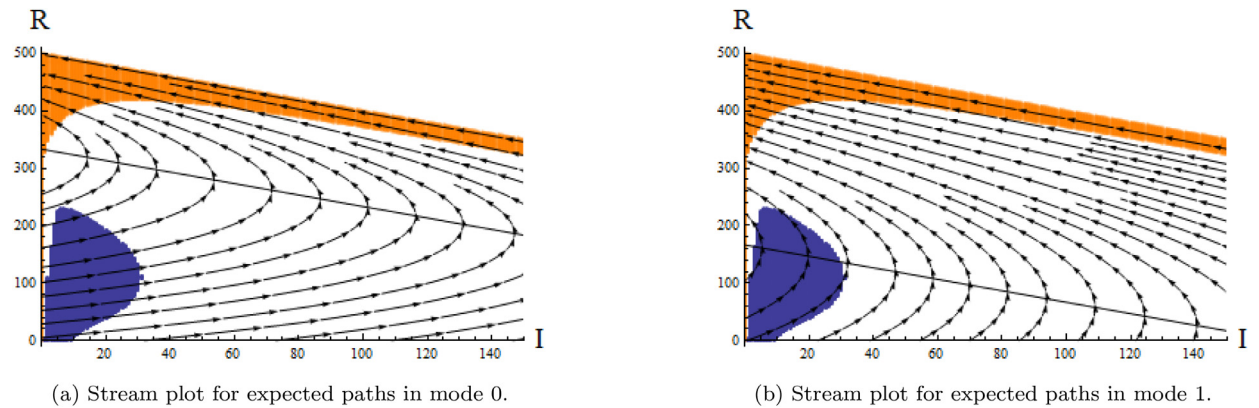


Fig. 26. Stopping sets with vector/stream plots. Here, γ_m has been multiplied by 3 to increase the volatility of the process. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.3$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 1500$, $K_{10} = 0$, $\delta_E = 4$, and $\delta_L = \frac{1200}{63.7}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 27. Stopping sets with stream plots with a different parametrization for K_{01} (Panel (a)) and δ_L (Panel (b)). Stream plot for expected paths in mode 1. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{10} = 0$, and $\delta_E = 4$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

cial planner is not proactive enough in detecting the virus. At the same time, the orange and brown areas have expanded. Although not visible, the brown area now entirely intersects with the orange area, which means that phasing out lockdown is no longer optimal. In short, lockdown is only optimal very early on in the epidemic and one comes out of lockdown entirely.

For subsequent lockdowns, we also find that phasing out lockdown is not optimal. Note that in the blue area an “inaction” region appears again as we found before in Section 4.7. The intuition is the same as before: for states where I_t is close to zero, a lock-

down of a short length can be imposed. However, when I_t is sufficiently positive and R_t is small, a small delay is optimal. This, thus, occurs even when costs are not increasing over time.

These illustrations confirm that, especially when considering situations where a state with a high number of infectives is reached, there may be a higher probability of lockdown duration being shorter in a scenario with a higher infection rate. The fact that this probability can increase underlines the importance of incorporating the stochastic nature of the development of the disease, i.e., effects as these cannot be observed in deterministic spec-

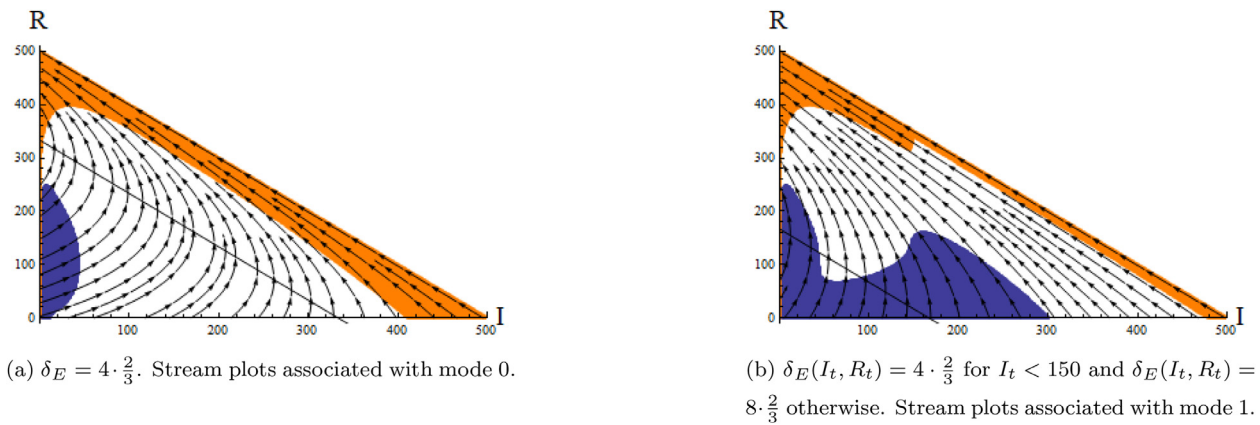


Fig. 28. Stopping sets with stream plots. Here, $\delta_L = 20$. Orange: switching from mode 1 to mode 0, blue: switching from mode 0 to 1. $\mathcal{R}_0^0 = 3$, $\mathcal{R}_0^1 = 1.5$, $\gamma_0 = \gamma_1 = 0.1$, $N = 500$, $\rho = 0.1/365$, $K_{01} = 2000$, $K_{10} = 0$, and $\delta_L = \frac{1200}{63.7}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

ifications. The most interesting result is the shrinkage of the blue area, so that it becomes even more important to discover the epidemic early.

Fig. 24 illustrates the opposite case. Effective policies during the epidemic, such as social distancing across all states, may lead to a situation where the R_0 is effectively below 1. Therefore, we now look at a case where $\mathcal{R}_0^{m=2} = 0.9$. The blue area has now expanded and the orange area has shrunk, which is in line with the intuition above. Also, in the blue area for the first lockdown the inaction region appears again.

C2. Recovery rate

In this section we provide an illustration of the effect of the recovery rate γ_m . Since $(E_x^m dI_t)^2 = 0$ for all $m = 0, 1$, we obtain

$$\begin{aligned} \text{Var}(dI_t) &= E(dI_t)^2 - 0 = [\lambda_m(I_t, R_t) + \mu_m(I_t, R_t)]dt \\ &= \gamma_m I_t \left(\mathcal{R}_0^m \frac{S(I_t, R_t)}{N} + 1 \right) dt. \end{aligned}$$

Similarly, $\text{Var}(dR_t) = E_x^m (dR_t)^2 - 0 = \mu_m(I_t, R_t)dt = \gamma_m I_t dt$. This means that increasing γ_m while keeping \mathcal{R}_0^m fixed leads the process to move quicker.

For our illustration we multiply all γ_m by a factor of 3 and reduce the sunk cost of entering lockdown to $K_{01} = 1500$. The latter change is needed because otherwise the stopping set would (almost) disappear and it would never be optimal to enter lockdown. To further illustrate the analysis in the main text, Fig. 25 reproduces Fig. 7 with the increased transition rates and illustrates that the shapes of the curves have remained qualitatively the same, but the time scale has changed. Also note that the share of host population that remains susceptible throughout has not (significantly) changed. This confirms that, at least in expectation, the process moves faster for higher γ_m without significantly changing the expected numbers of infectives and susceptibles.

An increase in γ_m has three consequences. First, an increase in the infection rate increases, in expectation, the number of infectives per period of time. Second, for any future point in time the standard deviation of susceptibles, infectives, and recovered goes up. These two reasons make the stopping set expand. On the other hand, the period over which costs are incurred is shorter, which should make switching to mode 1 less attractive as the investment required to realize a lockdown and consequently the economic damage might not be outweighed by the gain from “flattening the curve”. Flattening the curve has two cost advantages. On the one hand the peak of infectives is reached later, which leads

to a lower present value in absolute terms. At the same time, the number of total infectives is (in expectation) lower as well, which directly reduces cost.

Fig. 26 shows that both effects are present but that they affect different parts of the stopping sets. First, the stopping set for mode 0 has considerably shrunk. The period over which costs are incurred is shorter, so that lockdown is optimal for a smaller range of states.

At the same time, the stopping set in mode 1 has expanded. Despite the increase in volatility, the speed with which the number of infectives goes to zero is sufficient to make it optimal to lift lockdown in a larger set of states. In addition, the period over which costs are incurred is shorter which makes switching more attractive.

C3. Lockdown-related costs

We now illustrate the sensitivity of the stopping sets to changes in parameter values associated with incurred costs, starting with the sunk cost of entering lockdown, K_{01} , which is illustrated by Fig. 27(a). Entering lockdown is optimal when the gain from flattening the curve outweighs the additional costs involved with entering lockdown. When K_{01} increases, the benefits might no longer outweigh the costs for certain states and, therefore, the stopping set in mode 0 shrinks. Since K_{01} has no effect on the value in mode 1, the stopping set in mode 1, i.e., the orange area, is not affected.

Fig. 27 (b) illustrates what happens when δ_L is increased. Since these costs are only incurred while being in mode 1, there are fewer cases where switching to mode 1 is optimal and more cases where switching from mode 1 back to mode 0 is optimal.

Fig. 28 illustrates the effect of changes in (epidemic related) per-patient cost $\delta_E(i, r)$, incurred in all modes. Here, δ_E is one-third lower than in the baseline parameterization. When the value of δ_E is lower for all states, the gain of lockdown is lower and, therefore, there is less of a necessity for the social planner to switch to mode 1. This leads to the stopping set in mode 0 shrinking as illustrated in Panel (a). When δ_E increases, the opposite applies. Panel (b) illustrates what happens to the stopping set if a lower cost per infective is combined with a jump at $I_t = 150$ similar to (13).

References

Abdin, A. F., Fang, Y., Caunhye, A., Alem, D., Barros, A., & Zio, E. (2021). An optimization model for planning testing and control strategies to limit the spread of a pandemic - the case of COVID-19. *European Journal of Operational Research*, forthcoming

- Acemoglu, D., Chernozhukov, V., Werning, I., & Whinston, M. D. (2020). *A multi-risk SIR model with optimally targeted lockdown*. National Bureau of Economic Research. Working paper
- Alvarez, F. E., Argente, D., & Lippi, F. (2020). *A simple planning problem for COVID-19 lockdown*. Working paper. National Bureau of Economic Research.
- Amaratunga, D., Cabrera, J., Ghosh, D., Katehakis, M. N., Wang, J., & Wang, W. (2021). Socio-economic impact on COVID-19 cases and deaths and its evolution in new jersey. *Annals of Operations Research*, forthcoming, 1–14.
- Anderson, R. M., Heesterbeek, H., Klinkenberg, D., & Hollingsworth, T. D. (2020). How will country-based mitigation measures influence the course of the COVID-19 epidemic? *The Lancet*, 395, 931–934.
- Atkeson, A. G. (2020). What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios. *Staff report*. Federal Reserve bank of Minneapolis.
- Avery, C., Bossert, W., Clark, A., Ellison, G., & Ellison, S. F. (2020). Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. *Working paper*. National Bureau of Economic Research.
- Bastos, S. B., & Cajueiro, D. O. (2020). Modeling and forecasting the early evolution of the COVID-19 pandemic in Brazil. *Scientific Reports*, 10, 19457.
- Bethune, Z., & Korinek, A. (2020). COVID-19 infection externalities: Trading off lives vs. livelihoods. NBER Working Paper No. w27009.
- Bliman, P.-A., Duprez, M., Privat, Y., & Vauchelet, N. (2021). Optimal immunity control and final size minimization by social distancing for the SIR epidemic model. *Journal of Optimization Theory and Applications*, 189, 408–436.
- Böckman, T., Fleten, S.-E., Juliussen, E., Langhammer, H. J., & Revdal, I. (2008). Investment timing and optimal capacity choice for small hydropower projects. *European Journal of Operational Research*, 190, 255–267.
- Bozkir, C. D. C., Ozmemis, C., Kurbanzade, A. K., Balci, B., Gunes, E. D., & Tuglular, S. (2021). Capacity planning for effective cohorting of hemodialysis patients during the coronavirus pandemic: A case study. *European Journal of Operational Research*, forthcoming
- Caulkins, J. P., Grass, D., Feichtinger, G., Hartl, R. F., Kort, P. M., Prskawetz, A., et al., (2021). The optimal lockdown intensity for COVID-19. *Journal of Mathematical Economics*, 93, 102489.
- Chang, J. T., & Kaplan, E. H. (2021). Modeling local coronavirus outbreaks. *European Journal of Operational Research*, forthcoming
- Chen, K., Pun, C. S., & Wong, H. Y. (2021). Efficient social distancing during the COVID-19 pandemic: Integrating economic and public health considerations. *European Journal of Operational Research*, forthcoming
- Choi, T.-M. (2021). Fighting against COVID-19: What operations research can help and the sense-and-respond framework. *Annals of Operations Research*, forthcoming
- Cowling, B. J., Chan, K. H., Fang, V. J., Cheng, C. K., Fung, R. O., Wai, W., et al., (2009). Facemasks and hand hygiene to prevent influenza transmission in households: A cluster randomized trial. *Annals of Internal Medicine*, 151, 437–446.
- d'Halluin, Y., Forsyth, P., & Vetzal, K. (2007). Wireless network capacity management: A real options approach. *European Journal of Operational Research*, 176(1), 584–609.
- Devalkar, S. K., Anupindi, R., & Sinha, A. (2011). Integrated optimization of procurement, processing, and trade of commodities. *Operations Research*, 59(6), 1369–1381.
- Di Domenico, L., Pullano, G., Sabbatini, C. E., Boëlle, P.-Y., & Colizza, V. (2020). Impact of lockdown on COVID-19 epidemic in Île-de-France and possible exit strategies. *BMC Medicine*, 18, 240.
- Djidjou-Demasse, R., Michalakos, Y., Choisy, M., Sofonea, M. T., & Alizon, S. (2020). Optimal COVID-19 epidemic control until vaccine deployment. *Working paper, medRxiv*. Cold Spring Harbor Laboratory Press.
- Doyle, S. (2021). Predicting the spread of COVID-19 in Ireland using an age-cohort SEIRD model. *SIAM Undergraduate Research Online*. Published online, 16 April 2021.
- Drakopoulos, K., & Randhawa, R. S. (2021). Why perfect tests may not be worth waiting for: Information as a commodity. *Management Science*, 67, 6678–6693.
- Ernstsen, R. R., & Boomsma, T. K. (2018). Valuation of power plants. *European Journal of Operational Research*, 266(3), 1153–1174.
- Farboodi, M., Jarosch, G., & Shimer, R. (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, 196, 105293.
- Federico, S., & Ferrari, G. (2020). Taming the spread of an epidemic by lockdown policies. *Journal of Mathematical Economics*, 102453. <https://doi.org/10.1016/j.jmateco.2020.102453>.
- Felix, B. J., & Weber, C. (2012). Gas storage valuation applying numerically constructed recombining trees. *European Journal of Operational Research*, 216(1), 178–187.
- Feng, X., Rong, Y., Shen, Z. M., & Snyder, L. V. (2020). Pricing during disruptions: Order variability versus profit. *Decision Sciences*, forthcoming
- Ferguson, N. M., Cummings, D. A. T., Cauchemez, S., Fraser, C., Riley, S., Meeyai, A., et al., (2005). Strategies for containing an emerging influenza pandemic in southeast Asia. *Nature*, 437, 209–214.
- Firoozi, F., & Merrifield, J. (2003). An optimal timing model of water reallocation and reservoir construction. *European Journal of Operational Research*, 145(1), 165–174.
- Fischetti, M., Fischetti, M., & Stoustrup, J. (2021). Safe distancing in the time of COVID-19. *European Journal of Operational Research*, forthcoming
- Garriga, C., Manuelli, R., & Sanghi, S. (2020). Optimal management of an epidemic: Lockdown, vaccine and value of life. *Working paper 2020-031*. Human Capital and Economic Opportunity Global Working Group.
- Ghaderi, M. (2022). Public health interventions in the face of pandemics: network structure, social distancing, and heterogeneity. *European Journal of Operational Research*, 298(3), 1016–1031.
- Hatchett, R. J., Mecher, C. E., & Lipsitch, M. (2007). Public health interventions and epidemic intensity during the 1918 influenza pandemic. *Proceedings of the National Academy of Sciences of the United States of America*, 104, 7582–7587.
- He, X., Lau, E. H. Y., Wu, P., Deng, X., Wang, J., Hao, X., et al., (2020). Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine*, 26, 672–675.
- Huberts, N. F. D., & Thijssen, J. J. J. (2021). Investment timing in the case of a Birth-Death process with an application to markets with consumer externalities. Working paper.
- Jia, L., & Chen, W. (2021). Uncertain SEIAR model for COVID-19 cases in China. *Fuzzy Optimization and Decision Making*, 20, 243–259.
- Kallio, M., Kuula, M., & Oinonen, S. (2012). Real options valuation of forest plantation investments in Brazil. *European Journal of Operational Research*, 217(2), 428–438.
- Kamigaki, T., & Oshitani, H. (2009). Epidemiological characteristics and low case fatality rate of pandemic (H1N1) 2009 in Japan. *PLoS Currents*, 1, RRN1139.
- Kantner, M., & Koprucki, T. (2020). Beyond just “flattening the curve”: Optimal control of epidemics with purely non-pharmaceutical interventions. *Journal of Mathematics in Industry*, 10, 23.
- Keogh-Brown, M., Wren-Lewis, S., Edmunds, W., Beutels, P., & Smith, R. (2010). The possible macroeconomic impact on the UK of an influenza pandemic. *Health Economics*, 19, 1345–1360.
- Kruse, T., & Strack, P. (2020). Optimal control of an epidemic through social distancing. Cowles Foundation Discussion Paper No. 2229.
- Kumar, A., Choi, T.-M., Wamba, S. F., Gupta, S., & Tan, K. H. (2021a). Infection vulnerability stratification risk modelling of COVID-19 data: A deterministic SEIR epidemic model analysis. *Annals of Operations Research*, forthcoming
- Kumar, S., Xu, C., Ghildayal, N., Chandra, C., & Yang, M. (2021b). Social media effectiveness as a humanitarian response to mitigate influenza epidemic and COVID-19 pandemic. *Annals of Operations Research*, 1–29.
- Leung, G. M., & Nicoll, A. (2010). Reflections on pandemic (H1N1) 2009 and the international response. *PLoS Medicine*, 7, journal.pmed.1000346
- Liu, Q., Shen, Q., Li, Z., & Chen, S. (2021). Stimulating consumption at low budget: Evidence from a large-scale policy experiment amid the COVID-19 pandemic. *Management Science*, 67(12), 7291–7307.
- Liu, Y., Gayle, A. A., Wilder-Smith, A., & Rocklöv, J. (2020). The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine*, 27(2).
- Manski, C. F., & Molinari, F. (2021). Estimating the COVID-19 infection rate: Anatomy of an inference problem. *Journal of Econometrics*, 220(1), 181–192.
- Muñoz, J. L., Contreras, J., Caamaño, J., & Correia, P. F. (2011). A decision-making tool for project investments based on real options: The case of wind power generation. *Annals of Operations Research*, 186(1), 465.
- Mutesa, L., Ndishimye, P., Butera, Y., et al., (2021). A pooled testing strategy for identifying SARS-CoV-2 at low prevalence. *Nature*, 589, 276–280. Accelerated Article Preview
- Nagurny, A. (2021). Supply chain game theory network modeling under labor constraints: Applications to the COVID-19 pandemic. *European Journal of Operational Research*, 293, 880–891.
- Nikolopoulos, K., Punia, S., Schäfers, A., Tsinopoulos, C., & Vasilakis, C. (2021). Forecasting and planning during a pandemic: COVID-19 growth rates, supply chain disruptions, and governmental decisions. *European Journal of Operational Research*, 290(1), 99–115.
- Nikzad, E., Bashiri, M., & Abbasi, B. (2021). A matheuristic algorithm for stochastic home health care planning. *European Journal of Operational Research*, 288(3), 753–774.
- Patterson-Lomba, O. (2020). Optimal timing for social distancing during an epidemic. *Working paper, medRxiv*. Cold Spring Harbor Laboratory Press.
- Pendharkar, P. C. (2010). Valuing interdependent multi-stage investments: A real options approach. *European Journal of Operational Research*, 201(3), 847–859.
- Pifarré i Arolas, H., Acosta, E., López-Casasnovas, G., Lo, A., Nicodemo, C., Riffe, T., et al., (2021). Years of life lost to COVID-19 in 81 countries. *Scientific Reports*, 11, article 3504
- Piguillem, F., & Shi, L. (2020). Optimal COVID-19 quarantine and testing policies. EIEF Working papers series 2004, Einaudi Institute for Economics and Finance (EIEF), Revised Apr 2020.
- Rezapour, S., Baghaian, A., Naderi, N., & Sarmiento, J. P. (2021). Infection transmission and prevention in metropolises with heterogeneous and dynamic populations. *European Journal of Operational Research*, forthcoming
- Roques, L., Klein, E. K., Papaix, J., Sar, A., & Soubeyrand, S. (2020). Impact of lockdown on the epidemic dynamics of COVID-19 in France. *Frontiers in Medicine*, 7, 274.
- Ru, H., Yang, E., & Zou, K. (2021). Combating the COVID-19 pandemic: The role of the SARS imprint. *Management Science*, 67, 5606–5615.
- Salje, H., Kiem, C. T., Lefrançois, N., Courtejoie, N., Bosetti, P., Paireau, J., et al., (2020). Estimating the burden of SARS-CoV-2 in France. *Science*, 369(6500), 208–211.
- Shin, I. (2020). An optimal policy for social resources allocation: When outbreak of infectious diseases. MPRA Working Paper No. 99936.
- Silal, S. P. (2021). Operational research: A multidisciplinary approach for the management of infectious disease in a global context. *European Journal of Operational Research*, 291(3), 929–934.

- Sinha, P., Kumar, S., & Chandra, C. (2021). Strategies for ensuring required service level for COVID-19 herd immunity in Indian vaccine supply chain. *European Journal of Operational Research*. forthcoming
- Smith, R., Keogh-Brown, M., Barnett, T., & Tait, J. (2009). The economy-wide impact of pandemic influenza and the UK: A computable general equilibrium modelling experiment. *British Medical Journal*, 339, b4571.
- Støre, K., Fleten, S.-E., Hagspiel, V., & Nunes, C. (2018). Switching from oil to gas production in a depleting field. *European Journal of Operational Research*, 271(2), 710–719.
- Steffen, B., & Weber, C. (2016). Optimal operation of pumped-hydro storage plants with continuous time-varying power prices. *European Journal of Operational Research*, 252(1), 308–321.
- Stern, A., & Markel, H. (2009). What Mexico taught the world about pandemic influenza preparedness and community mitigation strategies. *Journal of the American Medical Association*, 302, 1221–1222.
- Stokey, N., & Lucas, R. (1989). *Recursive methods in economic dynamics*. Harvard University Press.
- Talic, S., Shah, S., Wild, H., Gasevic, D., Maharaj, A., Ademi, Z., et al., (2021). Effectiveness of public health measures in reducing the incidence of COVID-19, SARS-CoV-2 transmission, and COVID-19 mortality: systematic review and meta-analysis. *British Medical Journal*, 375:e068302.
- Taylor, J. W., & Taylor, K. S. (2021). Combining probabilistic forecasts of COVID-19 mortality in the United States. *European Journal of Operational Research*. forthcoming
- Thijssen, J. (2022). Optimal investment and abandonment decisions for projects with construction uncertainty. *European Journal of Operational Research*, 298, 368–379.
- Thijssen, J., & Bregantini, D. (2017). Costly sequential experimentation and project valuation with an application to health technology assessment. *Journal of Economic Dynamics and Control*, 77, 202–229.
- Thompson, M., Davison, M., & Rasmussen, H. (2004). Valuation and optimal operation of electric power plants in competitive markets. *Operations Research*, 52(4), 546–562.
- Thul, L., & Powell, W. (2021). Stochastic optimization for vaccine and testing kit allocation for the COVID-19 pandemic. *European Journal of Operational Research*. forthcoming
- Tippong, D., Petrovic, S., & Akbari, V. (2021). A review of applications of operational research in healthcare coordination in disaster management. *European Journal of Operational Research*. forthcoming
- Toda, A. A. (2020). Susceptible-infected-recovered (SIR) dynamics of COVID-19 and economic impact. *CEPR Press, Issue on Covid Economics*, 1, 43–63.
- Toxvaerd, F. (2020). Equilibrium social distancing. *CEPR Press, Covid Economics: Vetted and Real-Time Papers*, 15, 110–133.
- Trigeorgis, L., & Tsekrekos, A. E. (2018). Real options in operations research: A review. *European Journal of Operational Research*, 270(1), 1–24.
- Truong, C., & Trück, S. (2016). It's not now or never: Implications of investment timing and risk aversion on climate adaptation to extreme events. *European Journal of Operational Research*, 253(3), 856–868.
- Tsekrekos, A. E., & Yannacopoulos, A. N. (2016). Optimal switching decisions under stochastic volatility with fast mean reversion. *European Journal of Operational Research*, 251(1), 148–157.
- Tseng, C.-L., & Barz, G. (2002). Short-term generation asset valuation: A real options approach. *Operations Research*, 50(2), 297–310.
- de Vaan, M., Mumtaz, S., Nagaraj, A., & Srivastava, S. B. (2021). Social learning in the COVID-19 pandemic: Community establishments' closure decisions follow those of nearby chain establishments. *Management Science*, 67, 4446–4454.
- de Véricourt, F., Gurkan, H., & Wang, S. (2021). Informing the public about a pandemic. *Management Science*, 67, 6350–6357.
- Wang, H., Wang, Z., Dong, Y., Chang, R., Xu, C., Yu, X., et al., (2020). Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China. *Cell Discovery*, 6(1), 10.
- Wang, T., Liu, B., Zhang, J., & Li, G. (2019). A real options-based decision-making model for infrastructure investment to prevent rainstorm disasters. *Production and Operations Management*, 28(11), 2699–2715.
- Wolfinger, D., Gansterer, M., Doerner, K. F., & Popper, N. (2021). A large neighbourhood search metaheuristic for the contagious disease testing problem. *European Journal of Operational Research*. forthcoming
- Wu, J. T., Cowling, B. J., Lau, E. H., Ip, D. K., Ho, L. M., Tsang, T., et al., (2010). School closure and mitigation of pandemic (H1N1) 2009, Hong Kong. *Emerging Infectious Diseases*, 16, 538–541.
- You, C., Deng, Y., Hu, W., Sun, J., Lin, Q., Zhou, F., et al., (2020). Estimation of the time-varying reproduction number of COVID-19 outbreak in China. *International Journal of Hygiene and Environmental Health*, 228, 113555.
- Zhang, W., Shi, X., Huang, A., Hua, G., & Teunter, R. H. (2021). Optimal stock and capital reserve policies for emergency medical supplies against epidemic outbreaks. *European Journal of Operational Research*. forthcoming
- Zhang, Z., & Enns, E. (2020). Optimal timing and effectiveness of COVID-19 outbreak responses in China: A modelling study. SSRN Working paper.
- Zhao, S., Lin, Q., Ran, J., Musa, S. S., Yang, G., Wang, W., et al., (2020). Preliminary estimation of the basic reproduction number of novel coronavirus (2019-nCoV) in China, from 2019 to 2020: A data-driven analysis in the early phase of the outbreak. *International Journal of Infectious Diseases*, 92, 214–217.
- Zhou, F., Yu, T., Du, R., Fan, G., Liu, Y., Liu, Z., et al., (2020). Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: A retrospective cohort study. *Lancet*, 395, 1054–1062.