Received: 19 April 2022 | Revised: 31 March 2023 | Accepted: 6 April 2023

DOI: 10.1002/nav.22114

RESEARCH ARTICLE

WILEY

Optimal interventions of infectious disease

Xu Sun[1] | **Yunan Liu**[2]

[1]Industrial and Systems Engineering, University of Florida, Gainesville, Florida, USA
[2]Industrial and Systems Engineering, North Carolina State University, Raleigh, North Carolina, USA

Correspondence
Yunan Liu, Industrial and Systems Engineering, North Carolina State University, Raleigh, NC, USA.
Email: yliu48@ncsu.edu

Abstract

The recent outbreak of novel coronavirus has highlighted the need for a benefit-cost framework to guide unconventional public health interventions aimed at reducing close contact between infected and susceptible individuals. In this paper, we propose an optimal control problem for an infectious disease model, wherein the social planner can control the transmission rate by implementing or lifting lockdown measures. The objective is to minimize total costs, which comprise infection costs, as well as fixed and variable costs associated with lockdown measures. We establish conditions concerning model primitives that guarantee the existence of a straightforward optimal policy. The policy specifies two switching points (a, b), whereby the social planner institutes a lockdown when the percentage of infected individuals exceeds b, and reopens the economy when the percentage of infected individuals drops below a. We subsequently extend the model to cases where the social planner may implement multiple lockdown levels. Finally, numerical studies are conducted to gain additional insights into the value of these controls.

KEYWORDS

benefit-cost analysis, infectious diseases, non-pharmaceutical interventions, optimal control

1 | INTRODUCTION

The COVID-19 pandemic has led to the implementation of various non-pharmaceutical interventions, commonly referred to as lockdowns, such as stay-at-home orders, curfews, quarantines, and other societal restrictions, in numerous countries and territories worldwide. The primary purpose of these measures is to control the spread of the coronavirus. As of April 2020, about half of the world's population was under some form of lockdown, with more than 90 countries or territories mandating or requesting over 3.9 billion people to stay at home. The extent and severity of these enforced lockdowns differ from one country or territory to another. Several Asian countries have effectively contained COVID-19 pandemics by combining large-scale testing, contact tracing, isolation, and quarantine with moderate (e.g., South Korea) or strict (e.g., China) social distancing measures. In recognition of the potential for a sudden surge in COVID-19 cases to overwhelm their healthcare systems, many European countries as well as the United States adopted aggressive social distancing measures to combat the pandemic.

Although lockdown restrictions were generally supported by public health experts and economists, there were concerns about the potential health, social, and economic consequences of such measures. The economic downturn induced by a lockdown can result in health problems such as "deaths of despair" and create strains on public-health budgets, which may lead to more non-COVID-19-related deaths than a lockdown would save from the pandemic. Additionally, severe social confinement can lead to social tensions and have a profoundly negative impact on people. Therefore, a well-designed lockdown entry-and-exit strategy is critical for countries that choose to implement a lockdown policy. The ongoing debate over the necessity of lockdown measures highlights the need for a decision-making framework for governments and public health agencies to make well-informed choices regarding unconventional interventions against infectious diseases. In this paper, we aim to provide preliminary answers to the following policy questions: (i) Is a lockdown necessary? (ii) If a lockdown is deemed economically beneficial, when is the optimal time to implement and remove it?

In order to gain analytical insights into the effects of socioeconomic factors on government intervention strategies, we introduce a stochastic control formulation that consists of a stylized disease-spread model and a specific type of governmental intervention. The disease-spread model incorporates three realistic aspects of infectious disease transmission: (i) the disease can be transmitted from an infected individual to a virus-free individual through random contact, (ii) infected individuals may recover from the disease after a certain period of time, and (iii) environmental factors beyond contact-based infection and self-healing can contribute to disease spread, making the trajectory of disease spread partially unpredictable. The governmental intervention under consideration is binary in structure: the social planner can choose when to lock down and open the economy at any given time. When a lockdown is imposed, contact is reduced, and the transmission rate decreases to a lower level. However, a lockdown also incurs socioeconomic costs. At a high level, the problem can be viewed as an optimal switching problem, where the objective is to balance the burden imposed by the disease with the economic cost of lockdowns in an optimal manner.

While there is a substantial amount of literature on infectious disease control, most of it is based on deterministic disease-spread models, which ignore the inherent randomness of disease transmission and focus mainly on pharmaceutical interventions like vaccination. In contrast, our paper focuses on unconventional government interventions and incorporates randomness into the disease-spread process. In this regard, our paper contributes to the literature on epidemic control. We do acknowledge that the disease spread process is complex, and more accurate infectious disease models may require the inclusion of several additional compartments, as seen in El Housni et al. (2022a). Nevertheless, our aim is to derive actionable insights using a relatively simple disease model.

Our key findings can be summarized as follows: First, a lockdown policy should not be implemented when its economic cost exceeds the long-term cost of allowing the disease to spread uncontrollably. Second, when a lockdown is economically feasible, a social planner should adopt a threshold-based strategy, where a lockdown is imposed when the number of infected individuals exceeds a certain threshold b, and the economy is reopened when the number of infections falls below another threshold $a < b$. Although this strategy seems intuitive, determining the optimal values of a and b requires a formal analysis, which we conduct in this paper. Third, we extend our framework to allow for multiple levels of lockdown, and under certain conditions, we demonstrate that a multi-level sequential switching policy is optimal.

The rest of the paper is structured as follows: Section 2 reviews the relevant literature. Section 2 introduces our disease-spread model and presents the stochastic control problem. Section 2 characterizes the optimal intervention strategy via the solution to the corresponding dynamic programming equation; in the same section, we conduct numerical explorations to gain additional insights into the structure of the optimal policy. Section 2 discusses a useful extension of the base model. Section 2 concludes by highlighting some limitations and suggesting future research directions.

2 | LITERATURE REVIEW

The paper contributes to the literature on stochastic models of infectious diseases, which has been previously explored in works such as (Cai et al., 2017; Gray et al., 2011; Mao et al., 2002; Tuckwell & Williams, 2007). However, this previous research primarily focused on studying the persistence and extinction of diseases. In contrast, our work focuses on developing control mechanisms for infectious diseases, specifically addressing the timing of government intervention.

The present work is related to several studies that have explored the optimal control of infectious diseases using deterministic infectious disease models. For example, Sethi and Staats (1978) used optimal control to minimize the cost of disease and medical treatment. Behncke (2000) studied a general *susceptible-infected-recovered* (SIR) model with control by vaccination, quarantine, or health campaigns and showed that it is optimal to exert maximum effort in all cases. Chehrazi et al. (2019) investigated an optimal control problem for a *susceptible-infected-susceptible* (SIS) model of an infectious disease with resistance and demonstrated that the optimal prescription policy is of the bang-bang type with a single switching time. Chen and Kong (2022) studied the effects of different hospital admission policies on the spread of infectious diseases using a modified *susceptible-exposed-infected-recovered* (SEIR) model under a static social distancing policy. Additionally, (El Housni et al., 2022a; El Housni et al., 2022b) developed new compartmentalized models to study the effects of testing capacity and social distancing measures on the pandemic in New York City and North Carolina. Finally, Chen et al. (2022) developed a deterministic SIR model with a finite number of test kits and studied the impact of test capacity management on the disease transmission process. In contrast to these studies, our paper employs a stochastic modeling framework, which adds an extra degree of realism.

Our literature search has revealed a limited number of papers on the control of stochastic epidemic models. Lefévre (1981) derived the optimal control policy for a birth-and-death epidemic model, assuming that both birth and death rates (representing quarantine and medical treatment, respectively) are subject to control. In addition, the author examined the effect of model parameters on the optimal strategy. Yaesoubi and Cohen (2011) focused on vaccination and transmission reduction measures, formulating a

Markov decision process to derive optimal vaccination and transmission-reducing interventions. The recent work by Lee et al. (2022) conducted comprehensive agent-based simulations to study how the COVID-19 pandemic responds to several non-pharmaceutical interventions. Our work is most closely related to Bai et al. (2022), where the authors studied an optimal public health intervention problem with the goal of balancing the cost of infection and socioeconomic losses. Bai et al. (2022) modeled disease dynamics by a deterministic SIR process, and a major consideration of their work is the response from strategic individuals, each of whom aims to maximize their own utility. Unlike these papers that considered rate control driven by vaccination and/or treatment, we formulate an optimal switching problem, motivated by the mass implementation of non-pharmaceutical interventions across the globe during the COVID-19 crisis. Furthermore, our work is differentiated because we consider a Brownian model, which enables us to derive an explicit solution and clear-cut insights.

There exists a stream of literature focused on studying stochastic control problems that involve sequential switching decisions. These types of problems have fixed switching costs, which are investments necessary to realize the benefits of an appropriate regime. This forces the controller to look beyond immediate advantages to ensure that a regime switch will provide sufficient long-term benefits to justify the fixed investment. The problems are formulated using either a discounted cost criterion (Duckworth & Zervos, 2001; Ly Vath & Pham, 2007; Zervos et al., 2013) or an average cost criterion (Wu & Chao, 2014). These authors' contributions focus on characterizing solutions to the corresponding Bellman equation, assuming that the underlying state process satisfies the standard Lipschitz continuity condition (i.e., the drift and volatility are Lipschitz continuous functions of the state process). However, our state process does not satisfy this Lipschitz continuity assumption due to the super-linear growth of disease spread. Unlike most papers in this line of work, our dynamic programming equation features a non-Lipschitz singularity at the boundary, presenting non-trivial technical challenges for our analysis. Thus, our paper contributes to the literature on optimal switching by developing a set of tools that we believe are methodologically novel.

In summary, our paper appears to be the first to analyze the impact of non-pharmaceutical interventions on a pandemic within a formal stochastic control framework, while accounting for the inherently random nature of epidemic growth and spread.

3 | MODEL

Notation and convention We assume that all random variables are defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The expectation with respect to $\mathbb{P}$ is denoted by $\mathbb{E}$. For a sufficiently smooth function f, we use f' and f'' to denote its first and second derivatives, respectively. The partial derivative of a multivariate function $f(x, y, \dots)$ with respect to the variable x is denoted by f_x.

In this paper, we consider a social planner who must make binary decisions regarding interventions for an infectious disease that affects a population. The two possible decisions are "lock down the economy" and "open the economy." To formulate the social planner's problem, we first establish a mathematical framework for disease spread in the absence of any intervention in Section 2. Our base formulation extends the classical SIS epidemic model from a deterministic framework to a stochastic one, where the number of infectious individuals follows a *stochastic differential equation* (SDE). In Section 2, we introduce the intervention mechanism, present relevant socioeconomic costs, and state the social planner's problem.

3.1 | Disease dynamics

In a closed population of fixed size, there are two categories of individuals: susceptible individuals and infected individuals. We simplify the assumption that all infected individuals exhibit symptoms and will report their status change (from susceptible to infected) upon contracting the disease.

The system state at any given time, denoted by X_t, represents the percentage of infected people, where $t \geq 0$. Thus, for all $t \geq 0$, X_t falls within the range of $[0, 1]$, and the percentage of the susceptible population at time t is given as $1 - X_t$. An infected individual transmits the disease to a susceptible one through contact, with a transmission rate determined by the host-population density and the infectiousness of the pathogens. Furthermore, all infected individuals can recover at a rate γ. Recovery means that the virus is eliminated from the body, and the parameter γ reflects the strength of the host's immune system. We assume that all individuals who recover return to the susceptible state, which effectively limits our analysis to infections with low mortality rates and no conferred immunity. This assumption is appropriate for viruses that can mutate continuously, allowing the virus to evade an individual's immune system. In the absence of any intervention, the dynamics can be described by the following stochastic differential equation (SDE):

$$\mathrm{d}X_t = (\beta(1 - X_t) - \gamma)X_t\mathrm{d}t + \sigma\sqrt{X_t(1 - X_t)}\mathrm{d}B_t, \qquad (1)$$

where the first drift term, $\beta X_t(1 - X_t)$, captures the disease spread dynamics, the second drift term, γX_t, models the natural recovery dynamics, and the third term, B_t, denotes a standard Brownian motion. Here, β can be interpreted as the rate of contact resulting from all activities, including both *essential activities* (e.g., grocery shopping and doctor visits) and *nonessential social interactions* (e.g., social gatherings and entertainment events), without any lockdown measures. The parameter γ represents the recovery rate of an infected individual, while σ is the volatility parameter that characterizes the magnitude of the stochastic fluctuations of the system.

To better motivate the dynamic model given by Equation (1), let us consider a continuous-time Markov chain model for disease spread. We have a closed population of size N that can be divided into two compartments, namely the susceptible and the infected. Let S_t and I_t represent the number of susceptible and infected individuals, respectively, such that $S_t + I_t = N$. A susceptible person becomes infected at the rate $\beta I_t/N$, indicating that a new infection occurs at a rate of $\beta S_t I_t/N$. On the other hand, an infected person recovers and becomes susceptible again at a rate of γ. We can observe that $I := I_t; t \geq 0$ is a birth-and-death process. Using the general framework presented by Kurtz (1971), we can approximate the evolution of the fraction of infected individuals in the population (i.e., I/N) by a diffusion process $\bar{I}$ with the following dynamics:

$$d\bar{I}_t = (\beta(1-\bar{I}_t)-\gamma)\bar{I}_t dt + \frac{1}{\sqrt{N}}\sqrt{\beta\bar{I}_t(1-\bar{I}_t)+\gamma\bar{I}_t}dB_t. \quad (2)$$

To obtain (1) from (2), we simply need to remove the term $\gamma\bar{I}_t$ under the square root in (2), set $\sqrt{\beta/N}$ as the volatility parameter, and replace $\bar{I}$ with X. There are two reasons why we drop the term $\gamma\bar{I}_t$. Firstly, the reproduction number, which is defined as β/γ, is high for COVID-19, making the term $\beta\bar{I}_t(1-\bar{I}_t)$ under the square root in (2) more dominant than $\gamma\bar{I}_t$. Secondly, the assumption that recovery times are exponentially distributed causes the term $\gamma\bar{I}_t$ to arise in the model in (2). However, recovery times from COVID-19 are less variable than an exponential distribution suggests. Thus, removing the term $\gamma\bar{I}_t$ from the model in (2) helps to correct to some extent the overestimation of the system's variability.

To summarize, with σ being $\sqrt{\beta/N}$, the process (1) is considered a diffusion approximation for the corresponding Markov chain model. The key distinction is that a Markov chain has a discrete state space, whereas a diffusion process has a continuous state space. However, we anticipate that the actual volatility will be greater than what $\sqrt{\beta/N}$ suggests in practice due to the contribution of other random factors, such as environmental factors, to the system's variability.

Remark 1 (Stochastic model vs. deterministic model). The disease model represented by (1) belongs to the class of Wright-Fisher processes (see, e.g., equation (1) in Jenkins and Spano (2017)). Moreover, by verifying the conditions of Theorem 5 in Feller (1954), it is evident that zero is an accessible and absorbing state. This implies that the process will hit the zero boundary (without intervention) in a finite (yet random) time and remain at zero thereafter. The stochastic model considered in this paper offers an extra degree of realism by capturing the stochastic nature of disease transmission and avoids the need to assume an artificial threshold, as was often done in deterministic models to capture eradication or extinction (see, e.g., Tebbens and Thompson (2009)). In fact, one can obtain the deterministic counterpart of (1) by removing the Brownian term, which yields

$$dx_t = (\beta(1-x_t)-\gamma)x_t dt. \quad (3)$$

It is worth noting that the solution to (3) cannot hit zero given $x_0 > 0$, indicating that a disease governed by (3) cannot go extinct. In contrast, our stochastic model in (1) yields a control policy that depends on X_t, enabling the decision maker to take action at t in response to the real-time state of the system (i.e., the number of infected individuals), which is not completely predictable before t. The solution to the deterministic problem in (3) ignores the random fluctuations in the system's future dynamics because it is a deterministic function of time t, rather than the system's state.

The disease model given in (1) makes an implicit assumption that the incubation period of the disease is zero. Additionally, the assumption of a constant rate parameter β ignores potential changes in behavior in response to the epidemic's growth, which is a well-known phenomenon. It should be noted that these simplifications are commonly made in order to achieve greater mathematical tractability. However, it is important to keep in mind that using static parameters can result in an overestimation of disease incidence or prevalence.

3.2 | Intervention mechanism and the social planner's problem

We consider the scenario where a social planner has the option to implement a lockdown to reduce the transmission rate of the disease. During a lockdown, the disease is transmitted at a potentially much slower rate $\tilde{\beta}$ compared to the baseline transmission rate β. Note that β accounts for the transmission rate due to both essential and nonessential activities, while $\tilde{\beta}$ only captures the transmission rate due to essential activities required for daily living, such as grocery shopping and medical visits. It is important to note that a reduced transmission rate can potentially result in a different volatility parameter. However, to keep the analysis simple, we assume that a lockdown only affects the transmission rate and not the volatility parameter.

To perform a formal cost-benefit analysis, let $\ell > 0$ be the marginal cost of infection, so the cost of infection accumulates at a rate of ℓX_t. This cost may include expenses related to symptom relief, hospitalizations, long-term infection complications, and lost productivity. In addition to the cost of infections, we consider the costs of implementing a lockdown. We assume that each time the social planner enacts a lockdown, there is a fixed cost of K, and that the lockdown

incurs economic tolls at a rate of κ for the duration of its implementation.

Given the possibility of the epidemic dying out without intervention, it is important to consider the cost-effectiveness of implementing a lockdown. The system can be in one of two modes at any given point in time: "on" or "off." The on mode refers to the implementation of a lockdown policy, while the off mode refers to the absence of a policy. The decision to switch between modes is made by the social planner and constitutes an intervention strategy. To model this strategy, we use an adapted, finite variation, càdlàg process Y with values in the set $0, 1$. Specifically, $Y_t = 0$ represents the off mode, and $Y_t = 1$ represents the on mode. Switching from the off to the on mode incurs a fixed cost of K, whereas switching from the on to the off mode is free. The objective of the social planner is to minimize the total cost, which includes both the cost of infections, accumulated at a rate of ℓX_t, the cost of implementing the lockdown policy, which incurs economic tolls at a rate of κ for the duration of its implementation, and the cost due to switching. More formally, the social planner aims to identify some Y to

$$\text{minimize} \quad J(x,y) := \mathbb{E}\left[\int_0^\tau \ell X_u \mathrm{d}u + \kappa \int_0^\tau Y_u \mathrm{d}u + K\sum_{u\le\tau}[\Delta Y_u]^+ \middle| X_0 = x, Y_0 = y\right], \quad (4)$$

where the stopping time τ denotes the time at which X_t hits zero for the first time, and X_t is the solution to the SDE

$$\mathrm{d}X_t = [b(Y_t)(1 - X_t) - \gamma]X_t\mathrm{d}t + \sigma\sqrt{X_t(1 - X_t)}\mathrm{d}B_t \quad (5)$$

for $b(y) := \beta 1_{\{y=0\}} + \tilde{\beta} 1_{\{y=1\}}$. See Table 1 for a summary of all notation.

To find the optimal strategy, we need to consider all possible policies that can be employed by the social planner. While the optimal strategy can be an arbitrary functional of the sample paths of X, the strong Markov property of the process X allows us to focus on a more restricted class of policies, namely Markov policies (A Markov policy is a function that maps the current state of the system to an action, and the choice of action depends only on the current state and not on the past history of the system.).

TABLE 1 Glossary of model parameters and notation.

Symbol	Definition
X_t	Fraction of infected population at t
Y_t	State of intervention at t
γ	Individual recovery rate
β	Rate of infection due to all activities (under no intervention)
$\tilde{\beta}$	Rate of infection due to essential activities (under lockdown)
σ	Volatility parameter
ℓ	Marginal cost of infection
κ	Rate of economic tolls under lockdown
K	Fixed cost per lockdown

The focus of this work is on the technical treatment of optimal control in an infectious disease model. However, it is important to recognize that real-world data can be utilized to calibrate key parameters of theoretical models. The World Health Organization (WHO)[1] is one of the primary sources of data on COVID-19 cases, deaths, and recoveries worldwide. Other sources of data include national and regional health departments, academic institutions, and research organizations. Increasingly, COVID-19 dashboards have gained popularity, such as the John Hopkins University COVID-19 Dashboard[2], the European Centre for Disease Prevention and Control (ECDC) COVID-19 Dashboard[3], and the Worldometer COVID-19 Dashboard[4]. These dashboards provide real-time information on cases, deaths, and recoveries in a specific region, presented in maps, graphs, and charts. Policymakers and the public can use these dashboards to understand the spread of the disease and its impact. Therefore, integrating real-world data into infectious disease models can significantly improve the accuracy of the model and the effectiveness of the control measures implemented.

Multiple recent studies have demonstrated how to incorporate dashboard data to calibrate models that predict disease trajectory and verify the accuracy of the model's predictions. For instance, Tatapudi et al. (2020) presented an agent-based simulation model for COVID-19 as a policy evaluation tool for public health decision-makers, which was calibrated using data from the Florida COVID-19 dashboard for Miami-Dade County. Similarly, Tatapudi and Das (2021) used an agent-based simulation model to investigate the impact of partial/full reopening of school/college campuses on pandemic spread, calibrated based on public information from the same dashboard. Hinch et al. (2021) developed an agent-based simulation of the epidemic, incorporating detailed age-stratification and realistic social networks, calibrated using observed data from the UK Government's COVID19 dashboard. Additionally, Bai et al. (2022) provided a pathway for efficiently estimating the value of β and γ in their model. However, estimating the disease burden parameter ℓ, as well as κ and K that characterize socio-economic losses due to lockdown is far less straightforward and requires ongoing research. Nonetheless, recent developments have been made in this area; see, for example, Zhang (2022) and reference therein.

4 | ANALYSIS

This sections is devoted to solving the control problem posited in the previous section. In Section 2, we describe the dynamic programming equation that will be used to characterize the

[1] https://covid19.who.int/.
[2] https://coronavirus.jhu.edu/map.html.
[3] https://www.ecdc.europa.eu/en/data/dashboards.
[4] https://www.worldometers.info/coronavirus/.

optimal strategy. In Section 2, we characterize the optimal intervention strategy via the solution to the dynamic programming equation.

4.1 | Dynamic programming equation

Let $V(x, y)$ denote the value function associated with the problem (4). With reference to the general control theory, the function $V(x, y)$ ought to solve some dynamic programming equation. To derive this equation, we begin with setting up a Bellman equation in a discrete-time setting with time step Δt and then take $\Delta \to 0$.

Suppose that the economy is locked down at time 0, that is, $Y_0 = 1$. In light of the problem setup, the social planner's immediate decisions consist of choosing between two options. The first is to continue enforcing the lockdown policy for a short time, Δt, and then continue optimally. With regard to Bellman's principle of optimality, we must have

$$V(x,1) \le \mathbb{E}\left[\int_0^{\Delta t} \ell X_u \mathrm{d}u + \kappa \int_0^{\Delta t} Y_u \mathrm{d}u + V(X_{\Delta t}, 1)\right],$$

where the inequality is due to the fact that continuing locking down the economy for the next moment may not be optimal. Assuming that $V(x, 1)$ is sufficiently smooth, we can apply Itô's lemma to the second term, divide both sides by Δt, and send $\Delta t \to 0$ to get

$$\frac{\sigma^2}{2}x(1-x)V_{xx}(x,1) + \left(\tilde{\beta}x - \tilde{\beta}x^2 - \gamma x\right)V_x(x,1) + \ell x + \kappa \ge 0.$$

The second option is to switch to the off mode and then continue optimally, resulting to

$$V(x,1) \le V(x,0).$$

Since these two are the social planner's only options, we can conclude that one of the two preceding inequalities must hold as an equality. Thus, the function $V(x, 1)$ must satisfy

$$\min\left\{\frac{\sigma^2}{2}x(1-x)V_{xx}(x,1) + \left(\tilde{\beta}x - \tilde{\beta}x^2 - \gamma x\right)V_x(x,1) + \ell x + \kappa,\ V(x,0) - V(x,1)\right\} = 0. \tag{6}$$

Arguing along similar lines, we can conclude that the function $V(x, 0)$ must satisfy

$$\min\left\{\frac{\sigma^2}{2}x(1-x)V_{xx}(x,0) + \left(\beta x - \beta x^2 - \gamma x\right)V_x(x,0) + \ell x,\ V(x,1) + K - V(x,0)\right\} = 0. \tag{7}$$

Combining (6) and (7) leads to the following *variational inequality* (VI) for $V(x, y)$:

$$\min\left\{\frac{\sigma^2}{2}x(1-x)V_{xx}(x,y) + \left(b(y)x - b(y)x^2 - \gamma x\right)V_x(x,y) + \ell x + \kappa y,\ V(x,1) + K - V(x,0),\ V(x,0) - V(x,1)\right\} = 0, \tag{8}$$

subject to the condition $V(0, y) = 0$ and the requirement that $\lim_{x\to 1} V_x(x, y)$ exists and is finite. In general, a VI formulation reveals structural properties of the optimal policy by reducing a dynamic decision problem to a point-wise optimization problem.

We now postulate the solution structure of the dynamic programming equation. Intuitively, if the costs associated with a lockdown are not prohibitively high, it would be beneficial for the social planner to implement a lockdown policy on a temporary basis. This intuition leads us to conjecture that the optimal control strategy is a sequential switching policy comprised of the following actions. If the system is currently operating in its off mode, then it is optimal to remain in that mode if X is below a threshold, say x_1^*, and switch to its on mode once X rises above x_1^*. On the other hand, if the system is currently operating in its on mode, then it is optimal to remain in that mode if X is above a certain level, say x_0^*, and switch to its off mode as soon as X drops below x_0^*. Clearly, this strategy is well-defined if $x_0^* < x_1^*$. Moreover, if this strategy, henceforth denoted as Y^*, is indeed optimal, we should be able to find $V_0(\cdot)$ and $V_1(\cdot)$ such that

$$\frac{\sigma^2}{2}x(1-x)V_0''(x) + \left(\beta x - \beta x^2 - \gamma x\right)V_0'(x) + \ell x = 0 \quad \text{for} \quad x \in [0, x_1^*) \quad \text{and} \tag{9}$$

$$\frac{\sigma^2}{2}x(1-x)V_1''(x) + \left(\tilde{\beta}x - \tilde{\beta}x^2 - \gamma x\right)V_1'(x) + \ell x + \kappa = 0 \quad \text{for} \quad x \in (x_0^*, 1] \tag{10}$$

subject to the boundary conditions:

$$V_0(x) = V_1(x) \quad \text{for} \quad x \in [0, x_0^*] \quad \text{and}$$
$$V_0(x) = V_1(x) + K \quad \text{for} \quad x \in [x_1^*, 1], \tag{11}$$

plus two optimality conditions derived from the "principle of smooth fit":

$$V_0'(x_0^*) = V_1'(x_0^*) \quad \text{and} \quad V_0'(x_1^*) = V_1'(x_1^*). \tag{12}$$

4.2 | Characterizing the optimal policy

Note that Equation (9) does not involve the unknown function V itself. Therefore, it is essentially a first-order differential equation. Indeed, if letting $U_0 := V_0'$, then U_0 solves

$$U_0'(x) + \frac{2}{\sigma^2}\left(\beta - \frac{\gamma}{1-x}\right)U_0(x) = -\frac{2}{\sigma^2}\frac{\ell}{1-x}. \tag{13}$$

To proceed, let

$$\phi(x,\iota) := e^{-2\beta x/\sigma^2}(1-x)^{-2\gamma/\sigma^2} \times \left(\iota - \frac{2\ell}{\sigma^2}\int_0^x e^{2\beta u/\sigma^2}(1-u)^{2\gamma/\sigma^2 - 1}\mathrm{d}u\right). \tag{14}$$

It is straightforward to verify that $\phi(\cdot, \iota)$ is a solution to (13) for each fixed ι, and $\phi(0, \iota) = \iota$.

Proposition 1. *There exists a unique $\bar{\iota}$ such that $\lim_{x\to 1}\phi(x, \bar{\iota})$ exists and is finite.*

Similarly, by letting $U_1 := V_1'$, from (10) we can see that U_1 solves the following differential equation:

$$U_1'(x) + \frac{2}{\sigma^2}\left(\tilde{\beta} - \frac{\gamma}{1-x}\right)U_1(x) = -\frac{2}{\sigma^2}\left(\frac{\ell}{1-x} + \frac{\kappa}{x(1-x)}\right). \tag{15}$$

Consider the class of functions $\{\psi(\cdot, c)\}$ indexed by c, where each $\psi(\cdot, c)$ is defined as

$$\psi(x,c) := e^{2\tilde{\beta}(1-x)/\sigma^2}(1-x)^{-2\gamma/\sigma^2} \times \left[\frac{2}{\sigma^2}\int_0^{1-x} e^{-2\tilde{\beta}u/\sigma^2}u^{2\gamma/\sigma^2-1}\left(\ell + \frac{\kappa}{1-u}\right)\mathrm{d}u + c\right].$$

It is straightforward to check that $\psi(\cdot, c)$ is a solution to (15) for each fixed c. It is also easy to verify that $\lim_{x\to 1}\psi(x,c)$ exits and is finite if and only if $c = 0$. Moreover, $\lim_{x\to 1}\psi(x,0) = (\ell+\kappa)/\gamma$. For simplicity, in what follows we will simply write $\psi(x,0)$ as $\psi(x)$. Thus,

$$\psi(x) = e^{2\tilde{\beta}(1-x)/\sigma^2}(1-x)^{-2\gamma/\sigma^2}\frac{2}{\sigma^2}\int_0^{1-x} \times e^{-2\tilde{\beta}u/\sigma^2}u^{2\gamma/\sigma^2-1}\left(\ell + \frac{\kappa}{1-u}\right)\mathrm{d}u. \tag{16}$$

The next result, which examines the number of intersections that $\phi(\cdot, \iota)$ and $\psi(\cdot)$ can have, is a key stepping stone towards the main result to be stated in Theorem 1.

Lemma 1. *For any $\iota \le \bar{\iota}$, $\phi(\cdot, \iota)$ and $\psi(\cdot)$ can intersect at most twice on the open interval $(0, 1)$.*

Assumption 1. The two functions, $\phi(\cdot, \bar{\iota})$ and $\psi(\cdot)$, intersect at two points, $\bar{x}_0$ and $\bar{x}_1$, and

$$K \le \overline{K} := \int_{\bar{x}_0}^{\bar{x}_1}\left(\phi(y,\bar{\iota}) - \psi(y)\right)\mathrm{d}y. \tag{17}$$

Theorem 1 (Sequential switching policy). *Suppose that Assumption 1 is satisfied. (i) There exists some $\iota^* \le \bar{\iota}$ such that $\phi(\cdot, \iota^*)$ and $\psi(\cdot)$ intersect exactly twice at some x_0^* and x_1^*, such that*

$$K = \int_{x_0^*}^{x_1^*}(\phi(y,\iota^*) - \psi(y))\,\mathrm{d}y.$$

(ii) A sequential switching policy characterized by (x_0^*, x_1^*) is optimal for Problem (4).

There are two possible ways Assumption 1 can be violated. It is possible that the function graphs of $\phi(\cdot, \iota^*)$ and $\psi(\cdot)$ never intersect. They could also intersect, but the area of intersection, $\overline{K}$ is less than the fixed cost K. In both cases, we interpret the violation to mean that intervention is too costly and therefore not worth considering; that is, it is optimal never to lock down.

We next develop some numerical examples to illustrate our main results. We first consider a base example with $\beta = 1$, $\tilde{\beta} = 0.2$, $\gamma = 1$, $\sigma = 0.5$, $\ell = 1$, $\kappa = 0.2$, and $K = 0.2$. Following Lemma 1, we compute $\overline{K} = 0.266$, $\bar{\iota} = 3.92$, $\iota^* = 3.86$, $x_0^* = 0.033$, and $x_1^* = 0.493$. In Figure 1, we graph three functions $\psi(x)$, $\phi(x, \bar{\iota})$ and $\phi(x, \iota^*)$ as functions of x. Here the area of the shaded region is $K = 0.2$.

To investigate how the model parameters affect our optimal decision, we perform sensitivity analysis by varying one parameter at a time while keeping the other parameters at the values specified in our base case (as given in the caption of Figure 1). Figure 2 displays the values of the switching thresholds $(x_0^\star, x_1^\star)$ as we vary β, γ, and ℓ. These plots reveal interesting patterns. For instance, we observe from the first plot that the gap between the two switching thresholds initially increases and then decreases as β increases. It is almost self-explanatory that $x_0^\star$ decreases with β. The threshold $x_1^\star$ increases as β rises from small to moderate levels because lockdown is relatively more costly than infection, and raising the upper threshold helps reduce the rate at which the system returns to on mode, given that the disease does not die out. However, as β increases to very high levels,

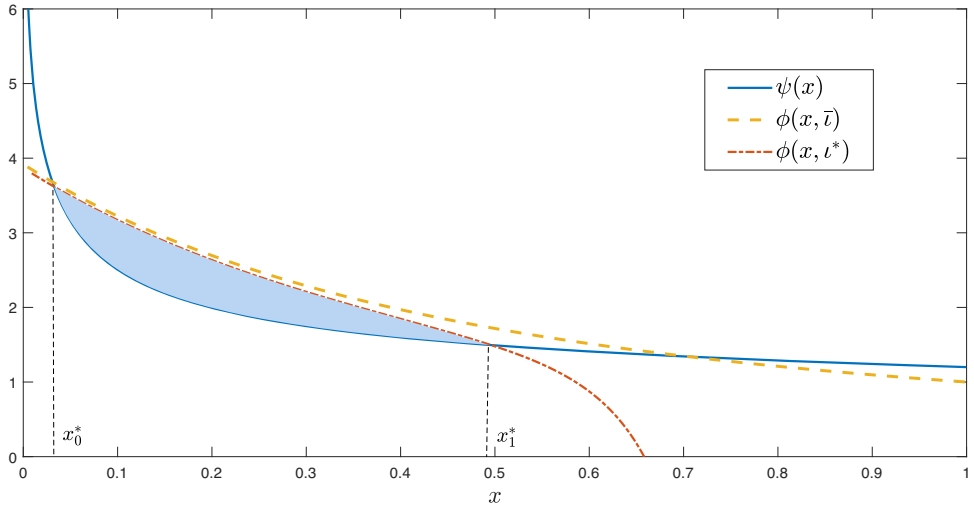

FIGURE 1 Illustration of the optimal switching policy: $x_0^* = 0.033$, $x_1^* = 0.493$, with $\beta = 1$, $\tilde{\beta} = 0.2$, $\gamma = 1$, $\sigma = 0.5$, $\ell = 1$, $\kappa = 0.2$, $K = 0.2$, $\overline{K} = 0.266$, $\bar{\iota} = 3.92$, $\iota^* = 3.86$.

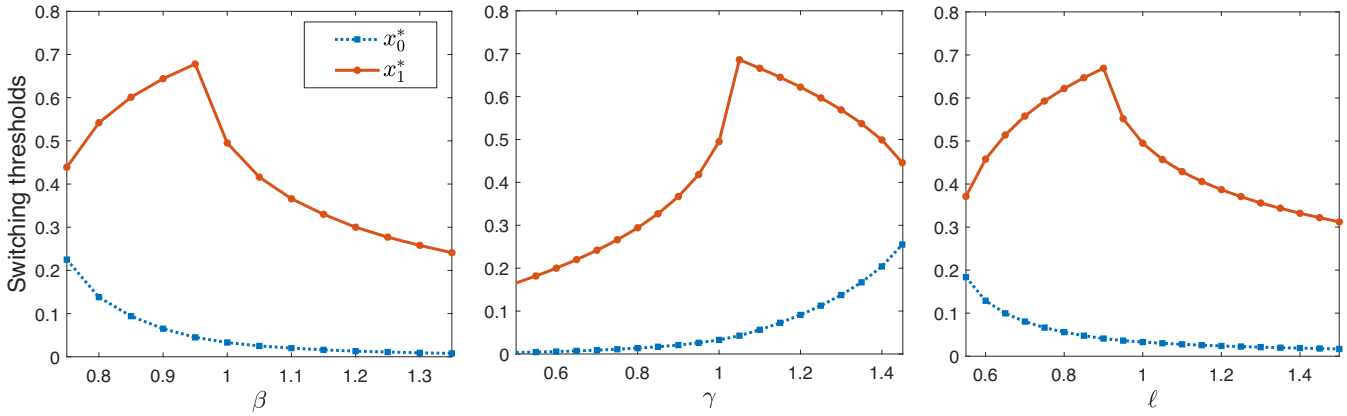

FIGURE 2 Impact of model parameters β, γ and ℓ on the values of the switching thresholds $(x_0^\star, x_1^\star)$.

suggesting that disease transmission is now devastating, it becomes more cost-effective to keep transmission under tight control and eliminate the disease with as few lockdowns as possible, explaining the decreasing trend in $x_1^\star$ at high β. The first and second plots exhibit symmetric behavior because the ratio β/γ reflects how deadly disease transmissions are in society. In other words, we can set the value of either rate parameter to one and normalize the other accordingly. The non-monotonic behavior of $x_1^\star$ and monotonic decrease of $x_0^\star$ with respect to ℓ follow similar reasoning. We observe similar non-monotonic patterns when we vary other model parameters; see Section C in the appendix for sensitivity analysis of $\tilde{\beta}$, κ, σ, and K.

5 | AN EXTENSION

We have so far considered scenarios where the social planner can only implement a binary lockdown policy. However, in real-world situations, there may be multiple levels of lockdown available, each with different levels of stringency and associated costs (e.g., full lockdown versus partial lockdown). For instance, many universities have adopted a "hybrid" teaching format, where students attend classes both in-person and online on alternate days, which can be viewed as a form of partial lockdown.

Motivated by this observation, we now extend our analysis to consider situations where the social planner can implement multiple levels of lockdown and make policy recommendations on how best to use them. Specifically, we assume that there are m levels of lockdown, indexed by $i = 1, \dots, m$, resulting in $m+1$ different modes. Under level i lockdown, the disease transmission rate is reduced from β to β_i, while incurring continuous economic tolls at the rate of κ_i. We assume that $\beta > \beta_1 > \cdots > \beta_m$ and $0 < \kappa_1 < \cdots < \kappa_m$, so that higher levels of lockdown are more stringent and more costly.

We model the social planner's decisions using an adapted, finite variation, càdlàg process Y with values in the set $0, 1, \dots, m$, where $Y_t = 0$ indicates that no lockdown measures are being implemented, and $Y_t = i$ indicates that level i lockdown is in effect. For simplicity, we refer to no lockdown as level 0 lockdown. We assume that switching between adjacent levels is possible, and upgrading from level i to level $i+1$ lockdown incurs a fixed cost $K_{i,i+1}$, while downgrading from level $i+1$ lockdown is free. This assumption is realistic since small and gradual changes are often more acceptable to the public. It should be noted that allowing a switch-over to occur between arbitrary pairs can lead to significant technical difficulties, as observed by Chernoff and Petkau (1978). In fact, for problems involving more than two control regimes, "the analytic approach becomes cumbersome".

Let $k(y) = \sum_{i=1}^m \kappa_i 1_{\{y=i\}}$ and define

$$K(u,v) := \begin{cases} K_{u,v} & \text{if} \quad (u,v) = (i, i+1) \\ & \text{for} \quad i = 0, \dots, m-1, \\ 0 & \text{otherwise.} \end{cases}$$

The objective of the social planner now becomes to

$$\text{minimize } \mathbb{E}\left[\int_0^\tau \ell X_u \mathrm{d}u + \int_0^\tau k(Y_u)\mathrm{d}u + \sum_{u\le\tau} K(Y_{u-}, Y_u)[\Delta Y_u]^+ \Big| X_0 = x, Y_0 = y\right]. \tag{18}$$

In the above, the stopping time τ again denotes the time at which X hits zero for the first time, while X now satisfies the SDE

$$\mathrm{d}X_t = [b(Y_t)(1 - X_t) - \gamma] X_t \mathrm{d}t + \sigma\sqrt{X_t(1 - X_t)}\mathrm{d}B_t,$$

where b is redefined as $b(y) := \beta 1_{\{y=0\}} + \sum_{i=1}^m \beta_i 1_{\{y=i\}}$.

As before, let $V(x, y)$ denote the value function associated with the problem. With reference to the general control theory, we anticipate the function V to satisfy

$$\min\Big\{ \frac{\sigma^2}{2} x(1-x) V_{xx}(x,y) + \left(b(y)x - b(y)x^2 - \gamma x\right) V_x(x,y) + \ell x + k(y), \ \min_i \{V(x,i+1) + K(i,i+1) - V(x,i), V(x,i) - V(x,i+1)\}\Big\} = 0, \tag{19}$$

subject to the condition $V(0, y) = 0$ and the requirement that $\lim_{x\to 1} V_x(x, y)$ exists and is finite.

Our analysis in the preceding section motivates us to consider the following functions:

$$\psi_i(x,c) := e^{2\beta_i(1-x)/\sigma^2}(1-x)^{-2\gamma/\sigma^2} \times \left[\frac{2}{\sigma^2}\int_0^{1-x} e^{-2\beta_i u/\sigma^2} u^{2\gamma/\sigma^2-1}\left(\ell + \frac{\kappa_i}{1-u}\right)\mathrm{d}u + c\right] \tag{20}$$

for $i = 1, \ldots, m$, plus the function ϕ as defined by (14). In particular, note that, for each fixed c, $\psi_i(\cdot, c)$ is a solution function to the following first-order differential equation:

$$U_i'(x) + \frac{2}{\sigma^2}\left(\tilde{\beta}_i - \frac{\gamma}{1-x}\right)U_i(x) = -\frac{2}{\sigma^2}\left(\frac{\ell}{1-x} + \frac{\kappa_i}{x(1-x)}\right).$$

Recall from the base-case scenario (i.e., $m = 1$) that in order for a sequential switching policy to be optimal, the fixed cost cannot be too high. Precisely, using the current set of notations, we would need to ensure that the area of intersection between $\phi(\cdot, \bar{\imath})$ and $\psi_1(\cdot, 0)$ should be greater than or equal to $K_{0,1}$, so that some $\imath^*$ exists to make the area of intersection between $\phi(\cdot, \imath^*)$ and $\psi_1(\cdot, 0)$ equal exactly $K_{0,1}$. This motivates the following condition as a natural extension of Assumption 1 to ensure that a sequential switching policy utilizing all possible lockdown levels is optimal.

Assumption 2. There exist $\imath^* \leq \bar{\imath}$ and $(c_1^*, \ldots, c_m^*)$ such that the following conditions hold:

(i) $\phi(\cdot, \imath^*)$ and $\psi_1(\cdot, c_1^*)$ intersect twice at $x_{1,0}$ and $x_{0,1}$ and

$$\int_{x_{1,0}}^{x_{0,1}} \left[\phi(x, \imath^*) - \psi_1(x, c_1^*)\right] \mathrm{d}x = K_{0,1};$$

(ii) for $i = 1, \ldots, m-1$, $\psi_i(\cdot, c_i^*)$ and $\psi_{i+1}(\cdot, c_{i+1}^*)$ intersect twice at $x_{i+1,i}$ and $x_{i,i+1}$, and

$$\int_{x_{i+1,i}}^{x_{i,i+1}} \left[\psi_i(x, c_i^*) - \psi_{i+1}(x, c_{i+1}^*)\right] \mathrm{d}x = K_{i,i+1};$$

(iii) $\lim_{x\to 1} \psi_m(x, c_m^*)$ is finite.

Some comments are in order. Part (*ii*) of Assumption 2 implies that $\psi_{m-1}(x, c_{m-1}^*) < \psi_m(x, c_m^*)$ for all $x > x_{m-1,m}$. This observation, in turn, implies that $c_{m-1}^* \leq 0$, because we have $\lim_{x\to 1} \psi_m(x, c_m^*) < \infty$ but $\lim_{x\to 1} \psi_{m-1}(x, c) = \infty$ for all $c > 0$. Using backward induction, we can deduce that $c_i \leq 0$ for all $i = 1, \ldots, m-1$. By the same token, we can conclude from part(i) of Assumption 2 that $\imath^\star \leq \bar{\imath}$. This is because $\lim_{x\to 1} \psi_1(x, c_1^*) < \infty$ whereas $\lim_{x\to 1} \phi(x, \imath) = \infty$ for all $\imath > \bar{\imath}$. It is easily verifiable from the explicit expressions given by (20) that part (*iii*) of Assumption 2 necessarily implies that $c_m^* = 0$.

Theorem 2 (Sequential switching under multiple lockdown levels). *If Assumption 2 holds with $x_{1,0} < x_{2,1} < \cdots < x_{m,m-1}$ and $x_{0,1} < x_{1,2} < \cdots < x_{m-1,m}$, then a sequential switching policy characterized by the two strongly ordered sequences is optimal for the problem described by (18). That is, the social planner switches from*

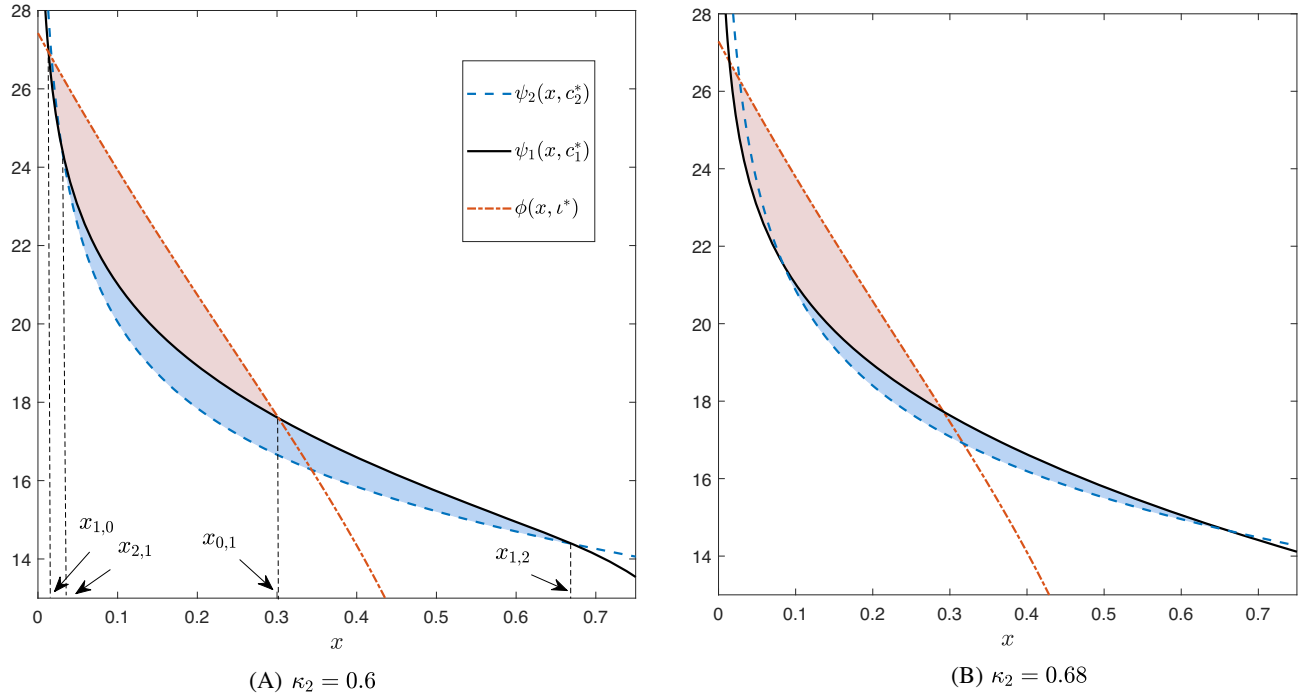

(A) $\kappa_2 = 0.6$ (B) $\kappa_2 = 0.68$

FIGURE 3 Optimal switching policies with multiple lockdown levels: $m = 2$, $\beta = 0.45$, $\beta_1 = 0.2$, $\beta_2 = 0.1$, $\gamma = 1$, $\sigma = 0.5$, $\ell = 6$, $K_{0,1} = 0.5$, $K_{1,2} = 0.45$, $\kappa_1 = 0.4$, $\kappa_2 = 0.6$ (left) and 0.68 (right).

level i to $i+1$ ($i+1$ to i) as soon as the infection level becomes above $x_{i,i+1}$ (below $x_{i+1,i}$).

It is worth pointing out that a violation of Assumption 2 does not necessarily mean that switching between different modes is not worthwhile. For example, it could be that all the conditions in Assumption 2 hold, but with m therein replaced by some $j(< m)$ but not with $j+1$. We interpret this case to mean that the social planner should consider adopting lockdown levels up to j but not to $j+1$. This observation motivates us to devise an iterative scheme to compute the optimal intervention strategy with $m+1$ modes. The idea is to examine one lockdown level at a time, starting from the lowest (i.e., level 1) to see if it is worth consideration, until a lockdown level is found to be not worth pursuing. A procedure to compute all switching thresholds is described in Appendix B.

We use a numerical example having three intervention levels (with $m = 2$) to illustrate our results in Theorem 2. Let $\beta = 0.45, \beta_1 = 0.2, \beta_2 = 0.1, \gamma = 1, \sigma = 0.5, \ell = 6, \kappa_1 = 0.4$, $\kappa_2 = 0.6$, and $K_{0,1} = 0.5, K_{1,2} = 0.45$. Using the iterative scheme described previously, we compute $\iota^* = 27.42$, $c_1^* = -0.0015$, $c_2^* = 0$, $x_{1,0} = 0.014$, $x_{0,1} = 0.301$, $x_{2,1} = 0.030$, $x_{1,2} = 0.778$. In the left-hand panel of Figure 3 we graph three functions: $\phi(x, \iota^*)$, $\psi_1(x, c_1^*)$ and $\psi(x, c_2^*)$ as functions of x. Here, the areas of the two shaded regions are equal to $K_{0,1}$ and $K_{1,2}$. Because $x_{1,0} < x_{2,1}$ and $x_{0,1} < x_{1,2}$, we conclude that the optimal sequential switching policy as described in Theorem 2 is optimal. On the other hand, when we increase κ_2 from 0.6 to 0.68 (so that level-2 lockdown becomes more costly, the optimal policy stipulates that we will never activate the level-2 lockdown because the bottom shaded area is smaller than $K_{0,1}$ as shown in the right-hand panel of Figure 3.

6 | CONCLUDING REMARKS

The emergence of new and highly contagious diseases, such as COVID-19, has brought to the forefront the importance of developing decision-support tools to aid governments in making unconventional interventions that do not rely solely on vaccines. Unlike previous studies that focus on pharmaceutical interventions, we propose an optimal control framework to determine the optimal timing for governments to intervene and withdraw measures. Our results demonstrate that the decision to intervene and when to act is heavily influenced by both the biological properties of the disease and the socioeconomic costs of the intervention. This highlights the importance of considering both the health and economic impacts when designing strategies to control infectious diseases.

There are several avenues for future research in this area. One potential direction is to consider a finite-time control problem, which involves minimizing the terminal disease prevalence by a fixed time rather than the cumulative disease costs, as proposed by Bai et al. (2022). In this context, the optimal intervention policy would become time nonstationary, depending not only on the current state of the system but also on the remaining time until the fixed time horizon. We leave this extension to future research.

While our work has focused on relatively simple disease dynamics, we acknowledge that COVID-19 has several realistic features, such as the incubation time, symptomatic and asymptomatic individuals, and isolated individuals, which are not captured by our model. Incorporating these features would significantly enlarge the state space and give rise to a control problem with partial observation. For example, the incorporation of a positive incubation period would require tracking the number of individuals who have been exposed but have not developed symptoms, and the number of new infections may not be directly observable to the social planner. Hence, relaxing such assumptions would result in a more complex problem than a simple sequential switching policy, which we leave to future studies.

Moreover, our analysis has been based on the assumption that the infection cost is proportional to the fraction of infected individuals. However, in reality, medical resources are limited, and the demand for medical treatment can overwhelm the healthcare system as the number of infected people grows, leading to people not receiving medical care on time or receiving low-quality care. As a result, the marginal social-economic cost associated with infection may increase as the number of infected people grows. This can be modeled by replacing the linear cost rate with a general cost rate function that varies with the number of infected individuals. We propose a specific function, $\ell(x) = \ell_1 x$ for $0 < x < \bar{x}$ and $\ell(x) = \ell_2 x$ for $\bar{x} \le x < 1$, where $\bar{x}$ denotes the maximum level of infected individuals that can be treated by the available medical resources, and ℓ_i $(0 < \ell_1 < \ell_2)$ denotes the marginal costs of infection when the total number of infections is below or above the threshold $\bar{x}$. We envision that the dynamic programming equation describing the optimal intervention strategy can be written directly, subject to appropriate conditions, and we leave the analysis of this extension to future research.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

Xu Sun https://orcid.org/0000-0003-2560-7370
Yunan Liu https://orcid.org/0000-0001-9961-2610

REFERENCES

Bai, M., Cui, Y., Kong, G., & Zhang, Z. (2022). No panic in pandemic: The impact of individual choice on public health policy and vaccine priority. *Working Paper* URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3763514

Behncke, H. (2000). Optimal control of deterministic epidemics. *Optimal Control Applications and Methods*, *21*(6), 269–285.

Cai, Y., Kang, Y., & Wang, W. (2017). A stochastic SIRS epidemic model with nonlinear incidence rate. *Applied Mathematics and Computation*, *305*, 221–240.

Chehrazi, N., Cipriano, L. E., & Enns, E. A. (2019). Dynamics of drug resistance: Optimal control of an infectious disease. *Operations Research*, *67*(3), 619–650.

Chen, N., Hu, M., & Zhang, C. (2022). Capacitated SIR model with an application to COVID-19. *Working paper* URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3692751

Chen, Z., & Kong, G. (2022). Hospital admission and social distancing: An SEIR model with constrained medical resources. *Forthcoming in Production and Operations Mangement* URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3704897

Chernoff, H., & Petkau, A. J. (1978). Optimal control of a brownian motion. *SIAM Journal on Applied Mathematics*, *34*(4), 717–731.

Duckworth, K., & Zervos, M. (2001). A model for investment decisions with switching costs. *Annals of Applied Probability*, *11*(1), 239–260.

El Housni, O., Sumida, M., Rusmevichientong, P., Topaloglu, H., & Ziya, S. (2022a). Can testing ease social distancing measures? Future evolution of COVID-19 in NYC. *Working paper* URL http://ziya.web.unc.edu/wp-content/uploads/sites/15166/2020/07/CovidModelandAnalysis-NYC-Version-May-28.pdf

El Housni, O., Sumida, M., Rusmevichientong, P., Topaloglu, H., & Ziya, S. (2022b). Future evolution of COVID-19 pandemic in North Carolina: Can we flatten the curve? *Working paper* URL http://ziya.web.unc.edu/wp-content/uploads/sites/15166/2020/07/CovidModelandAnalysis-NC-Version-July-2.pdf

Feller, W. (1954). Diffusion processes in one dimension. *Transactions of the American Mathematical Society*, *77*(1), 1–31.

Gray, A., Greenhalgh, D., Hu, L., Mao, X., & Pan, J. (2011). A stochastic differential equation SIS epidemic model. *SIAM Journal on Applied Mathematics*, *71*(3), 876–902.

Hinch, R., Probert, W. J., Nurtay, A., Kendall, M., Wymant, C., Hall, M., Lythgoe, K., Bulas Cruz, A., Zhao, L., Stewart, A., Ferretti, L., Montero, D., Warren, J., Mather, N., Abueg, M., Wu, N., Legat, O., Bentley, K., Mead, T., … Fraser, C. (2021). Openabm-COVID-19—an agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing. *PLoS Computational Biology*, *17*(7), e1009146.

Jenkins, P. A., & Spano, D. (2017). Exact simulation of the Wright–Fisher diffusion. *The Annals of Applied Probability*, *27*(3), 1478–1509.

Kurtz, T. G. (1971). Limit theorems for sequences of jump markov processes approximating ordinary differential processes. *Journal of Applied Probability*, *8*(2), 344–356.

Lee, S., Zabinsky, Z., Kofsky, S., & Liu, S. (2022). COVID-19 pandemic response simulation: Impact of non-pharmaceutical interventions on ending lockdowns. *Working paper* URL https://www.medrxiv.org/content/10.1101/2020.04.28.20080838v1

Lefévre, C. (1981). Optimal control of a birth and death epidemic process. *Operations Research*, *29*(5), 971–982.

Ly Vath, V., & Pham, H. (2007). Explicit solution to an optimal switching problem in the two-regime case. *SIAM Journal on Control and Optimization*, *46*(2), 395–426.

Mao, X., Marion, G., & Renshaw, E. (2002). Environmental brownian noise suppresses explosions in population dynamics. *Stochastic Processes and their Applications*, *97*(1), 95–110.

Sethi, S. P., & Staats, P. W. (1978). Optimal control of some simple deterministic epidemic models. *Journal of the Operational Research Society*, *29*(2), 129–136.

Tatapudi, H., Das, R., & Das, T. K. (2020). Impact assessment of full and partial stay-at-home orders, face mask usage, and contact tracing: An agent-based simulation study of COVID-19 for an urban region. *Global Epidemiology*, *2*, 100036.

Tatapudi, H., & Das, T. K. (2021). Impact of school reopening on pandemic spread: A case study using an agent-based model for COVID-19. *Infectious Disease Modelling*, *6*, 839–847.

Tebbens, R. J. D., & Thompson, K. M. (2009). Priority shifting and the dynamics of managing eradicable infectious diseases. *Management Science*, *55*(4), 650–663.

Tuckwell, H. C., & Williams, R. J. (2007). Some properties of a simple stochastic epidemic model of SIR type. *Mathematical Biosciences*, *208*(1), 76–97.

Wu, J., & Chao, X. (2014). Optimal control of a Brownian production/inventory system with average cost criterion. *Mathematics of Operations Research*, *39*(1), 163–189.

Yaesoubi, R., & Cohen, T. (2011). Dynamic health policies for controlling the spread of emerging infections: Influenza as an example. *PloS One*, *6*(9).

Zervos, M., Johnson, T. C., & Alazemi, F. (2013). Buy-low and sell-high investment strategies. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, *23*(3), 560–578.

Zhang, E. (2022). Optimal timing and effectiveness of COVID-19 outbreak responses in china: a modelling study. *BMC Public Health*, *22*, 1–15.

How to cite this article: Sun, X., & Liu, Y. (2023). Optimal interventions of infectious disease. *Naval Research Logistics (NRL)*, 1–14. https://doi.org/10.1002/nav.22114

APPENDIX A: PROOFS

Proof of Proposition 1.

Let

$$\bar{\iota} := \frac{2\ell}{\sigma^2}\int_0^1 e^{2\beta u/\sigma^2}(1-u)^{2\gamma/\sigma^2-1}\mathrm{d}u. \tag{A1}$$

Substituting (A1) into (14) yields

$$\phi(x,\bar{\iota}) = e^{-2\beta x/\sigma^2}(1-x)^{-2\gamma/\sigma^2}\frac{2\ell}{\sigma^2} \times \int_x^1 e^{2\beta u/\sigma^2}(1-u)^{2\gamma/\sigma^2-1}\mathrm{d}u.$$

By applying l'Hospital rule, we obtain

$$\begin{aligned}\lim_{x\to 1}\phi(x,\bar{\iota}) &= \lim_{x\to 1}\frac{-\frac{2\ell}{\sigma^2}e^{2\beta x/\sigma^2}(1-x)^{2\gamma/\sigma^2-1}}{\frac{2\beta}{\sigma^2}e^{2\beta x/\sigma^2}(1-x)^{2\gamma/\sigma^2}-\frac{2\gamma}{\sigma^2}e^{2\beta x/\sigma^2}(1-x)^{2\gamma/\sigma^2-1}}\\ &= \ell/\gamma.\end{aligned}$$

It is easy to see that $\lim_{x\to 1}\phi(x,\iota) = \infty$ for $\iota > \bar{\iota}$ and $\lim_{x\to 1}\phi(x,\iota) = -\infty$ for $\iota < \bar{\iota}$, meaning that $\iota = \bar{\iota}$ is the only

initial value for which $\lim_{x\to 1}\phi(x,\iota)$ is finite. The proof is thus complete. □

Proof of Lemma 1.

Define $\varphi(\cdot,\iota) := \phi(\cdot,\iota) - \psi(\cdot)$. Suppose, by way of contradiction, that $\phi(\cdot,\iota)$ and $\psi(\cdot)$ cross more than twice. Then for each fixed ι, the function graph of $\varphi(\cdot,\iota)$ must have crossed the horizontal line at least four times, two times from below and two times from above. On the other hand, since ϕ and ψ are solutions to Equations (13) and (15), we can use the two equations to deduce that

$$\frac{\sigma^2}{2}x(1-x)\varphi'(x,\iota) = \kappa - x(1-x)\psi(x) \quad \text{whenever} \quad \varphi(x,\iota) = 0. \tag{A2}$$

A tedious yet straightforward calculation involving (16) will show that $x(1-x)\psi(x)$ is strictly quasi concave. This, however, implies that the function graph of φ can the horizontal line at most three times thanks to (A2), leading to a contradiction. Therefore, $\phi(\cdot,\iota)$ and $\psi(\cdot)$ can cross at most twice, as desired. □

Proof of Theorem 1.

To establish part (i), we know from Proposition 1 that $\phi(x,\iota)$ decreases to negative infinity for any fixed $x \in [0,1)$ as ι approaches negative infinity. Therefore, there exists some $\underline{\iota} \in (-\infty, \bar{\iota})$ such that $\phi(x,\iota)$ and $\psi(\cdot)$ do not intersect for any $\iota < \underline{\iota}$, intersect but do not cross (i.e., touch each other) for $\iota = \underline{\iota}$, and intersect at least twice for any $\iota \in (\underline{\iota}, \bar{\iota}]$. On the other hand, part (i) of the lemma reveals that the number of points of intersection that $\phi(\cdot,\iota)$ and $\psi(\cdot)$ have can be at most two. It follows that $\phi(x,\iota)$ and $\psi(\cdot)$ must intersect exactly twice for $\iota \in (\underline{\iota}, \bar{\iota}]$. For such ι, let $x_0(\iota)$ and $x_1(\iota)$ denote the corresponding two points at which $\phi(\cdot,\iota)$ and $\psi(\cdot)$ intersect. Then

$$x_1(\iota) - x_0(\iota) \to 0 \quad \text{and} \quad \int_{x_0(\iota)}^{x_1(\iota)} [\phi(x,\iota) - \psi(x)]\,\mathrm{d}x \to 0 \quad \text{as} \quad \iota \to \underline{\iota}.$$

When $\iota = \bar{\iota}$, $x_0(\iota) = \bar{x}_0$ and $x_1(\iota) = \bar{x}_1$. Since $x_0(\iota)$ and $\phi(x,\iota)$ are monotonically increasing in ι, and $x_1(\iota)$ is monotonically decreasing in ι, $\int_{x_0(\iota)}^{x_1(\iota)} [\phi(x,\iota) - \psi(x)]\,\mathrm{d}x$ increases from 0 to $\overline{K}$, which is greater than or equal to K by our hypothesis. Hence, we conclude that there exists some $\iota^* \in (\underline{\iota}, \bar{\iota}]$ such that

$$\int_{x_0(\iota^*)}^{x_1(\iota^*)} \left[\phi(x,\iota^*) - \psi(x)\right]\mathrm{d}x = K.$$

Denoting $x_0^* = x_0(\iota^*)$ and $x_1^* = x_1(\iota^*)$, we complete the proof of part (i).

Towards establishing part (ii), let $V(x,y)$ be such that (a) $V(0,0) = 0$, (b) $V(x_0^*,0) = V(x_0^*,1)$, (c) $V_x(x,0) = \phi(x,\iota^*)1_{\{x<x_1^*\}} + \psi(x)1_{\{x\geq x_1^*\}}$ and $V_x(x,1) = \phi(x,\iota^*)1_{\{x\leq x_0^*\}} + \psi(x)1_{\{x>x_0^*\}}$. By construction, $V(x,y)$ satisfies the variational inequality (8). For each $n \geq 1$, define $\tau_n := \tau \wedge n$, where we recall that τ is the random time when X hits zero for the first time. Using the Itô-Tanaka formula, we can find that

$$\begin{aligned} V(X_t,Y_t) = V(x,y) &+ \int_0^t \left[\frac{\sigma^2}{2}X_u(1-X_u)V_{xx}(X_u,Y_u) \right. \\ &\left. + (b(Y_u)(1-X_u) - \gamma)X_u V_x(X_u,Y_u)\right]\mathrm{d}u \\ &+ \int_0^t \sqrt{X_u(1-X_u)}V_x(X_u,Y_u)\mathrm{d}B_u \\ &+ \sum_{u\leq t}[V(X_u,Y_{u+}) - V(X_u,Y_u)]. \end{aligned}$$

This implies that

$$\begin{aligned} &\int_0^t \ell X_u \mathrm{d}u + \kappa\int_0^t Y_u\mathrm{d}u + K\sum_{u\leq t}[\Delta Y_u]^+ \\ &= V(x,y) - V(X_t,Y_t) + \int_0^t \left[\frac{\sigma^2}{2}X_u(1-X_u)V_{xx}(X_u,Y_u) \right. \\ &\quad \left. + (b(Y_u)(1-X_u) - \gamma)X_u V_x(X_u,Y_u) + \ell X_u + \kappa Y_u\right]\mathrm{d}u \\ &\quad + \int_0^t \sqrt{X_u(1-X_u)}V_x(X_u,Y_u)\mathrm{d}B_u \\ &\quad + \sum_{u\leq t}[V(X_u,Y_{u+}) - V(X_u,Y_u) + K]\,[\Delta Y_u]^+ \\ &\quad + \sum_{u\leq t}[V(X_u,Y_{u+}) - V(X_u,Y_u)]\,[\Delta Y_u]^-. \end{aligned}$$

Because $V(x,y)$ satisfies (8), we know that

$$\begin{aligned} &\int_0^t \ell X_u \mathrm{d}u + \kappa\int_0^t Y_u\mathrm{d}u + K\sum_{u\leq t}[\Delta Y_u]^+ \\ &\geq V(x,y) - V(X_t,Y_t) + \int_0^t \sqrt{X_u(1-X_u)}V_x(X_u,Y_u)\mathrm{d}B_u. \end{aligned} \tag{A3}$$

Letting $t = \tau_n$, taking expectations on both sides, and noting that τ_n is bounded, we obtain

$$\begin{aligned} &\mathbb{E}\left[\int_0^{\tau_n} \ell X_u \mathrm{d}u + \kappa\int_0^{\tau_n} Y_u\mathrm{d}u + K\sum_{u\leq \tau_n}[\Delta Y_u]^+\right] \\ &\geq V(x,y) - \mathbb{E}\left[V(X_{\tau_n},Y_{\tau_n})\right]. \end{aligned} \tag{A4}$$

Since $\tau_n \to \tau$ almost surely as $n \to \infty$, we can apply the monotone convergence theorem to conclude

$$\begin{aligned} &\mathbb{E}\left[\int_0^{\tau_n} \ell X_u \mathrm{d}u + \kappa\int_0^{\tau_n} Y_u\mathrm{d}u + K\sum_{u\leq \tau_n}[\Delta Y_u]^+\right] \\ &\to \mathbb{E}\left[\int_0^{\tau} \ell X_u \mathrm{d}u + \kappa\int_0^{\tau} Y_u\mathrm{d}u + K\sum_{u\leq \tau}[\Delta Y_u]^+\right] \end{aligned}$$

Similarly, we can apply the bounded convergence theorem to obtain

$$\mathbb{E}\left[V(X_{\tau_n},Y_{\tau_n})\right] \to \mathbb{E}\left[V(X_\tau,Y_\tau)\right] = 0,$$

where the equality uses the fact that $X_\tau = 0$ and the boundary condition that $V(0,0) = 0$. Thus, by passing to the limit $n \to \infty$ in (A4), we obtain

$$V(x,y) \leq \mathbb{E}\left[\int_0^{\tau} \ell X_u \mathrm{d}u + \kappa\int_0^{\tau} Y_u\mathrm{d}u + K\sum_{u\leq \tau}[\Delta Y_u]^+\right]. \tag{A5}$$

If Y^* is the sequential switching policy with switching boundaries x_0^* and x_1^*, then (A3) holds with equality. By adopting the same argument as before, we can conclude that (A5) holds with equality, which completes the proof of part (ii). □

Proof of Theorem 2.

The key to the proof is to construct a solution to (19). For this purpose, let $V(x,y)$ be such that (a) $V(0,0) = 0$, (b) $V(x^*_{i,i+1}, i) = V(x^*_{i,i+1}, i+1)$ for $i = 0, \dots, m-1$, and (c)

$$V_x(x,y) = \begin{cases} \phi(x,\iota^*)1_{\{x<x^*_{0,1}\}} + \sum_{j=1}^{m}\psi(x,c_j^*)1_{\{x^*_{j-1,j}<x<x^*_{j+1,j}\}} \\ \qquad \text{if } y = 0, \\ \phi(x,\iota^*)1_{\{x\le x^*_{1,0}\}} + \sum_{j=1}^{i-1}\psi(x,c_j^*)1_{\{x^*_{j,j-1}<x\le x^*_{j+1,j}\}} \\ \qquad + \psi(x,c_i^*)1_{\{x^*_{i,i-1}<x<x^*_{i+1,i}\}} \\ \qquad + \sum_{j=i+1}^{m}\psi(x,c_j^*)1_{\{x^*_{j-1,j}\le x<x^*_{j,j+1}\}} \\ \qquad \text{if } y \neq 0. \end{cases}$$

By construction and applying Assumption 2, it is easy to see that $V(x,y)$ satisfies (19). Given this, the remaining steps towards the desired conclusion will mimic those in the proof of Theorem 1. Hence, we omit the details. □

APPENDIX B: AN ITERATIVE PROCEDURE TO COMPUTE SWITCHING THRESHOLDS

Suppose level l lockdown is considered worthwhile with $l < m$. This would imply that (i) there exist $\iota^{(l)}$ and $(c_1^{(l)}, \dots, c_l^{(l)})$ that collectively satisfy a modified version of Assumption 2 (changing m therein to l), and (ii) the corresponding points of intersection are such that $x_{1,0}^{(l)} < x_{2,1}^{(l)} < \cdots < x_{l,l-1}^{(l)}$ and $x_{0,1}^{(l)} < x_{1,2}^{(l)} < \cdots < x_{l-1,l}^{(l)}$. To examine if level $l+1$ lockdown is worthwhile, one computes the area of intersection between $\psi_l(\cdot, c_l^{(l)})$ and $\psi_{l+1}(\cdot, 0)$, denoted as $K_{l,l+1}^{(l)}$. ($K_{l,l+1}^{(l)} = 0$ if the two functions graphs touch or do not intersect).

- If $K_{l,l+1}^{(l)} \le K_{l,l+1}$, then level $l+1$ lockdown is not worth consideration. Hence, the procedure terminates and concludes that a sequential switching policy with lockdowns up to level l is optimal. In particular, the switching policy is characterized by two strongly ordered sequences: $x_{1,0}^{(l)} < x_{2,1}^{(l)} < \cdots x_{l,l-1}^{(l)}$ and $x_{0,1}^{(l)} < x_{1,2}^{(l)} < \cdots x_{l-1,l}^{(l)}$.

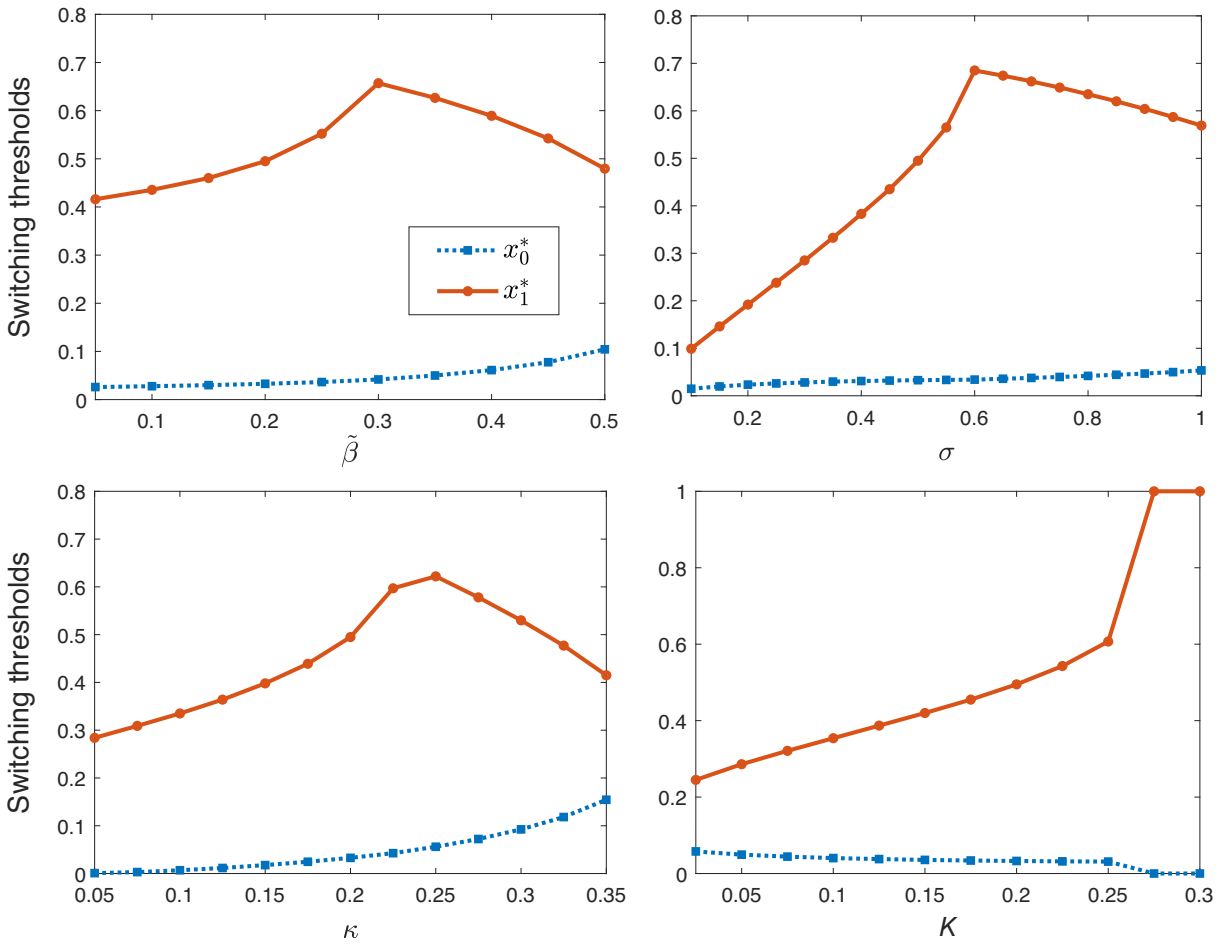

FIGURE C1 Impact of model parameters $\tilde{\beta}$, σ, κ and K on the values of the switching thresholds $(x_0^\star, x_1^\star)$. In the last plot on K, $(x_0^\star, x_1^\star) = (0,1)$ when K is sufficiently large (in this case $K > \overline{K} = 0.266$), meaning that the system stays "off" and never switches "on."

- If $K^{(l)}_{l,l+1} > K_{l,l+1}$, then one looks for $\iota^{(l+1)}$ and $(c^{(l+1)}_1, \dots, c^{(l+1)}_{l+1})$ that collectively satisfy a modified version of Assumption 2 (changing m therein to $l+1$).
  - If the points of intersection are such that $x^{(l+1)}_{1,0} < x^{(l+1)}_{2,1} < \cdots x^{(l+1)}_{l+1,l}$ and $x^{(l+1)}_{0,1} < x^{(l+1)}_{1,2} < \cdots x^{(l+1)}_{l,l+1}$, then level $l+1$ lockdown is considered worthwhile, in which case the procedure continues by setting $l \leftarrow l+1$ (unless $l+1=m$);
  - otherwise, the procedure stops and concludes the previously identified switching policy with lockdowns up to level l is optimal.

APPENDIX C: SENSITIVITY ANALYSIS

Next, we supplement the numerical experiments in the main paper by conducting sensitivity analyses for the other model parameters, including $\tilde{\beta}$, σ, κ and K; See Figure C1. Similar to our findings in Figure 2, the switching thresholds (x^*_0, x^*_1) exhibits non-monotonic behavior.