



Research article

Optimal policy for control of epidemics with constrained time intervals and region-based interactions

Xia Li^{1,*}, Andrea L. Bertozzi^{2,*}, P. Jeffrey Brantingham³ and Yevgeniy Vorobeychik⁴

¹ Department of Mathematics University of California, Los Angeles, USA

² Departments of Mathematics and Mechanical and Aerospace Engineering, University of California, Los Angeles, USA

³ Department of Anthropology, University of California, Los Angeles, USA

⁴ Department of Mathematics, Washington University in St. Louis, USA

* **Correspondence:** Email: xli51@g.ucla.edu, bertozzi@ucla.edu.

Abstract: We introduce a policy model coupled with the susceptible–infected–recovered (SIR) epidemic model to study interactions between policy-making and the dynamics of epidemics. We considered both single-region policies as well as game-theoretic models involving interactions among several regions and hierarchical interactions among policy-makers modeled as multi-layer games. We assumed that the policy functions are piece-wise constant with a minimum time interval for each policy stage, considering that policies cannot change frequently in time or be easily followed. The optimal policy was obtained by minimizing a cost function that consists of an implementation cost, an impact cost, and, in the case of multi-layer games, a non-compliance cost. We show, in a case study of COVID-19 in France, that when the cost function is reduced to the impact cost and parameterized as the final epidemic size, the solution approximates that of the optimal control in Bliman et al, (2021) for a sufficiently small minimum policy time interval. For a larger time interval, however, the optimal policy is a step down function, quite different from the step up structure typically deployed during the COVID-19 pandemic. In addition, we present a counterfactual study of how the pandemic would have evolved if herd immunity was reached during the second wave in the county of Los Angeles, California. Finally, we study a case of three interacting counties with and without a governing state.

Keywords: SIR model policy; epidemic dynamics

1. Introduction

In the course of battling COVID-19, public health policies sought to enforce non-pharmaceutical interventions to slow or halt the spread of the pandemic. Common policies included “safer-at-home”

“social distancing”, and “mask wearing” mandates, which were seen as crucial during the early stages of the pandemic prior to the availability of vaccines. The global and local timelines of COVID-19 ([7, 26]) indicate that policy evolution affected the spread of the pandemic and vice versa. For example, in the county of Los Angeles, social distancing was first mandated [9] on March 21, 2020, about a month after the first reported COVID-19 case in LA. Around that time, the Los Angeles Mayor’s Office released the safer-at-home policy [1]. One week later, beaches, hiking trails, dog parks, skate parks, and other public sites and facilities were temporarily closed. On April 15th, as infected cases continued to increase, facial coverings were mandated in many indoor places [19]. In hindsight, it is important to ask: Were policies enforced in an optimal way? What can we learn by using mathematical modeling to understand the interplay between policy and disease spread? This paper introduces a policy model coupled to a susceptible–infected–recovered (SIR) epidemic model to study interactions between policy-making and the dynamics of epidemics. There have been several studies on the relationship between policies and epidemics [4–6, 18, 21]. In a study analyzing data from 16 US cities during the 1918 pandemic [5], Bootsma and Ferguson analyzed specific outcomes related to the impact of delaying lockdown policies on the total deaths and on the occurrence of second outbreak waves due to reopening too early. The analysis was done by fitting available data to an SEIR model. They also considered optimal control for the simpler SIR and the end-state of the pandemic, noting that there exists an optimal control level with fewer deaths and no second wave. More recently, Bliman et al [3] developed a theoretical study of the optimal control of a classical SIR outbreak. They did not consider the possibility of vaccines or pharmaceutical interventions. Rather, focusing exclusively on non-pharmaceutical interventions, they designed an optimal policy that achieves an end state as close to the herd immunity threshold as possible. This is the same problem considered briefly in a section of [5]. Bliman et al proved the existence and uniqueness of their solution and showed the optimal social distancing policy is a bang-bang controller [2], generalizing the results of [3] by modeling without prescribing the starting date of the policy.

The substantial theoretical insights of Bliman et al’s model are limited in their practical implications by a few assumptions. First, Bliman et al. assumed that a policy can change continuously in time, which would imply, for example, the ability to shift in three successive instants between no restrictions, perfect “lockdown”, and back to no restrictions. As observed during the COVID-19 pandemic, policies that change frequently in time cannot be easily followed. Moreover, policies must be relatively easy to interpret, with a small number of different intensity levels (see Figure 3b). A practical implementation also requires a minimum time duration for a particular stage of the policy. These practical constraints can be modeled together as a piece-wise constant function of time with a minimum time interval for each well-defined policy level (i.e., not continuous). With this idea in mind, we aim to re-examine the optimal practical policy among all possible piece-wise constant policies with minimal time duration.

Second, Bliman et al., assumed that the only outcome to manage is the final epidemic size. This so-called “impact cost” is clearly a central concern (see below). However, as seen during the pandemic, there are real trade-offs between decreased infections and the negative impact of strict policies on other aspects of society such as remote learning for young students, employment curtailment in certain job sectors, and lack of key services provided to the public. In the present work, we modify Bliman et al’s model to take into account these other practical “implementation costs”. Specifying a short minimal time interval during which policies must remain constant (e.g., one week), we find that our results resemble Bliman et al.’s bang-bang controller [3] despite the more complex cost structure that includes

both impact and implementation costs. With a larger minimal time interval during which time policies must remain constant (e.g., 28 days), optimal policies depart from the bang-bang solution.

Finally, Bliman et al., also assumed a pandemic spreading in a single population pool overseen by a single policy-making entity. The reality of the COVID-19 pandemic is that there are policy makers at several (nested) hierarchical scales that oversee different population pools. For example, within the United States, policies may be set at Federal, State, County, and local levels, not to mention finer-grained institutional and family scales. And populations at any one scale (e.g., counties) may interact to varying degrees. Inspired by the work of Jia et al. [14], we introduce a hierarchical version of Bliman et al.'s model with sequential (Stackleberg) policy-making. Specifically, levels higher in a jurisdictional hierarchy make policy decisions, while levels lower in the hierarchy make their decisions with full knowledge of the policy recommendations from above. We find that a hierarchical structure can make the policies converge in all regions using the right weight for a non-compliance cost.

The remainder of this paper is organized as follows. We first introduce the work in [3] and reproduce the results using our methods. We discuss how different optimal policies result from different parameter choices for model constraints and costs. Next, we discuss an empirical case study of the so-called "second wave" of the pandemic (November 6, 2020–May 12, 2021) in Los Angeles County, California. Last, we use a simulation to study optimal control of the pandemic in three counties with and without a governing state as an example of the multi-layer multi-regional model.

2. Policy model using optimal control

A policy function is a continuous function that has a range of $[0, 1]$. As the numerical value increases, the strictness of the policy decreases. The numerical value 0 denotes a total lockdown and 1 denotes no control. We assume a policy $u(t)$ directly influences the level of a lockdown, which affects the rate of the population transport from compartment S to I . We use the following policy-incorporated SIR:

$$\left\{ \begin{array}{l} \frac{dS(t)}{dt} = -u(t)\beta\frac{I(t)S(t)}{N}, \\ \frac{dI(t)}{dt} = u(t)\beta\frac{I(t)S(t)}{N} - \gamma I(t), \\ \frac{dR(t)}{dt} = \gamma I(t), \\ S(0) = S_0, \quad I(0) = I_0, \quad R(0) = R_0. \end{array} \right. \quad (2.1)$$

Like the traditional SIR model, the reproduction number $\mathcal{R}_0 = \frac{\beta}{\gamma}$. Herd immunity occurs when a large proportion of the population has become immune to the infection. Mathematically, it is defined as the value of S below which the number of infected decreases and can be calculated as $S_{\text{herd}} = \frac{N}{\mathcal{R}_0}$. In [3], a policy $u(t)$ is assumed to belong to the admissible set $\mathcal{U}_{\alpha_{\max}, T_0}$ defined by

$$\{u \in \mathcal{L}^\infty([0, +\infty)), \alpha_{\max} \leq u(t) \leq 1 \text{ if } t \in [0, T_0], u(t) = 1 \text{ if } t > T_0\}.$$

The constant T_0 characterizes the duration of the policy, and α_{\max} characterizes its maximal intensity. In [3], Theorem 2.1 states that no finite time intervention is able to stop the epidemics before or exactly at the herd immunity. However, one may stop arbitrarily close to herd immunity by having a sufficiently

long intervention of sufficient intensity. To determine the closest state S to this threshold attainable by control of maximal intensity α_{\max} on the interval $[0, T_0]$, one is led to consider the following optimal control problem:

$$\sup_{u \in \mathcal{U}_{\alpha_{\max}, T_0}} S_{\infty}(u). \quad (2.2)$$

Furthermore, Bliman et al proved the existence and uniqueness of the optimal solution to problem 2.2 and that the solution is a bang-bang controller (a control that switches from one extreme to the other). More specifically, they have the following theorem:

Theorem 2.1. (Theorem 2.1 in [3]) *Let $\alpha_{\max} \in [0, 1)$ and $T_0 > 0$. Problem 2.2 admits a unique solution u^* . Furthermore,*

(i) *the maximal value $S_{\infty, \alpha_{\max}, T_0}^* := \{\max S_{\infty}(u) : u \in \mathcal{U}_{\alpha_{\max}, T_0}\}$ is non-increasing with respect to α_{\max} and non-decreasing with respect to T_0 .*

(ii) *there exists a unique $T_0^* \in [0, T_0)$ such that $u^* = u_{T_0^*} := \mathbb{1}_{[0, T_0^*]} + \alpha_{\max} \mathbb{1}_{[T_0^*, T_0]} + \mathbb{1}_{[T_0, +\infty)}$ (in particular, the optimal control is bang-bang).*

3. Single region case

We use the same policy-incorporated SIR model for the epidemic dynamics as in [3]. Instead of minimizing the final epidemic size alone, we adopt a similar policy-making process as in [14] by using a cost function that takes into account the cost of implementing the policy, the impact of the infection, and a penalty for being non-compliant. The latter cost only applies in hierarchical models where a lower-level unit can choose not to follow the policy recommendation of a higher-level unit.

We also consider practical implementation constraints, namely that the policy can only be implemented using a finite number of discrete levels of control and with a minimal time interval during which a policy must remain constant. As an example, consider the policy implementation in France during the year 2020 and 2021 shown in Figure 1 ([27]). Implemented policies were discrete both in terms of the small number of intervention types and the fixed time intervals of enforcement, the shortest of which was approximately 15 days in duration, with the longest lasting more than a year. A discrete policy model is realistic given the empirical pattern of real-world interventions. Such a model also simplifies the computation problem of optimal policy discovery by searching through a discrete set of potential policies rather than a continuum of policies.

3.1. The policy-incorporated SIR model

To model the evolution of the pandemic, we discretize the system of ODE using forward Euler's method with a time step of 1:

$$\begin{cases} S(t) = S(t-1) - \alpha \beta \frac{I(t-1)S(t-1)}{N}, \\ I(t) = I(t-1) + \alpha \beta \frac{I(t-1)S(t-1)}{N} - \gamma I(t-1), \\ R(t) = R(t-1) + \gamma I(t-1), \\ S(0) = S_0, \quad I(0) = I_0, \quad R(0) = R_0. \end{cases} \quad (3.1)$$

Where $\alpha = u(t-1)$. Equation (3.1) can be seen as a first-order approximation of the system of ODE in (2.1).

Timeline of measures

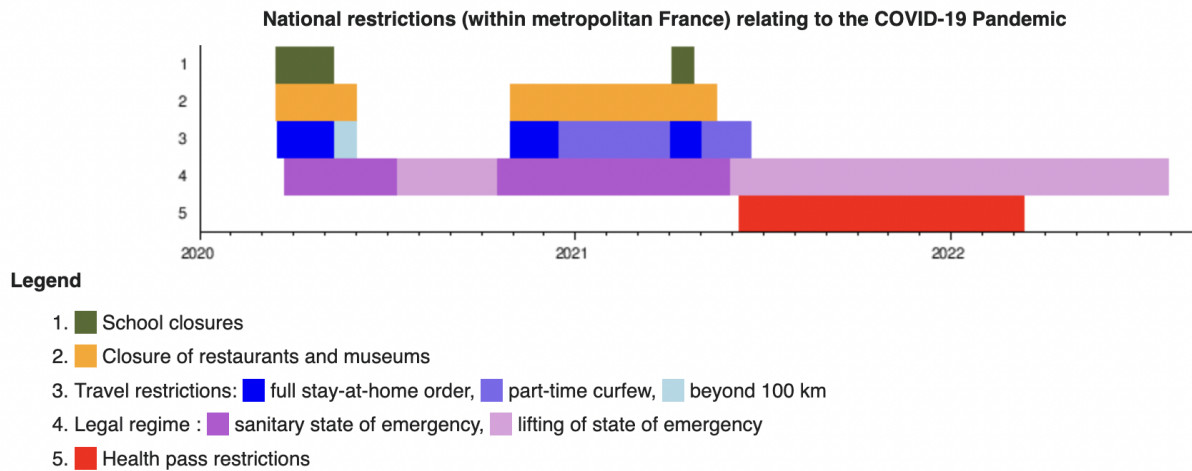


Figure 1. Timeline of COVID-19 restrictions in France.

3.2. The policy model

Policy function Instead of assuming continuous policy functions, we consider a more realistic set of policies with a discrete number of different stages and intensity levels. Therefore, policy functions form a subset of the admissible set $\mathcal{U}_{\alpha_{\max}, T_0}$ in [3].

We define the minimal policy time interval (MPTI) as the minimal duration time during which a policy remains constant or unchanged. This notion assumes that there is a minimal duration time of different stages of a policy. In addition to $u \in \mathcal{U}_{\alpha_{\max}, T_0}$, we assume that every policy u has a minimal policy time interval Δt and, in our simulations, the duration of each stage is a multiple of the MPTI. We denote this subset of policy functions as $\mathcal{U}_{\alpha_{\max}, T_0}^{\Delta t}$. In the past, many public health agencies enforced policies for time periods that corresponded to the work week (e.g., seven days) or multiples of this (e.g. one month). For the purpose of this paper, we assume the MPTI is an integer multiple of seven days.

We additionally assume that policy functions take values from a finite number of intensity levels $A \subset [\alpha_{\max}, 1]$, corresponding to different stages of the policy. In the simulations, we use $A = \{\alpha_{\max}, \frac{\alpha_{\max}+1}{2}, 1\}$. As a result, the policies we consider are piece-wise functions.

Cost function At time t , let $u(t) = \alpha$. The cost at time t is defined by:

$$c(\alpha) = \kappa c^{\text{implementation}}(\alpha) + \eta c^{\text{impact}}(\alpha) + (1 - \kappa - \eta) c^{\text{non-compliance}}(\alpha) \quad (3.2)$$

The cost function is a linear combination of three parts:

- (i) the implementation cost, which represents the consequences of policies meant to curtail the pandemic on individuals and the broader economic and social systems.
- (ii) the impact cost, which represents the consequences of people getting sick both on individuals and the broader economic and social systems.

- (iii) the non-compliance cost, which is a penalty imposed by a policy-maker upon an agent within its jurisdiction for deviating from its recommendation (e.g., a fine or litigation costs).

The implementation cost is a non-increasing function of α and the impact cost function is a non-decreasing function of α . The coefficients $\kappa, \eta, \kappa + \eta \in [0, 1]$. The cost from time t_1 to t_2 is defined as the averaged integral of the cost function over a total time period T :

$$c_{t_1 t_2}(u) = \frac{1}{T} \int_{t_1}^{t_2} c(\alpha(t)) dt. \quad (3.3)$$

There are different ways to parameterize the cost function. In this paper, the cost function is parameterized in the following way:

$$c_{t_1 t_2}(u) = \kappa \left(1 - \frac{\int_{t_1}^{t_2} u(t) dt}{T} \right) + \eta R_{t_2}(u) + (1 - \kappa - \eta) \frac{1}{T} \int_{t_1}^{t_2} (u(t) - \pi(u(t)))^2 dt, \quad (3.4)$$

where $R_{t_2}(u)$ is the fraction of the recovered population at time t_2 if policy u is adopted during $[t_1, t_2]$, and $\pi(u)$ is the policy of the agent one level above. The parameterization of the implementation cost and the non-compliance cost are adopted from [14]. The impact cost is parameterized as the recovered population at time t_2 to approximate the impact on the medical system since a fraction of the recovered represents the hospitalized population. If the cost function u is fixed at constant value α over time interval $[t_1, t_2]$, the cost can be written as:

$$c_{t_1 t_2}(u) = \kappa \frac{t_2 - t_1}{T} (1 - \alpha) + \eta R_{t_2}(\alpha) + (1 - \kappa - \eta) \frac{t_2 - t_1}{T} (\alpha - \pi(\alpha))^2. \quad (3.5)$$

An example of cost functions with different weights using the above parameterization is shown in Figure 6. In our simulation for a single region, we use a averaged total cost over a time period T as the following:

$$\begin{aligned} c_{\text{total}}(u) &= \frac{1}{T} \int_0^T c(u(t)) dt \\ &= \kappa \left(1 - \frac{\int_0^T u(t) dt}{T} \right) + \eta R_T(\alpha) + (1 - \kappa - \eta) \int_0^T (u(t) - \pi(u(t)))^2 dt. \end{aligned} \quad (3.6)$$

If at time T , the SIR model has reached the equilibrium, we can use $R_T(\alpha)$ to approximate R_∞ , the fraction of the final size of the recovered population. To find the optimal policy, we solve for the following optimization problem:

$$u(t) = \arg \min_{u'} c_{\text{total}}(u') \quad (3.7)$$

Algorithm 1 SINGLE-REGION POLICY SIR

```

1: Input: Time  $T$ , initial infected population  $I_0$ , initial susceptible population  $S_0$ , intensity
   levels  $A$ , minimal policy time interval  $\Delta t$ , policy end time  $T_0$ , Tol  $\epsilon$ 
2: Initialize county policies, minimal cost  $c_{\min} = 9999$ , current cost  $c = 0$ 
3:  $N = \frac{T}{\Delta t}$ ,  $n = 1$ 
4: Initialize policy  $u \in \mathbb{R}^N$ , optimal policy  $u_{\text{opt}} \in \mathbb{R}^N$ 
5: if  $n == N$  then
6:   for intensity level  $\alpha \in A$  do
7:     calculate  $S_{\text{final}}$  using the intensity level  $\alpha$ ,  $S_{N-1}$ ,  $I_{N-1}$ , and the update rule (3.1)
8:     if  $S_{\text{final}} > S_{\text{herd}} - \epsilon$  then
9:       calculate the cost  $c_{\text{temp}} = C(\alpha)$  for  $N$ -th time interval
10:       $c += c_{\text{temp}}$ ,  $u(N) = \alpha$ 
11:      if  $c \leq c_{\min}$  then
12:         $c_{\min} = c$ ,  $u_{\text{opt}} = u$ 
13:      end if
14:       $c -= c_{\text{temp}}$ 
15:    end if
16:  end for
17: else
18:  for intensity level  $\alpha \in A$  do
19:    calculate the cost  $c_{\text{temp}} = C(\alpha)$  for the  $n$ -th time interval
20:    calculate the susceptible  $S_n$  and the infected  $I_n$  at time  $n\Delta t$  using  $\alpha$ ,  $S_{n-1}$ ,  $I_{n-1}$  and the update
    rule 3.1
21:     $c += c_{\text{temp}}$ ,  $u(n) = \alpha^*$ 
22:     $n+ = 1$ 
23:    repeat line 5–22 until  $n = N$ 
24:     $n- = 1$ 
25:     $c -= c_{\text{temp}}$ 
26:  end for
27: end if
28: return  $c_{\min}$ ,  $u_{\text{opt}}$ 

```

3.3. Algorithm

We discretize time by MPTI Δt and the policy intensity into multiple levels. Let T be the total time and A be the set of possible policy intensities (e.g., $A = \{0, 0.5, 1\}$). We search for all the policies that lead to S_{final} being close to S_{herd} , i.e., $S_{\text{final}} > S_{\text{herd}} - \epsilon$, for some sufficiently small ϵ using a depth-first search algorithm [23]. The depth-first search algorithm stores the cost up to the current time interval and reuses this result to obtain the total cost for each policy function through backtracking. Let $N = \frac{T}{\Delta t}$ and N denote the number of stages of a policy. In total, there are $|A|^N$ policies. We initialize the minimal cost c_{\min} to be 9999. Assume that the initial susceptible and infected population are S_0 and I_0 ,

* $u(i)$ represents the i -th entry of vector u .

respectively. For n -th time interval ($n < N$), we choose a value from the set intensity levels A that has not been used before, calculate the cost for the policy intensity, add it to the previous cost, and calculate the susceptible and the infected at the end of n -th time interval using the chosen intensity. Then, we move to $(n + 1)$ -th time interval. If the end time interval is reached, we check if $S_{\text{final}} > S_{\text{herd}} - \epsilon$. If so, we calculate the cost for the final time interval and add it to the previous cost to get the current total cost c . If the total cost c is smaller than c_{min} , we update c_{min} with the total cost c , and the optimal policy u_{opt} with u . Next, we go back to the previous time interval and repeat the same procedure. After searching over all policies, the policy with the lowest cost is the optimal policy. The detailed algorithm is presented in Algorithm 1.

3.4. Simulations


In this section, we present the results for both single-region and multiple-region cases. We first compare the results of our discretized method of the COVID-19 in France with the results [3]. Next, we study the second wave (November 2020–May 2021) in Los Angeles County.

3.4.1. Optimal policy in France


We compare the results from [3] to our model with the same cost function but only three possible levels of policy intensity α . As in [3], the general cost function (3.6) reduces to the impact cost and is parameterized as the final epidemic size R_{∞} . Bliman et al assumed that the paths considered all reach herd immunity. Therefore, in our search for the optimal policy, we exclude cases that do not reach herd immunity. Note that without this exclusion, the optimal solution is to adopt and hold the strictest possible policy starting from the beginning of the pandemic. This results in the least number of infections. For ease of computation, we consider three levels of policy intensity: 0, 0.5, 1 and fixed time intervals for the MPTI. We use the same set of parameters for the SIR model as in Bliman et al [3]: $N = 6.7 \times 10^7$, $I_0 = 10^3$, $S_0 = N - I_0$, $\mathcal{R}_0 = 2.9$. Following [3], we also choose the policy end time T_0 as close as possible to 100, thus setting $T_0 = 98$ since the time interval needs to be a multiple of the MPTI of 7 days. We show the result our algorithm produces in Figure 2a, which we visually compare with the result from [3], shown in Figure 2b. Note that we normalized curves by the total population. Both solutions are bang-bang controllers. The solution using our model starts the control on day 63 (a multiple of 7) rather than day 61.9 (continuous). Slightly more people are infected under a policy that is forced to use seven day intervals compared with continuous time as used by Bliman et al.

Using a larger minimal policy time interval of 28 days and $T_0 = 112$, the optimal solution is no longer a bang-bang controller, as shown in Figure 2c with a larger $S_{\infty} = 0.32$. The optimal policy starts with a looser “intermediate” policy phase followed by a stricter phase. Interestingly, in practice, during COVID-19 it was common for policies to begin with the most stringent restrictions followed by partial opening [9, 27]. Thus, it is interesting to contrast the optimal policy with a policy in which the two stages are flipped in time; see Figure 2d. The flipped policy is a sub-optimal solution—it results in a larger pandemic size and a second wave of infections, as was often seen during the first two years of the COVID-19 pandemic. Nevertheless, the policy in Figure 2d, while infecting more people, divides the impacted population into two distinctive waves, which could decrease daily hospital demand over the course of the outbreak. Our policy model does not optimize for hospital demand. Since many public health agencies (including Los Angeles County) considered hospital demand when making


policy decisions, it could be important to consider that in future studies.




(a) Optimal policy with the minimal policy time interval $\Delta t = 7$ days, $S_\infty = 0.296$



(b) Optimal policy in [3], $S_\infty = 0.31$



(c) Optimal policy with the minimal policy time interval $\Delta t = 28$ days, $S_\infty = 0.32$



(d) Flipped policy from panel (c), $S_\infty = 0.174$

Figure 2. Optimal policy and the SIR model of France from March 17 to May 11, 2020. $S_{herd} \approx 0.345$.

Table 1. Parameters.

Figure	T_0	Δt	N	I_0	S_0	\mathcal{R}_0	$S_{s\infty}$
A	98	7	6.7×10^7	10^3	$N - I_0$	2.9	0.296
B	100	Not applicable	6.7×10^7	10^3	$N - I_0$	2.9	0.31
C	112	28	6.7×10^7	10^3	$N - I_0$	2.9	0.32
D	112	28	6.7×10^7	10^3	$N - I_0$	2.9	0.174

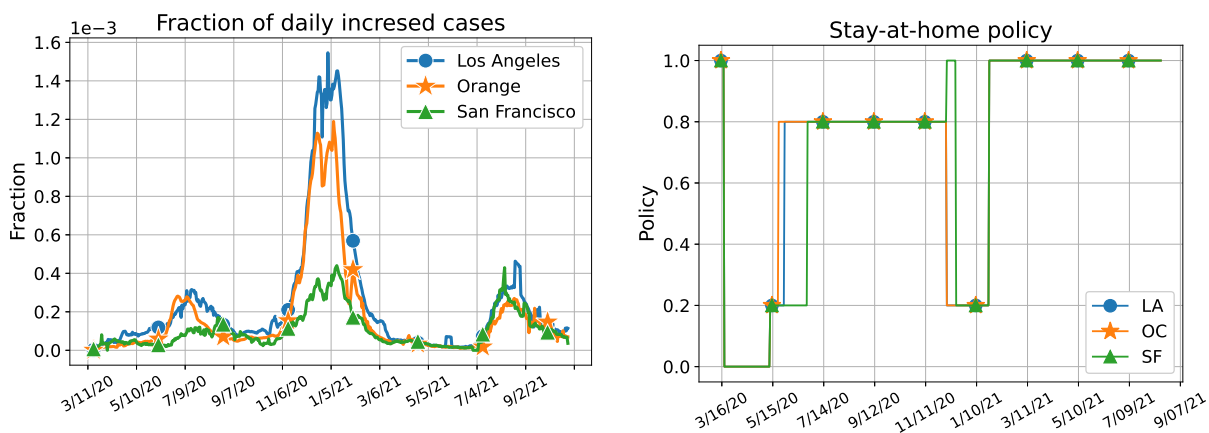
3.4.2. Case study—second wave in Los Angeles

We first present the course of infections in three counties in California and their corresponding stay-at-home policy changes from March 2020 to September 2021. Figure 3a shows the 7-day rolling average of the fraction of daily increased infected cases based on the data from [10] in three counties with the largest population density in California, namely, San Francisco, Orange, and Los Angeles.

There were three major outbreaks during the given time interval. For the first and second waves, Orange and Los Angeles followed similar trajectories, while San Francisco stayed more contained. Due to substantial holiday travel in winter 2020-21, the second wave was much larger than the first. In [8], the US Centers for Disease Control and Prevention described six levels of stay-at-home policy. The intensity of the policy decreases as the numerical value increases. The exact descriptions of the five levels of policies and their numerical representation are shown in Table 2. Figure 3b shows the change of the intensity of the stay-at-home policy during the same period. Policy during the first wave was *proactive*, whereas for the second wave it was more *reactive*. This may reflect some hesitancy on the part of policy-makers as well as lesser compliance by the population at large by the time the second wave emerged. During the second wave, with a relatively strict policy, the regions all stayed below herd immunity. With vaccination available in early 2021, the pandemic in all three regions tapered off.

Table 2. CDC stay-at-home policies. There are six levels of policies and we map the levels linearly onto the interval $[0, 1]$ for simplicity. The numerical value on the left is used to graph actual policies over time in Figure 3b.

Numerical value	Stay-at-home policy
0	Mandatory for all individuals
0.2	Mandatory only for all individuals in certain areas of the jurisdiction
0.4	Mandatory only for at-risk individuals in the jurisdiction
0.6	Mandatory only for at-risk individuals in certain areas of the jurisdiction
0.8	Advisory/Recommendation
1	No order for individuals to stay home



(a) The fraction of the daily increase of the infected with a 7-day rolling average.

(b) Stay-at-home policy

Figure 3. The fraction of the infected and stay-at-home policy over time in Los Angeles, San Francisco, and Orange County.

Now, we consider a counterfactual study of how the pandemic would have evolved had herd immunity been reached during the second wave, controlled by our policy model, using parameters measured from the Los Angeles data. We choose to study the period of the second wave for several reasons. First, the data reporting scheme improved for the second wave compared to the first wave. In

addition, with the experience and knowledge gained from the first wave, authorities were in a better position to make optimal decisions. Given that there was no complete lockdown during the second wave, we consider the policy intensity levels $A = \{0.2, 0.6, 1\}$, and use the minimal policy time interval $\Delta t = 7$. We choose 0.2 as our maximal policy intensity, because full lockdown was not desirable during this period. We choose a second policy level of 0.6 as the midpoint between 0.2 and 1. In all simulations, we optimize for final pandemic size and compare the optimal controls found.

In Figure 4, the left column (Figures 4a, 4c, 4e) is the simulated SIR with the optimal policy when the basic reproduction number $R_0 = 2.5$ and the initial recovered $r_0 = 0.1, 0.2, 0.3$. The right column (Figures 4b, 4d, 4f) is the simulated SIR with the optimal policy when the reproduction number $R_0 = 2.15$ and the initial recovered $r_0 = 0.1, 0.2, 0.3$. This value $R_0 = 2.5$ is estimated from the early COVID-19 infected data (Jan 22—Mar 15, 2020 [10]) and $R_0 = 2.15$ is estimated using the infected data from September 16 to November 15, 2021 ([10]), prior to the second wave. All optimal policies have a bang-bang-like shape. The policy started approximately around the peak of the infected curve, and the resulting dynamics approach herd immunity. For larger values of r_0 , we expect that a shorter period of high-intensity policy is needed to reach herd immunity; our results confirm this. Once enough of the population is infected and recovered, a shorter control policy is needed.

4. Multi-layer multi-regional case

In this section, we present a multi-regional model with multiple policy-making layers, extending the model proposed by Jia et al., [14] to consider a dynamic epidemic model and the control policies discussed above (see Section 3). Specifically, we propose a game-theoretic model in which regions are combined into layers, with the top layer corresponding to the highest-level decision maker (e.g., a federal government), the next layer comprised of the next level of decision making (e.g., states or provinces), and so on (see Figure 5). The top decision maker chooses the policy first, then all the decision makers in the next layer do so simultaneously, and so on. Additionally, we consider a special case in which there are multiple decision makers (e.g., states, counties, etc.) choosing their epidemic control policies simultaneously in one layer. We use a form of hierarchical best response dynamics to compute approximate equilibria in this multi-layer game [14], performing this computation independently for each time interval (essentially assuming that the players do not reason explicitly about future dynamics when making instantaneous policy decisions at a particular point in time).

The multi-region case naturally has a competition between regions to optimize their strategy with respect to the choices made by other regions. For this reason, the single-region model does not directly extend. There are two main differences between our work and that of Jia et al. [14]. First, their model is based on Nash equilibrium, where agents make decisions with other agents' possible actions in mind. We use the idea of in-game learning [11]. We assume that the agents gradually evolve toward the best decisions instead of being optimal instantly. In practice, each region in the game assumes that other regions' policies (at the same level) stay the same when optimizing its own cost function. Second, we focus on the dynamics, instead of snapshot in time considered by Jia et al.

Network SIR In practice, counties can hardly be treated as independent. People travel across county borders to work and socialize. The majority of the literature of network-style SIR models focus on the individuals as nodes and study the effects of interpersonal network on the pandemics [16, 17, 20]. For

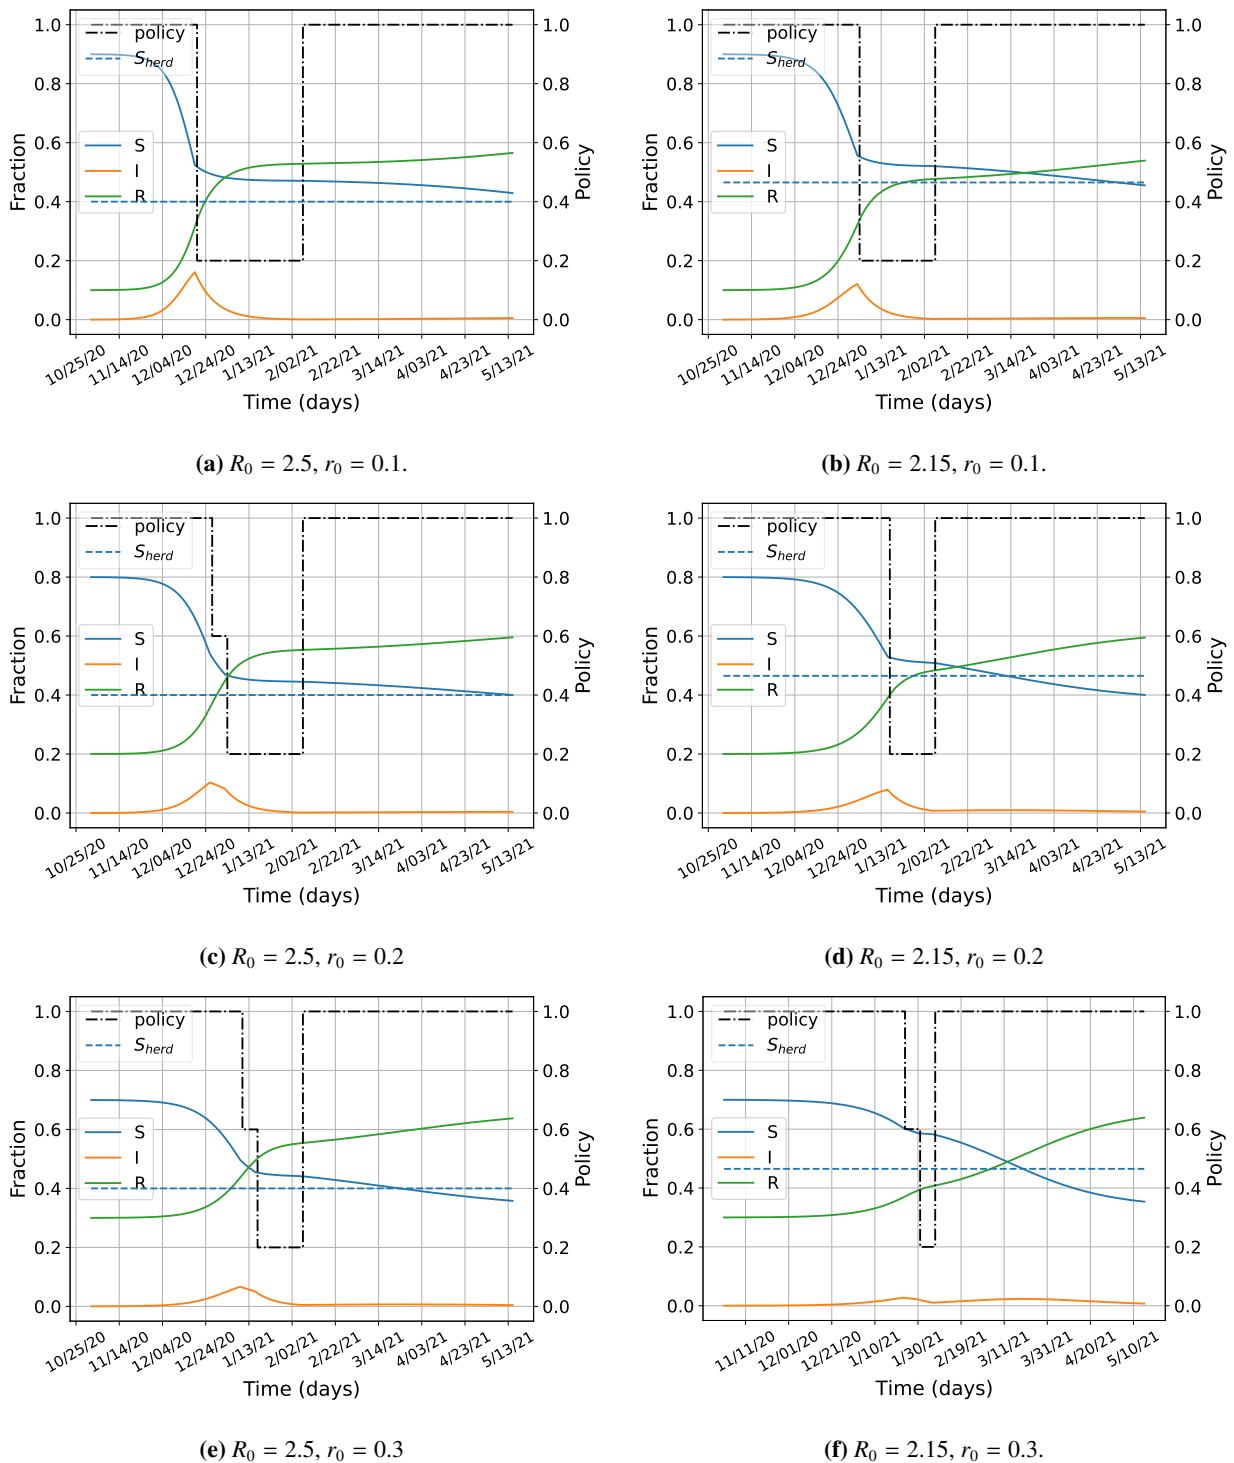


Figure 4. Optimal policy in Los Angeles with the basic reproduction number $R_0 = 2.5, 2.15$, $S_{\text{herd}} \approx 0.4, 0.465$, and the fraction of the initial recovered population $r_0 = 0.1, 0.2, 0.3$, respectively.

example, B. Macdonald et al [20] empirically studied how well various centrality measures perform

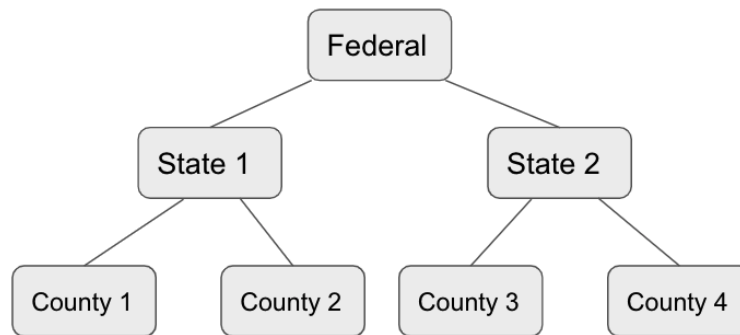


Figure 5. An example of a three-layer hierarchical structure.

at identifying which individuals in a network will be the best spreaders of disease. In [25], authors explains why most COVID-19 infection curves are linear after the first peak in the context of the contact network using a network SIR model. There are a few works that study the interplay between different geographical regions rather than the interpersonal contact network. In [12], a kernel-modulated SIR model was introduced to model the spread of COVID-19 across counties. The kernel is based on the spatial proximity between regions. Metapopulation epidemic models are based on the spatial structure of the environment and detailed knowledge of transportation infrastructure and movement patterns. The metapopulation dynamics of infectious diseases has generated a wealth of models and results using mechanistic approaches taking explicitly into account the movement of individuals ([13, 15, 22]). For example, in [22], the authors proposed a multi-regional compartmental model using medical geography theory (central place theory) and studied the effect of the traveling of individuals (especially those infected and exposed) between regions on the global spread of severe acute respiratory syndrome (SARS). Another way to account for the interplay between regions is to use a cross excitation matrix [28]. This scheme assumes a uniform mixing of the population across regions that and the infected population in one region can trigger the infection in another. The entries of the matrix record the pair-wise cross excitation from one region to another. In this paper, we assume uniform mixing in the population and use an excitation matrix $K = \{K_{aa'}\}$ to model the travel and infections across counties. Our network-style SIR is the following:

$$\left\{ \begin{array}{l} \frac{dS_a(t)}{dt} = -\alpha_a \beta \sum_{a'} K_{aa'} \frac{I_{a'}(t) S_a(t)}{N_a}, \\ \frac{dI_a(t)}{dt} = \alpha_a \beta \sum_{a'} K_{aa'} \frac{I_{a'}(t) S_a(t)}{N_a} - \gamma I_a(t), \\ \frac{dR_a(t)}{dt} = \gamma I_a(t), \\ S(0) = S_0, \quad I(0) = I_0, \quad R(0) = R_0. \end{array} \right. \quad (4.1)$$

For any county a , the rate of change from S_a to I_a triggered by $I_{a'}$ depends on $K_{aa'}$, the current fraction of the susceptible S_a in county a and the current fraction of the infected $I_{a'}$ in county a' . Note that $K_{aa} = 1$. When $K = I$, the network SIR is the independent SIR.




Figure 6. Different cost functions versus policy intensity α .

Cost function Consider the i -th time interval $[i\Delta t, (i+1)\Delta t]$ and $u(t) = \alpha$ for $t \in [i\Delta t, (i+1)\Delta t]$. A region a adopts the following cost function:

$$c_{i\Delta t, (i+1)\Delta t}^a(\alpha) = \kappa_a(1 - \alpha)\Delta t/T + \eta_a R_{a,T}(\alpha) + (1 - \kappa_a - \eta_a)(\alpha - \pi(\alpha))^2 \Delta t/T,$$

where $R_{a,T}(\alpha)$ is the epidemic size of region a at time T . For a top-layer region f , there is no non-compliance cost and the cost function is

$$c_{\Delta t}^f(\alpha) = \kappa_f(1 - \alpha)\Delta t/T + \eta_f R_{f,T}(\alpha),$$

where $\kappa_f + \eta_f = 1$ and $R_{f,T}(\alpha)$ is the number of the recovered of region f , which is an aggregation of the epidemic size of its leaf nodes.

4.1. Algorithms

The single-region algorithm minimizes over all admissible piece-wise functions, while the multiple-region algorithm only minimizes over every time interval. We assume there are up to three layers: federal government, the states, and the counties. At n -th time interval, we first determine the optimal policy intensity that minimizes the cost $C_{n\Delta t, (n+1)\Delta t}^f$ for the federal layer. After obtaining the optimal federal policy, each state optimizes its own cost function $C_{n\Delta t, (n+1)\Delta t}^s$ for the period $[n\Delta t, (n+1)\Delta t]$ unilaterally, i.e., assuming other states follow their previous policies. Next, we choose the optimal policy intensity for the counties in the same manner. Note that the federal layer does not pay the non-compliance cost as it is not subject to any higher-level policy making. The states and counties may pay a non-compliance cost. The full details of the three-layer model is in Algorithm 2.

4.2. Simulations

In this section, we present results for a three-county example of the multiple-regions game and a three-county example with a state. First, we discuss when one layer exists (i.e., only counties). The

Algorithm 2 GAME POLICY SIR

```

1: Input: Time  $T$ , excitation matrix  $K$ , intensity levels  $A$ , time interval  $\Delta t$ 
2: Initialize state, county policies
3: Number of policy stages  $N = \frac{T}{\Delta t}$ ,  $n = 1$ 
4: while  $n \leq N$  do
5:    $t = n\Delta t$ 
6:   while  $t < T$  do
7:     for every state  $s$  do
8:       for every county  $a$  in state  $s$  do
9:         update  $S_a, I_a, R_a$  according to the current policy  $\alpha_a$  and the excitation matrix  $K$ :
10:         $S_a(t) = S_a(t-1) - \alpha_a \beta \sum_{a'} K_{aa'} \frac{I_{a'}(t-1)S_a(t-1)}{N_a}$ 
11:         $I_a(t) = I_a(t-1) + \alpha_a \beta \sum_{a'} K_{aa'} \frac{I_{a'}(t-1)S_a(t-1)}{N_a} - \gamma I_a(t-1)$ 
12:         $R_a(t) = R_a(t-1) + \gamma I_a(t-1)$ 
13:       end for
14:     end for
15:      $t += 1$ 
16:   end while
17:    $\alpha_f = \arg \min_{\alpha' \in A} C_{n\Delta t, (n+1)\Delta t}^f(\alpha')$ 
18:   for every state  $s$  do
19:      $\alpha_s = \arg \min_{\alpha' \in A} C_{n\Delta t, (n+1)\Delta t}^s(\alpha')$ 
20:     for every county  $a$  in state  $s$  do
21:        $\alpha_a = \arg \min_{\alpha' \in A} C_{n\Delta t, (n+1)\Delta t}^a(\alpha')$ 
22:     end for
23:   end for
24:    $n += 1$ 
25: end while

```

game between the counties is through cross excitation of infection among the counties. Next, we study the case when a governing state is added.

We consider three interacting counties with the excitation matrix K :

$$K = \begin{bmatrix} 1 & 0 & 0 \\ 0.1 & 1 & 0 \\ 0 & 0.1 & 1 \end{bmatrix}$$

We set the reproduction number $\mathcal{R}_0 = 2$ and therefore, $S_{\text{herd}} = 0.5$. Counties 1, 2, and 3 have initial fractions of the infected population as $i_0 = 0.2, 0.1, 0.1$, respectively. This implies that county 1 has a bigger outbreak initially, and part of the infection in county 2 is excited from county 1 and part of the infection in county 3 is excited from county 2. The cost functions for all counties consist of an implementation cost and an impact cost with equal weights ($\eta_a = \kappa_a = 1/2$, for all a). The minimal policy time interval Δ is set to be 7 (days).

The left column (Figures 7a, 7c, 7e) are simulations for the counties without any intervention, and the right column (Figures 7b, 7d, 7f) are simulations with interventions. Without intervention, we see propagation of waves of infection from county 1 to county 2 and then to county 3. All of the counties

reach herd immunity eventually. With interventions, policy restrictions started on day 7 and, for county 2 and 3, the infected curves decrease before reaching their peaks. With control, county 1 contained the pandemic and the final S_∞ is close to herd immunity level S_{herd} . With a smaller infected population to begin with, county 2 and 3 contained the pandemic before reaching herd immunity.

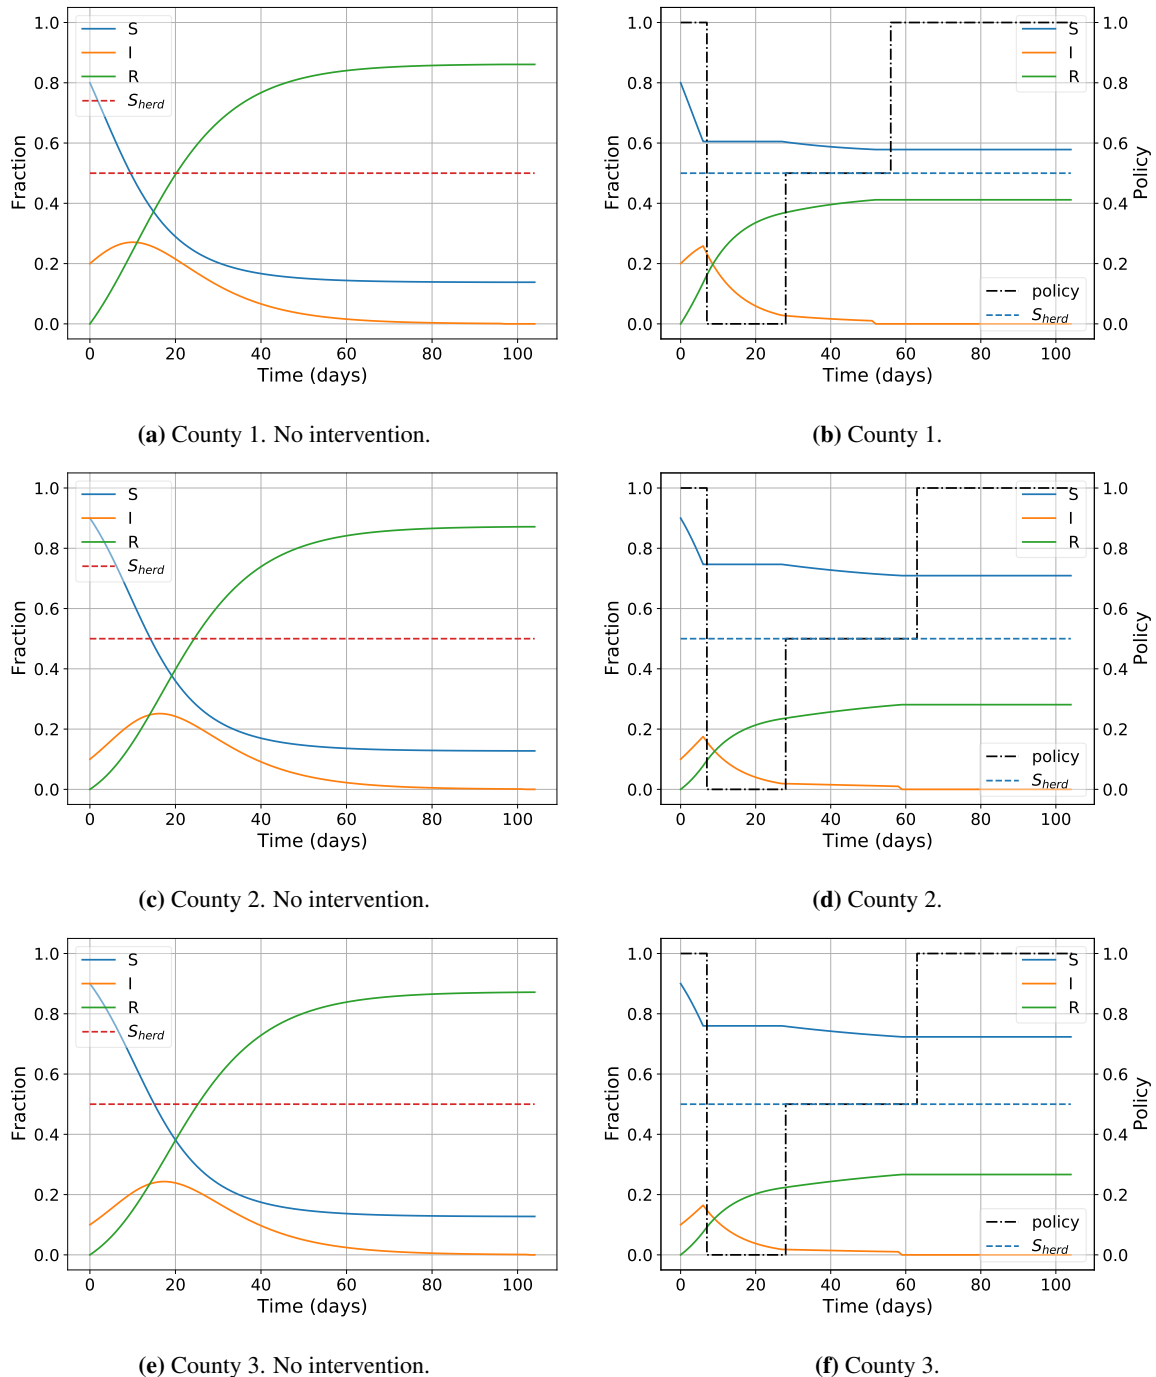


Figure 7. An example of three dependent counties without and with interventions. With intervention, for all counties, the coefficients for the implementation cost $\kappa = \frac{1}{2}$ and the coefficients for the impact cost $\eta = \frac{1}{2}$. The minimal policy time interval $\Delta t = 7$.

Figure 8 shows the results of adding a governing state on top of the county layer. We keep the ratio of the weights for the implementation cost and the impact cost to be 1:1, the same as in the no-state case in Figure 7. The state has slightly different weights, with the ratio of the weights for the implementation cost and the impact cost being 1:2. Compared to Figure 7, by adding a state, the three counties ended up with the same policy. In this case, the non-compliance cost results in each county choosing the same policy as the state rather than different policies.

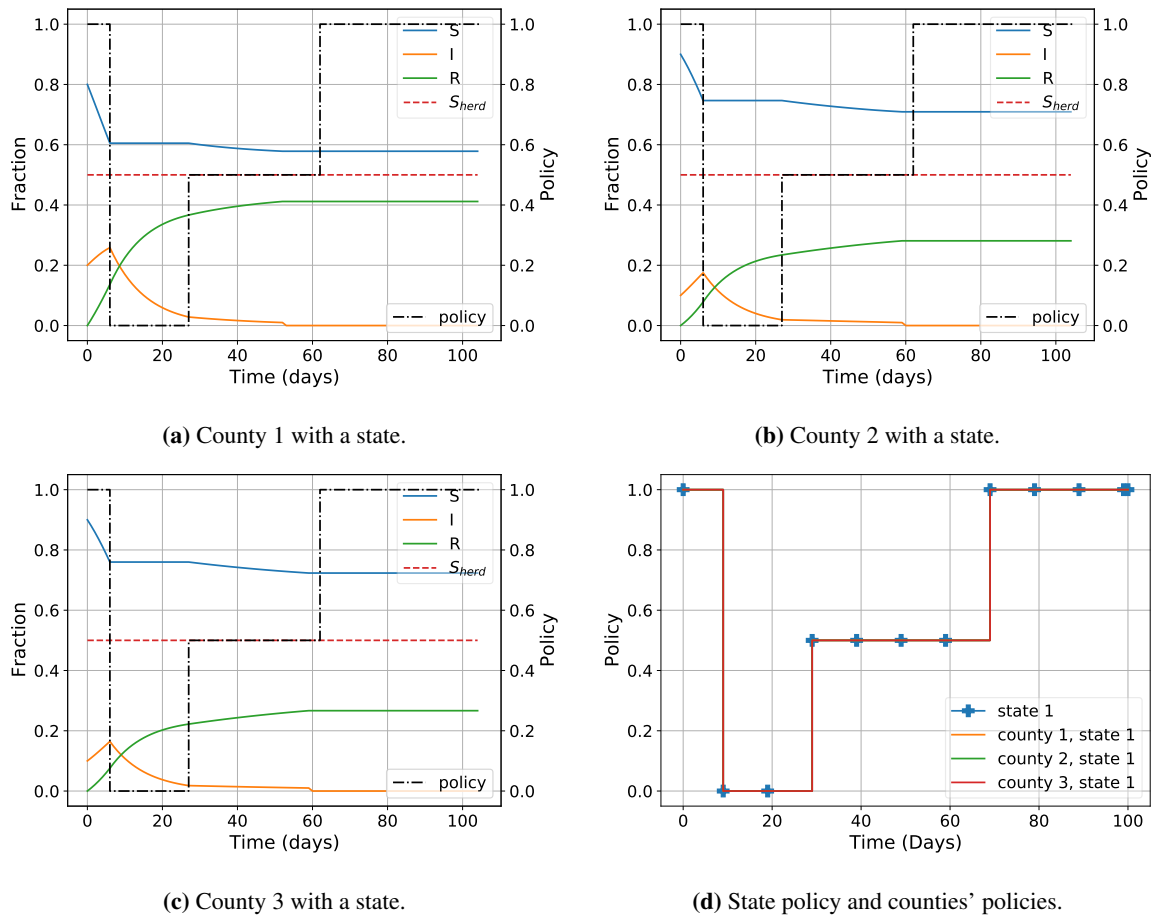


Figure 8. An example with three counties and a governing state. For all counties, the coefficients for the implementation cost $\kappa = \frac{1}{6}$, the coefficients for the impact cost $\eta = \frac{1}{6}$, and the coefficients for the non-compliance cost $1 - \kappa - \eta = \frac{2}{3}$. For the state, the coefficients for the implementation cost $\kappa = \frac{1}{3}$, the coefficients for the impact cost $\eta = \frac{2}{3}$. The minimal policy time interval $\Delta t = 7$.

5. Discussion and future work

We propose a policy-making model coupled with the SIR model to study a single region and game-like interactions between multiple regions. The model demonstrates its ability to model real-life situations with different sets of parameters in both one-region and multiple-region scenarios. One can extend the model to a hierarchical structure by building multiple layers of the multiple regions model

and study the cross-layer effects.

In the search for an optimal policy, we used a naive depth-first search algorithm for the one-region model. One can speed up the algorithm by removing some of the obvious non-optimal paths.

In our model, the policy intensity α is a heuristic representation of the lockdown, social distancing, and mask policy. It remains to be discussed how other policies, for example, vaccination policies, affect the spreading in the different stages of the pandemic. The model ignores some important features like the limitation of hospital capacity [24], which could be added as constraints when minimizing the cost function. Figure 3b shows that the policy for the first wave is proactive while the one for the second wave is reactive. One possible effect is from fatigue of following policies, which increases in time and has a memory. So far, the model does not have the capability of modeling this fatigue. In the future, one could consider an adaptive term in the cost function to model it. The network example considered was rather simplistic, with just three counties within one state. One could consider more complex systems with multiple layers. The computational method used here would likely need to be improved to address the computational complexity of the search space. In addition, a potentially important generalization is to capture implementation and impact costs with more refined cross-layer dependencies, but this is potentially non-trivial from both a modeling and computational perspective.

Author contributions

All authors contributed conceptualization, writing, review, and editing. Xia and Andrea worked on the methodology and wrote the original draft. Andrea, Jeffrey, and Yevgeniy supervised the project and acquired the funding, and Xia implemented the experiments.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of Interest

The authors declare there is no conflict of interest.

Acknowledgement

This work is supported by ARO MURI grant W911NF1810208, National Science Foundation grants DMS-2027277 and DMS-2318817 through the Algorithms for Threat Detection (ATD) program, and a Simons Foundation Math + X Investigator Award # 510776, and NSF (IIS-1905558, IIS-2214141, CNS-2310470). The code can be found: <https://github.com/xli07/PolicySIR/tree/main>.

References

1. LA County Department of Public Health, COVID-19: Keeping Los Angeles safe. Los Angeles Mayor's Office, 2020. Available from: <https://coronavirus.lacity.gov/>.
2. P. A. Bliman, M. Duprez, How best can finite-time social distancing reduce epidemic final size? *J. Theor. Biol.*, **511** (2021), 110557. <https://doi.org/10.1016/j.jtbi.2020.110557>

3. P. A. Bliman, M. Duprez, Y. Privat, N. Vauchelet, Optimal immunity control and final size minimization by social distancing for the SIR epidemic model, *J Optim Theory Appl*, **189** (2021), 408–436. <https://doi.org/10.1007/s10957-021-01830-1>
4. S. Blower, K. Koelle, J. Mills, Health policy modeling: Epidemic control, HIV vaccines, and risky behavior, *Quantitative evaluation of HIV prevention programs*, Yale University Press, New Haven, 2008, 260–289.
5. M. C. J. Bootsma, N. M. Ferguson, The effect of public health measures on the 1918 influenza pandemic in U.S. cities, *Proc. Natl. Acad. Sci.*, **104** (2007), 7588–7593. <https://doi.org/10.1073/pnas.0611071104>
6. D. S. Burke, J. M. Epstein, D. A. Cummings, J. I. Parker, K. C. Cline, R. M. Singa, et al., Individual-based computational modeling of smallpox epidemic control strategies, *Acad Emerg Med*, **13** (2006), 1142–1149. <https://doi.org/10.1197/j.aem.2006.07.017>
7. Centers for Disease Control and Prevention, CDC Museum COVID-19 Timeline. Centers for Disease Control and Prevention, 2022. Available from: <https://www.cdc.gov/museum/timeline/covid19.html>.
8. Centers for Disease Control and Prevention, U.S. State and Territorial Stay-At-Home Orders: March 15, 2020–August 15, 2021 by County by Day. Environmental Public Health Tracking, 2022. Available from: <https://data.cdc.gov/Policy-Surveillance/U-S-State-and-Territorial-Stay-At-Home-Orders-Marc/y2iy-8irm>.
9. Department of Public Health, Los Angeles, Department of Public Health, Los Angeles, 2022. Available from: <http://publichealth.lacounty.gov/>.
10. E. Dong, H. Du, L. Gardner, An interactive web-based dashboard to track COVID-19 in real time, *Lancet Infect. Dis.*, **20** (2020), 533–534. [https://doi.org/10.1016/S1473-3099\(20\)30120-1](https://doi.org/10.1016/S1473-3099(20)30120-1)
11. D. Fudenberg, D. K. Levine, *The Theory of Learning in Games*, Cambridge: The MIT Press, 1998.
12. X. Geng, G. G. Katul, F. Gerges, E. Bou-Zeid, H. Nassif, M. C. Boufadel, A kernel-modulated SIR model for COVID-19 contagious spread from county to continent, *Proc. Natl. Acad. Sci.*, **118** (2021), e2023321118. <https://doi.org/10.1073/pnas.2023321118>
13. R. F. Grais, J. H. Ellis, G. E. Glass, Assessing the impact of airline travel on the geographic spread of pandemic influenza, *Eur. J. Epidemiol.*, **18** (2003), 1065–1072. <https://doi.org/10.1023/A:1026140019146>
14. F. Jia, A. Mate, Z. Li, S. Jabbari, M. Chakraborty, M. Tambe, et al., *A game-theoretic approach for hierarchical policy-making*, arXiv: 2102.10646, [Preprint], (2012) [cited 2024 September 02]. Available from: <http://161.97.88.200/studii/Studiu%20vaccin%20sars%20cov2/studii%20efecte%20adverse%20vaccin/2102.10646.pdf>.
15. M. J. Keeling, P. Rohani, Estimating spatial coupling in epidemiological systems: a mechanistic approach, *Ecol. Lett.*, **5** (2002), 20–29. <https://doi.org/10.1046/j.1461-0248.2002.00268.x>
16. I. Z. Kiss, J. C. Miller, P. Simon, *Mathematics of Epidemics on Networks: from Exact to Approximate Models*, Schweiz: Springer International Publishing, 2017.

17. P. Kunwar, O. Markovichenko, M. Chyba, Y. Mileyko, A. Koniges, T. Lee, A study of computational and conceptual complexities of compartment and agent based models, *Netw. Heterog. Media*, **17** (2022), 359–384. <https://doi.org/10.3934/nhm.2022011>
18. T. H. Lee, B. Do, L. Dantzing, J. Holmes, M. Chyba, S. Hankins, et al., Mitigation planning and policies informed by COVID-19 modeling: A framework and case study of the state of Hawaii, *Int. J. Environ. Res. Public Health*, **19** (2022), 6119. <https://doi.org/10.3390/ijerph19106119>
19. R. G. Lin II, L.A. county now requires residents to wear face coverings. Los Angeles Times, 2020. Available from: <https://www.latimes.com/california/story/2020-04-16/l-a-county-now-requires-residents-to-wear-face-coverings-here-are-the-details>.
20. B. Macdonald, P. Shakarian, N. Howard, G. Moores, *Spreaders in the network SIR model: an empirical study*, arXiv:1208.4269, [Preprint], (2012) [cited 2024 August 02]. Available from: <https://arxiv.org/abs/1208.4269>.
21. A. L. Pitt, K. Humphreys, M. L. Brandeau, Modeling health benefits and harms of public policy responses to the us opioid epidemic, *Am. J. Public Health*, **108** (2018), 1394–1400. <https://doi.org/10.2105/AJPH.2018.304590>
22. S. Ruan, W. Wang, S. A. Levin, The effect of global travel on the spread of SARS, *Math Biosci Eng*, **3** (2006), 205–218. <https://doi.org/10.3934/mbe.2006.3.205>
23. R. Tarjan, Depth-first search and linear graph algorithms, *12th Annual Symposium on Switching and Automata Theory (swat 1971)*, IEEE Computer Society, 1971, 114–121.
24. R. Thompson, C. Gilligan, N. Cunniffe, Will an outbreak exceed available resources for control? Estimating the risk from invading pathogens using practical definitions of a severe epidemic, *J R Soc Interface*, **17** (2020), 20200690. <https://doi.org/10.1098/rsif.2020.0690>
25. S. Thurner, P. Klimek, R. Hanel, A network-based explanation of why most COVID-19 infection curves are linear, *Proc. Natl. Acad. Sci.*, **117** (2020), 22684–22689. <https://doi.org/10.1073/pnas.2010398117>
26. Wikipedia, California COVID-19 Timeline. Available from: https://en.wikipedia.org/wiki/Timeline_of_the_COVID-19_pandemic_in_California.
27. Wikipedia, COVID-19 pandemic in France. Available from: https://en.wikipedia.org/wiki/COVID-19_pandemic_in_France#Timeline_of_measures.
28. B. Yuan, H. Li, A. L. Bertozzi, P. J. Brantingham and M. A. Porter, Multivariate spatiotemporal Hawkes processes and network reconstruction, *SIAM J. Math. Data Sci.*, **1** (2019), 356–382. <https://doi.org/10.1137/18M1226993>



AIMS Press

©2024 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)