CoSIR: Managing an Epidemic via Optimal Adaptive Control of Transmission Rate Policy

Harsh Maheshwari*
Shreyas Shetty*
harsh.maheshwari@flipkart.com
shreyas.shetty@flipkart.com
Flipkart Internet Private Limited
Bangalore, India

Nayana Bannur
Wadhwani AI
Mumbai, India
nayana.bannur@wadhwaniai.org

Srujana Merugu
Independent
Bangalore, India
srujana@gmail.com

ABSTRACT

Shaping an epidemic with an adaptive contact restriction policy that balances the disease and socioeconomic impact has been the holy grail during the COVID-19 pandemic. Most of the existing work on epidemiological models focuses on scenario-based forecasting via simulation but techniques for explicit control of epidemics via an analytical framework are largely missing. In this paper, we consider the problem of determining the optimal control policy for transmission rate assuming SIR dynamics, which is the most widely used epidemiological paradigm. We first demonstrate that the SIR model with infectious patients and susceptible contacts (i.e., product of transmission rate and susceptible population) interpreted as predators and prey respectively reduces to a Lotka-Volterra (LV) predator-prey model. The modified SIR system (LVSIR) has a stable equilibrium point, an "energy" conservation property, and exhibits bounded cyclic behaviour similar to an LV system. This mapping permits a theoretical analysis of the control problem supporting some of the recent simulation-based studies that point to the benefits of periodic interventions. We use a control-Lyapunov approach to design adaptive control policies (CoSIR) to nudge the SIR model to the desired equilibrium that permits ready extensions to richer compartmental models. We also describe a practical implementation of this transmission control method by approximating the ideal control with a finite, but a time-varying set of restriction levels. We provide experimental results comparing with periodic lockdowns on few different geographical regions (India, Mexico, Netherlands) to demonstrate the efficacy of this approach.

CCS CONCEPTS

• **Applied computing** → *Life and medical sciences*; • **Computing methodologies** → **Control methods**.

KEYWORDS

SIR, Lotka-Volterra system, Optimal control, Epidemic, Contact Restrictions, COVID-19, Transmission rate

1 INTRODUCTION

The COVID-19 situation and its immense toll on human lives has highlighted the enormous public health challenges associated with managing a pandemic. In the absence of a vaccine, there are primarily three control levers available for public health officials, namely, (a) contact restrictions, (b) testing, tracing and isolation, and (c) provisioning for additional medical capacity. Of these, contact restrictions via lockdowns and social distancing have emerged as the most powerful policy instrument especially in low and middle income countries that cannot afford to scale up testing or medical capacity. Choosing the optimal level of restrictions, however, has been highly non-trivial not only because it involves a complex trade-off between the yet to be understood COVID-19 disease impact and other socioeconomic disruptions, but also because of the rapidly evolving situation on the ground.

Public health interventions related to the COVID-19 pandemic have largely been driven by scenario-based epidemiological forecasting studies [22, 27]. Current epidemiological models [11, 15, 54] incorporate spatiotemporal variations and predictive signals such as mobility to provide high fidelity forecasts. However, the decision making on contact restrictions is still fairly sub-optimal as it is based on the comparison of a few enumerated scenarios for a limited time horizon. Furthermore, forecasts based on a constant transmission rate (i.e., avg. number of new infections from an infectious person per time unit) convey the impression that the epidemic progression corresponds to a bell curve [54] regardless of empirical evidence to the contrary (see Figure 1). Flattening the curve till herd immunity is seen as the only choice. Epidemiological analysis is often centred around determining the height and timing of the caseload peak as well as the time to attain herd immunity. Though highly valuable, this scenario-based decision-making approach leans towards a limited reactive role for public health agencies.

In contrast, despite the potentially far-reaching impact, relatively less attention has been devoted to developing a mathematical control framework to support proactive decision-making based on the target disease & economic outcomes, and the state of the epidemic. Multiple studies [18, 21, 36, 56] point to the benefits of periodic lockdowns and staggered mobility among population groups, but these dynamic interventions are based on forecast simulations of a limited number of scenarios and are not adaptive in nature. Some recent works [7, 49] explore optimal control policies but the primary focus is on the modeling of the socioeconomic and disease impact.

In this paper, we explore the problem of optimal adaptive control of transmission rate for a desired bound on infectious population. We focus on epidemiological models based on compartmental (SIR and SEIR) dynamics [35] because of their wide applicability, parsimonious & interpretable encoding of the disease dynamics, and amenability for data-driven calibration to yield accurate forecasts.

*Both authors contributed equally to this research.

Contributions:

- We demonstrate that the SIR dynamics map to the well-known Lotka-Volterra (LV) system [12] on interpreting infectious patients as predators and susceptible contacts (i.e., the product of transmission rate and susceptible population) as the prey under specific conditions on the transmission rate. The resulting system (LVSIR) has a well-defined stable equilibrium point and an "energy" conservation property. It exhibits a bounded cyclic trend for active infections and a steady decline of the susceptible population.
- We derive optimal control policy for transmission rate (CoSIR) using control-Lyapunov functions [61] based on the energy of the system, that is guaranteed to converge to the desired equilibrium, i.e., target infectious levels from any valid initial state. We also discuss extensions to compartmental model variants that involve an incubation period (e.g., delayed SIR, SEIR) as well as control of the infectious period that is influenced by testing and quarantine policy.
- We propose a practical approximate implementation (CoSIR-approx) of the transmission rate control via discrete, but time-varying restriction levels. Simulation results demonstrate the efficacy of this CoSIR-based approach in stabilizing infections and adaptability to perturbations [1].

The rest of the paper is organized as follows. Section 2 presents a formulation of the restriction control problem. Section 3 provides background on compartmental models, LV systems and relevant aspects of control theory. Sections 4, 5, 6, 7, 8 present the SIR to LV system mapping, the transmission rate control mechanism, practical restriction control policy, simulation results, and extensions respectively. Section 9 presents the concluding remarks. Figure 2 lists the different models of interest for clarity.

Notation: x_t and $x(t)$ interchangeably denote the value of a variable x at time t. Time derivative of x is denoted by $\dot{x} = \frac{dx}{dt}$. $[x_t],\ [t]_{t_1}^{t_2}$ denotes the series of x as t varies from t_1 to t_2.

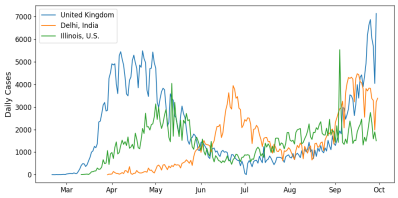

Figure 1: Daily reported infections in various regions of the world where more than one peak has been observed.

2 PROBLEM FORMULATION

During an epidemic, a key concern for public health officials is to determine the right level and schedule of contact restrictions that balances the disease and socioeconomic burdens. Strict lockdown conditions for a short time period suppress the infection levels, but infections tend to flare up again on easing restrictions unless the epidemic is completely wiped out. On the other hand, prolonged restrictions with no intermittent easing hinder economic activity and impose heavy costs on vulnerable population groups. Furthermore, progressive reduction in the susceptible population offers a chance for relaxation of restrictions that needs to be exploited.

[1]http://cosir.herokuapp.com/, https://github.com/dsindiavscovid/CoSIR.

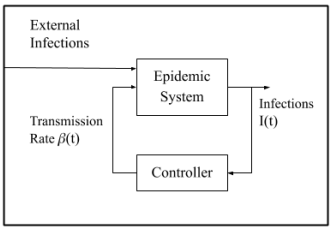

Model	Transmission rate (β)
SIR	constant
Lotka Volterra- SIR (LVSIR)	Varies with time following Eqn. 2.
Controlled SIR (CoSIR)	Varies with time following control mechanism in Eqn. 3.
Approximate Control (CoSIR-approx)	β from CoSIR is approximated by k discrete levels

Figure 2: Proposed SIR model variants and transmission rate feedback control.

Modelling the multi-faceted impact of contact restrictions requires accounting for region-specific cultural and economic constructs as well as the available medical capacity, a highly complex task. For tractability, we assume that the public health goal is to limit active infections to a certain target level determined via an independent impact analysis [3]. The controls available to the public health authorities can be viewed as multiple knobs that can be set to different levels (e.g., public transport at 50% occupancy). However, the need to communicate the policy to the general public and ensure compliance entails a simpler strategy centred around a few discrete restriction levels [1, 6], (e.g. Table 1) and a preset schedule for a future time horizon, which is often longer than the intervals at which the epidemic observations are collected. For example, the infection levels might be monitored at a daily frequency, but the restriction guidelines (e.g., Level 2 on weekdays and Level 1 on weekends) might be chosen for a monthly period.

Restriction Policy Optimization. For a given region, let N, S_{curr}, I_{curr} be the total, current susceptible, and infectious populations respectively. Let $(I_{avg}^{target}, I_{max}^{target})$ be the target average and maximum infectious levels. Let $\mathcal{A}$ be the set of restriction levels for which the transmission rate is known or can be estimated[2] and T, the decision horizon duration, then the goal is to identify restriction levels $[a_t],\ [t]_{curr+1}^{curr+T},\ a_t \in \mathcal{A}$ such that the infectious level averages at I_{avg}^{target} but does not exceed I_{max}^{target}.

Example. On Oct 1, 2020, a hypothetical city has a population of 13M of which 0.2M are currently infectious and 0.2M are post-infectious with the rest still susceptible. Assuming ten restriction levels with equispaced transmission rates between 0.1 to 0.55, the objective is to figure out a restriction policy schedule (possibly varying across days) for the upcoming month so that the infectious count averages $I_{avg}^{target} = 150,000$ with maximum of $I_{max}^{target} = 200,000$.

Since our primary focus is on an analytical control framework, we make simplifying assumptions on the observability, (i.e., accurate estimation of the infectious population is possible via a mix of sero-surveys and diagnosis tests) and the infection dynamics (region isolation, homogeneous interactions, negligible incubation period, and constant infectious period). Section 8 describes extensions when some of these assumptions are relaxed.

[2]Mapping restriction policies to transmission rates is itself the focus of multiple studies [24, 38] but not the primary concern in the current work.

Restriction level	Effective R value	COVID-19 Transmission rate (β)	Description
Early event (Pre-restrictions)	2.52	0.4	No policy interventions.
Lockdown	0.84	0.12	Residents are either not allowed to leave their residence or may leave only for essential functions.
Stay-at-home	0.95	0.14	Stay-at-home order includes closure of schools and private sector, and restrictions on mass gatherings.
Safer-at-home	1.26	0.19	Relaxed stay-at-home order which includes closure of schools, but restrictions on mass gatherings and private sector may be relaxed.
New normal	1.83	0.29	Stay-at-home lifted or relaxed with some restrictions on mass gatherings and private sector. Schools may or may not be re-opened. Safeguards such as face coverings encouraged.

Table 1: Example of restriction levels recommended by public health authorities [1] and the corresponding effective R value and transmission rate (β) for COVID-19.

3 BACKGROUND AND RELATED WORK

Our work builds on three research areas: (a) compartmental epidemiological models, (b) Lotka-Volterra systems, and (c) optimal control of non-linear dynamical systems.

3.1 Compartmental Models

Infectious diseases are commonly modeled using compartmental models where the population (N) is divided into various compartments corresponding to different disease stages and the inter compartment transitions are governed by the model dynamics.

The SIR model [28, 35] is the simplest and most widely used one. The model comprises three compartments: Susceptible (**S**), Infectious (**I**) and Removed (**R** - includes immune & post-infectious persons) with the dynamics in Figure 3(b). Here β is the rate of disease transmission from infectious to susceptible individuals, which largely depends on the contact restriction policy. γ is the inverse of the average infectious period, which depends on the testing and quarantine policy but is largely invariant when testing volumes are low [14]. In this model, the effective reproduction number (avg. number of direct infections from each infection) is $R_{eff} = \frac{\beta S}{\gamma N}$. Existing restriction control approaches [60] are often guided by the principle of ensuring $R_{eff} \simeq 1$.

Certain infectious diseases have a significant incubation period when the individuals are infected but are not spreading the disease (non-infectious). The SEIR model [30] includes an additional **E** (exposed) compartment to model the incubation phase. Though we focus on SIR model in this paper for brevity of the presentation, we discuss extensions of our control techniques to SEIR model in 8.

3.2 Lotka-Volterra Systems

Lotka-Volterra (LV) systems [12, 32, 65] model the population dynamics of predator-prey interactions in a biological ecosystem. These models form a special case of Kolmogorov systems [29] that capture the evolution of a stochastic process over time. In a simple 2-species LV system, the population of prey (p) interacts with that of predator (q). The growth rate of prey depends on its reproduction rate (r) and the rate of consumption by predator (e). The change in predator population depends on the nourishment-based birth rate b and its death rate d. The system has two fixed points: (a) a saddle point that maps to extinction $(p, q) = (0, 0)$, and (b) a stable equilibrium at $(p, q) = (\frac{d}{b}, \frac{r}{e})$. Typically, the system exhibits oscillations resulting in a closed phase plot that corresponds to conservation of an "energy" function. Figure 3(a) presents the dynamics of an LV system and the oscillations of the prey and predator populations. Due to the criticality of ecological population control, there has been considerable research on multiple variants of LV systems [29, 63] and their Hamiltonian dynamics [46].

3.3 Optimal Control Strategies

Optimal control of dynamical systems has rich connections to multiple fields [52] that deal with optimizing sequential decisions to maximize a desired objective such as reinforcement learning [20], multi-armed bandits [16], and stochastic control. Given a set of control variables, the optimal control policy describes the time derivatives of these variables that minimize the cost function and can be derived using Pontryagin's maximum principle [52] or the Hamilton-Jacobi-Bellman equations [48]. Though there exist comprehensive techniques for control of linear dynamical systems, the control strategies for non-linear dynamical systems rely heavily on the existence of control-Lyapunov functions, which are typically identified using conservation laws of the associated physical systems. Once a suitable Lyapunov function is identified, there exist multiple control design strategies such as feedback linearization, backstepping and sliding mode control that are guaranteed to converge using Artstein's theorem [59]. In the case of the SIR model, a suitable Lyapunov function is not readily evident. On the other hand, Lyapunov stability and practical control strategies of LV systems have been extensively studied [26, 39, 42, 47].

3.4 Epidemiological Modeling

Most of the existing literature on epidemiological modeling focuses on the following aspects: (a) Design of models to capture the disease dynamics [11, 28], (b) Accurate forecasting of future case counts [54], (c) Estimation of model parameters corresponding to non-pharmaceutical interventions (NPIs) [23] (d) Optimization of public-health policy (especially NPIs such as quarantine and lockdown policy) based on economic impact and disease burden [7, 9, 49, 51, 55]. Our current work is aligned with the last area, often referred to as economic epidemiology, in terms of the objective being the identification of optimal NPIs. However, most of the research in this space is focused on the variation of disease dynamics across subgroups using SIR or SEIR-based compartmental models

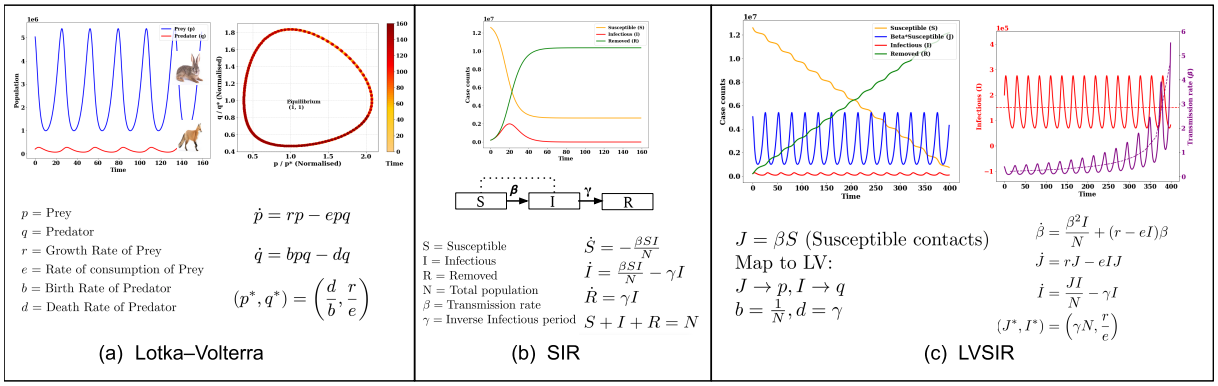

Figure 3: (a) Lotka-Volterra system and evolution of predator (fox) and prey (hare) population over time, (b) SIR compartmental dynamics for epidemiological modeling, (c) Mapping of SIR model to LV system and the behaviour of case counts S, J, I, R and transmission rate β.

and the explicit modeling of the economic and disease impact along with the trade-offs. The optimization problems are often formulated in terms of the cumulative outcomes across a horizon and not readily tractable due to the non-linear dependencies among the factors. Due to the subjective nature of the socioeconomic modeling assumptions and the computationally expensive solutions [7], this work is more suited for generating static recommendations rather than implementing an automated system.

Our work, on the other hand, focuses on dynamic adaptive control of the transmission rate of an epidemic assuming simple SIR dynamics with the goal of maintaining a steady acceptable target level of infections. This target level could be derived from any of the existing socioeconomic impact models or directly specified by a public health authority. Our primary contribution is a theoretical result that can be used to implement an automated decision support and alert system for a region of interest conditioned on the availability of a map between policy choices and transmission rate, and the applicability of SIR dynamics. As in the case of typical automated control systems (e.g., flight control), the control policy adapts based on observations to stay close to the equilibrium. Section 8 briefly discusses extensions of proposed theoretical results to other complex compartmental models and other control variables, which could be permit a similar adaptive control strategy for optimizing NPIs based directly on the overall impact.

4 MAPPING SIR TO LOTKA-VOLTERRA SYSTEM (LVSIR)

Our primary goal is to solve the contact restriction control problem described in Section 2. We focus on SIR dynamics because it captures the core disease spread mechanism of most epidemiological models. Since existing work on stability analysis of SIR models [8] does not address controllability, we first establish a connection between the SIR model and LV system, which is more amenable to principled control. In Section 5, we leverage the properties of the LV system to propose strategies for restriction control.

The problem of stabilizing infection levels assuming SIR dynamics has a direct analogy with population control in LV predator-prey systems where it is desirable to maintain the predator and prey population at certain target levels suitable for the ecosystem. Comparing the SIR and LV dynamics in Figure 3, we observe that the behaviour of the infectious people (I) is similar to that of the "predators" (q). There is an inflow (birth) $\beta SI/N$ that depends on β as well as the current infectious and susceptible population. There is also an outflow (death) γI from the **I** to the **R** compartment. However, the counterpart for the "prey" is not readily apparent.

An intuitive choice for "prey" is the "susceptible contacts" (i.e., the product of susceptible people and β, the number of "contacts" of a person per day) since this acts as "nourishment" to the predators and contributes to the inflow into the **I** compartment. Denoting the susceptible contacts by $J = \beta S$, we note that equivalence with the LV system requires

$$\dot{J} = \dot{\beta}S + \dot{S}\beta = (r - eI)J = (r - eI)\beta S, \tag{1}$$

where r and e correspond to the reproduction rate and consumption rate of an LV system as described in Section 3.2. Since $\dot{S} = -\frac{\beta IS}{N}$, we require the transmission rate β to follow

$$\dot{\beta} = (r - eI)\beta + \frac{\beta^2 I}{N}. \tag{2}$$

This modified version of SIR model (LVSIR) maps to a Lotka-Volterra system. Comparing the model parameters, we note that the inverse of the infectious period (γ) corresponds to the predator death rate (d) and the inverse of population ($1/N$) to the predator birth rate (b). The parameters (r, e) which correspond to reproductive rate and prey consumption rate are additional degrees of freedom that control the system dynamics.

Theorem 1 asserts the existence of a steady state for an epidemic following LVSIR dynamics. if the system were to be initialized in this state, the size of infectious population and the susceptible contacts would remain constant throughout the duration while the susceptible population goes down linearly.

Theorem 1. [3] *For the LVSIR model in Figure 3(c), the following holds true:*

(a) There exists a stable equilibrium at $(J^*, I^*) = (\gamma N, r/e)$.

(b) When the initial state $(J_0, I_0) = (J^*, I^*)$, then (J, I) remain constant while S, R take a linear form and β increases till the

[3] Please see Appendix A for the proofs.

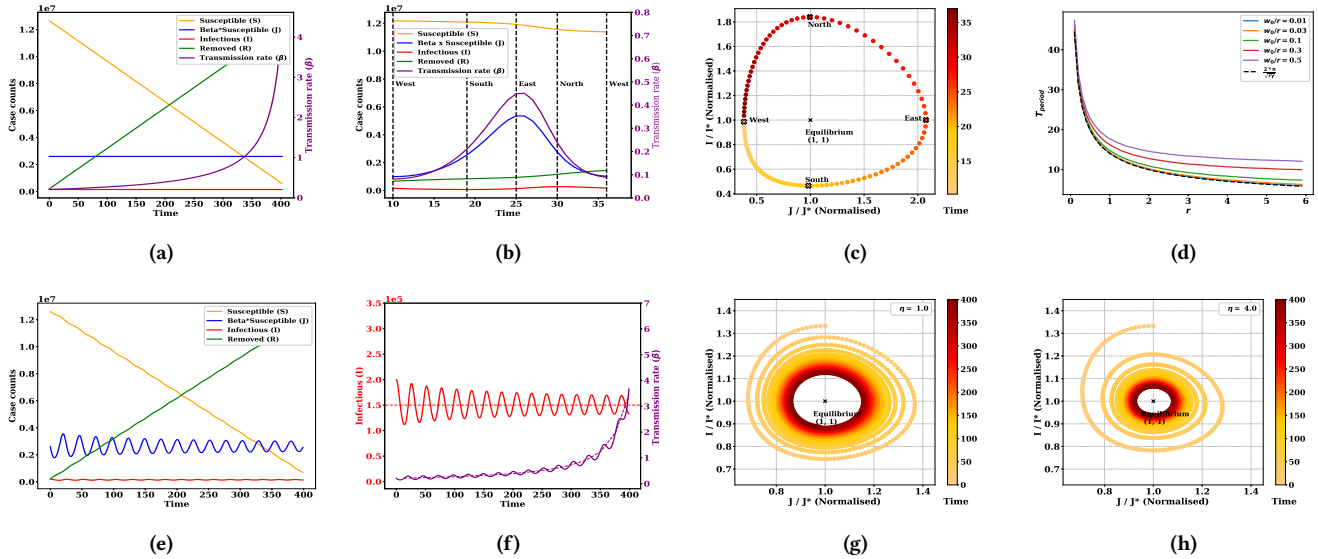

Figure 4: Plots (a-d) and (e-h) present the dynamics of LVSIR and CoSIR models respectively. (a) System (S, J, I, R, β) evolution when initialized at equilibrium, (b-c) Single period of oscillations and normalized phase plot for the configuration in Figure 3(c) annotated with the extreme points, (d) Dependence of T_{period} on r for different choices of w_0/r. (e-f) System evolution of the CoSIR model for learning rate $\eta = 1$. (g-h) Normalized phase plots of CoSIR model for $\eta = 1$ and 4. Any other parameter choices can be tried on the webapp.[1]

susceptible population reaches 0 at T_{end}.
(i) $S(t) = S_0 - \gamma I^* t$; $R(t) = R_0 + \gamma I^* t$
(ii) $\beta(t) = \frac{\gamma N}{S_0 - \gamma I^* t}$; $T_{end} = \frac{S_0}{\gamma I^*}$

When the LVSIR system is initialized at a non-equilibrium state, it follows an oscillatory behavior that is characterized in Theorem 2 along with the associated conservation property.

Theorem 2. [3] *For the LVSIR model in Figure 3(c), if the initial state is not at equilibrium, i.e., $(J_0, I_0) \neq (J^*, I^*)$, it exhibits a cyclic behaviour with the following properties.*

(a) The system "energy" $w(J, I) = \gamma(x - \log(x) - 1) + r(y - \log(y) - 1)$ (where $x = \frac{J}{J^*}, y = \frac{I}{I^*}$) remains constant at $w_0 = w(J_0, I_0) \geq w(J^*, I^*) = w^* = 0$ till termination.
(b) The I, J curves exhibit periodic oscillations resulting in a closed trajectory. The normalized phase plot has four extreme points $\{(x_{min}, 1), (1, y_{min}), (x_{max}, 1), (1, y_{max})\}$ where (x_{min}, x_{max}) are the roots of the equation $x - \log(x) = 1 + w_0/\gamma$ and (y_{min}, y_{max}) are the roots of the equation $y - \log(y) = 1 + w_0/r$.
(c) The cyclic period is given by

$$T_{period} = \int_{\log(x_{min})}^{\log(x_{max})} \left(\frac{1}{F_1^{-1}(G(z))} - \frac{1}{F_2^{-1}(G(z))}\right) dz,$$

where (x_{min}, x_{max}) are defined as above, $G(z) = \gamma(e^z - z - 1) - w_0$, and $F_1(s)$, $F_2(s)$ are restrictions of $F(s) = s + r\log(1 - \frac{s}{r})$ to positive and negative ranges. In general, T_{period} has the form $g(r, \gamma, \frac{w_0}{r})$ with approximation via linearization yielding $T_{period} \simeq \frac{2\pi}{\sqrt{(r\gamma)}}$.
(d) In each cyclic period, S reduces by a fixed amount $\Delta S = \gamma I^* T_{period}$. When the susceptible population $S < \Delta S$ at the beginning of a cycle, it reaches 0 during the cycle and the epidemic terminates.

Figure 3(c) depicts the oscillatory behaviour of the LVSIR model for the configuration in Table 3. Similar to an LV system, the "energy" which corresponds to a weighted Itakura-Saito distance [33] between (I, J) and the equilibrium (I^*, J^*) is conserved. The infectious population I (and susceptible contacts J) oscillates between $[y_{min}I^*, y_{max}I^*]$ (and $[x_{min}J^*, x_{max}J^*]$) during the entire period with an average value of I^* (and J^*) while the susceptible population reduces steadily in a staircase-like fashion. The transmission rate β also exhibits periodic oscillations but the average steadily goes up to compensate for the reduction in the susceptible population. Figure 4(c-b) shows the phase plot and the variation of the key quantities during a single period with the four extreme points (South, East, North, West) marked explicitly. Figure 4(d) shows how T_{period} [58] depends on r for different choices of w_0/r. This mapping can be used to identify suitable values of r for practical policy making. For the special case where the initial state is at equilibrium, the system behaviour is steady as in Figure 4(a).

5 TRANSMISSION RATE CONTROL (COSIR)

Let us consider an epidemic system from a control theory perspective with the infectious population as the output and the external injections of infections as input as shown in Figure 2. The SIR system has a natural positive feedback loop because of the infection spread mechanism (since $\dot{S} = \frac{\beta SI}{N}$) which leads to an exponential-like behavior in early stages of the epidemic. The LVSIR model neutralizes this feedback loop by varying β over time following Eqn. 2, but it does not necessarily convergence to a desired equilibrium. We now explore the problem of controlling the transmission rate β for the

LVSIR model (Figure 3(c)) to nudge the infectious levels to a desired equilibrium.

As discussed in Section 3, control of non-linear dynamical systems is typically achieved via Control-Lyapunov functions (CLFs). Our approach is to exploit the mapping from the SIR model to LV system and use CLFs derived from the "Lotka-Volterra energy" function $w(J, I)$.

Definition 5.1. Given a dynamical system $\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}, \mathbf{u})$ with state vector $\mathbf{z} \in D \subset \mathbb{R}^n$, control $\mathbf{u} \in \mathbb{R}^m$, and equilibrium state $\mathbf{z}^* = \mathbf{0}$, a control-Lyapunov function (CLF) is a function $V : D \mapsto \mathbb{R}$ that is continuously differentiable, positive-definite s.t. $\forall \mathbf{z} \neq 0, \exists\, \mathbf{u}$, $\dot{V}(\mathbf{z}, \mathbf{u}) = \langle \nabla V(\mathbf{z}), \dot{\mathbf{z}} \rangle = \langle \nabla V(\mathbf{z}), \mathbf{f}(\mathbf{z}, \mathbf{u}) \rangle < 0$.

The CLF $V(\cdot)$ can be viewed as a generalized energy function with $\dot{V}(\cdot)$ being a dissipation function. Artstein [10] proved that as long as there is a CLF, there exists a control $\mathbf{u}$ to ensure the reduction of energy at every non-equilibrium state and eventual convergence to the zero energy equilibrium.

Theorem 3 (Artstein [59]). *For any non-linear dynamical system with affine control, $\dot{\mathbf{z}} = \mathbf{f}(\mathbf{z}, \mathbf{u}) = \mathbf{f_0}(\mathbf{z}) + \sum_{i=1}^{m} \mathbf{f_i}(\mathbf{z}) u_i$ with state $\mathbf{z} \in D \subset \mathbb{R}^n$, control $\mathbf{u} \in \mathbb{R}^m$, has a CLF if and only if it admits a regular stabilizing control feedback $\mathbf{u}$, that is a locally Lipschitz function on $\mathbb{R}^n \setminus \{0\}$.*

Once a CLF is identified, it is relatively straightforward to design an appropriate control function $\mathbf{u}$ as described in [10, 62]. For our scenario, we rely on the conservation law of the LV system as well as the existing literature on its Lyapunov functions [53].

Let $\mathbf{z} = (\frac{J}{J^*} - 1, \frac{I}{I^*} - 1)$ so that the equilibrium $\mathbf{z}^* = (0, 0)$. Let $L(a_1, a_2) : \mathbb{R}_+^2 \mapsto \mathbb{R}$ be a continuously differentiable divergence such that $|\frac{dL}{da_1}| > 0$ and $L(a_1, a_2) > 0$, $\forall a_1 \neq a_2$ and $L(a_1, a_2) = 0 \iff a_1 = a_2$. Then, the function $V(\mathbf{z}) = L(w, w^{target})$ where $w = w(J, I)$ can be used as a CLF. Examples of $L(\cdot, \cdot)$ include L_p norms and Bregman divergences. We focus on the case where $w^{target} = w^* = 0$, and propose a controlled SIR model (CoSIR).

Theorem 4 (CoSIR). [3] *For the SIR model, a proportional additive control on β,*

$$\dot{\beta} = (r - eI)\beta + \frac{\beta^2 I}{N} + u\beta \tag{3}$$

converges to the equilibrium (J^*, I^*) when

$$u = -\eta(t) \frac{dL}{dw} \left(\frac{J}{J^*} - 1 \right)$$

with the learning rate $\eta(t) > 0$, $\forall t$.

A special case of the above construction with $V(\mathbf{z}) = L(w, w^*) = \frac{1}{2}(w - w^*)^2$ and a constant positive learning rate η referred to as the speed-gradient method was proposed by [47] in the context of population control. The resulting control policy is given by

$$u = -\eta(w(J, I) - w^*) \left(\frac{J}{J^*} - 1 \right). \tag{4}$$

Figure 4(e-f) show the behavior of the CoSIR model using Eqn. 4 and parameters from Table 3. It can be observed that the oscillations reduce with time and the I, J curves slowly converge to the equilibrium with the convergence rate dependent on η. The design of the control also makes it robust to perturbations in the infectious population as the system re-calibrates β as appropriate.

The β-control policy (Eqn. 3) can be interpreted as follows. The first term $\frac{\beta^2 I}{N}$ corresponds to the relaxation possible due to the decreasing susceptible population while the second term $(r - eI)\beta$ leads to oscillatory behavior, and the last term $u\beta = -\eta(t)\beta \frac{dL}{dw}(J/J^* - 1)$ ensures dissipation of energy and convergence to the equilibrium.

6 PRACTICAL TRANSMISSION RATE CONTROL

We now describe a practical solution to the public health restriction control problem in Section 2. Algorithm 1 outlines a holistic approach to obtain a restriction policy schedule using the optimal β-control in Theorem 4. There are four key steps.

Algorithm 1 CoSIR: Transmission Rate Control

Input:
N – Population,
$\mathcal{A}$ – Restriction policy levels,
$\rho : \mathcal{A} \mapsto \mathbb{R}_{++}$ – Map of levels to transmission rate,
$[S_t, I_t, R_t, a_t], [t]_0^{curr}$ – Case counts & restriction level history,
$I_{max}^{target}, I_{avg}^{target}$ – max and avg., target infection levels,
T – Decision horizon,
T_{period} – Cyclic period for restriction schedule
Output:
$\mathbf{a}^* = [a_t],\ [t]_{curr+1}^{curr+T},\ a_t \in \mathcal{A}$ – Near optimal restriction levels for next T time units s.t. infection levels conform to I_{avg}^{target}.
Method:
Data-driven Calibration:
$\gamma \leftarrow$ ComputeGamma$([S_t, I_t, R_t],\ [t]_0^{curr}, N)$
$\rho \leftarrow$ RefineBetaMap$([S_t, I_t, R_t, a_t],\ [t]_0^{curr}, N)$
Choosing CoSIR Parameters:
$\hat{\beta}_{curr} \leftarrow \frac{\gamma N}{S_{curr}}$
$(J^*, I^*) \leftarrow (\gamma N, I_{avg}^{target})$
$\alpha \leftarrow (y - \log(y) - 1)$ where $y = \frac{I_{max}^{target}}{I_{avg}^{target}}$
$r \leftarrow \frac{(2\pi)^2}{\gamma T_{period}^2}$ for $\alpha \simeq 0$
(or chosen s.t. $g(r, \gamma, \alpha) = T_{period}$ for $\alpha >> 0$ where $g(\cdot)$ is as defined in Theorem 2(c).)
$e \leftarrow \frac{r}{I^*}$
$\eta \leftarrow$ desired convergence rate
Determine Optimal Restrictions:
$[\beta_t^{ideal}],\ [t]_{curr+1}^{curr+T} \leftarrow$ CoSIR$(N, S_{curr}, I_{curr}, R_{curr}, \hat{\beta}_{curr}, r, e, \gamma, \eta, T)$

$a_t = \text{argmin}_{a \in \mathcal{A}}\ (\beta_t^{ideal} - \rho(a))^2,\ [t]_{curr+1}^{curr+T}$
return $\mathbf{a}^* = [a_t],\ [t]_{curr+1}^{curr+T}$

Input Collection. Infection level targets $(I_{avg}^{target}, I_{max}^{target})$, periodicity of the restriction schedule (T_{period}), decision horizon (T) need to be determined based on a careful assessment of public health and socioeconomic considerations. Historical case counts and restrictions $([S_t, I_t, R_t, a_t],\ [t]_0^{curr})$ also need to be collected to enable accurate optimization.

Algorithm	Description
`No-Restrictions`	Constant $\beta = 0.44$
`PL-high`	Periodic Lockdown of 60 days with $\beta = 0.16$ and a relaxation period of 30 days with $\beta = 0.44$
`PL-low`	Periodic Lockdown of 60 days with $\beta = 0.1$ and a relaxation period of 30 days with $\beta = 0.44$
`CoSIR`	Follows Eqn. 3
`CoSIR-approx`	Approximation of CoSIR β with 10 levels starting from 0.1 to 0.55 in steps of 0.05.

Table 2: Transmission control policies used for simulation.

Data-driven Calibration. The next step is to use SIR-calibration methods [11, 15, 54] along with historical data to estimate a static γ (ComputeGamma), time-varying β, and the state of the epidemic (S_{curr}, I_{curr}, R_{curr}). The restriction level to transmission map, ρ, can be initially chosen from public health guidelines [1] and refined using the observed β for past restrictions in the region of interest (RefineBetaMap).

Choosing CoSIR Parameters. The free parameters of the CoSIR model need to be chosen based on the control requirements. Algorithm 1 lists the choices derived from Theorem 2. These include the equilibrium state, the normalized energy level based on ratio of the maximum and average target infectious levels, the reproductive rate r and the consumption rate e. There is flexibility on the choice of $\hat{\beta}_{curr}$ and η. Choosing the immediate transmission rate to be $\hat{\beta}_{curr} = \frac{J^*}{S_{curr}} = \frac{\gamma N}{S_{curr}}$ (equivalent to forcing effective reproduction number $R_{eff} = 1$) ensures a maximal reduction in the system "energy" and faster convergence to the desired equilibrium, but dampens fluctuations. However, fluctuations might be necessary for economic activity. When nearly steady infection levels are desired, $I_{avg}^{target} \simeq I_{max}^{target}$, then $r = \frac{(2\pi)^2}{\gamma T_{period}^2}$ and high η are appropriate.

Computing Near Optimal Restrictions. Determining the restriction policy can be split into two phases. The first involves estimating the ideal β control from Eqn. 4 while the second involves identifying the "closest" restriction level for the ideal β at each time step with "closeness" based on a suitable divergence such as the squared loss.

Note that in a real-world implementation, there are likely to be lags in data collection as well as communicating the policy to public. The map between restriction policy and transmission rate might also vary with time because of changes in public compliance. Unlike the CoSIR model (Eqn. 3), the approximation is not guaranteed to converge. However, the behavior of the approximate algorithm approaches that of the original control policy as number of (uniformly distributed) restriction levels increases.

7 SIMULATION RESULTS

We present simulation results on the COVID-19 pandemic spread on multiple diverse geographical regions to demonstrate the efficacy and adaptability of the proposed `CoSIR`-based approaches relative to other widely cited approaches such as periodic lockdowns [21]. Since modeling of the socioeconomic costs of mobility restrictions is a challenging problem in itself, we primarily focus on assessing the efficacy in terms of the medical capacity utilization and the severity distribution of mobility restrictions. Note that even though COVID-19 disease progression involves a significant incubation period (4-5 days) that is ideally modeled via variants of SEIR compartmental models, forecasting techniques using SIR dynamics have also been fairly accurate for coarser temporal granularity [54].

7.1 Experimental Setup

The key elements of our experimental setup are described below.

Data. Multiple geographical regions that are diverse in population size, medical capacity, and reporting dynamics were considered. We present simulation results for three regions (India, Mexico, and Netherlands) from April 2020 to April 2021 along with real available data (till 29th December 2020) of the ongoing COVID-19 pandemic [2, 25]. To assess the adaptability of the CoSIR approach, we also use synthetically generated data from a hypothetical city.

Algorithms. We consider six variants of the transmission control policies assuming a no restriction level corresponding to $\beta = 0.44$. Two of these policies `PL-low`, `PL-high`, are based on dynamic periodic interventions [21] alternating between no restrictions for a period of 30 days and a constant, but different levels of restriction ($\beta = 0.1$ and $\beta = 0.16$) for a period of 60 days. In addition, we also study a completely non-restrictive policy `No-Restrictions`, the CoSIR approach `CoSIR` and an approximate variant `CoSIR-approx` based on $k = 10$ discrete levels. Finally, we also include the real outcome as indicative of the actual public health transmission control that was adopted. Table 2 enumerates the details of these algorithms.

Parameter Choices. Table 3 enumerates the various parameter choices and the relevant assumptions for each region.

(a) Target Hospitalization & Infectious levels. Choosing an optimal control policy requires consideration of multiple factors. In this study, we focus primarily on the medical capacity constraints. We derive an acceptable level for infectious population I_{avg}^{target} such that the available medical capacity can meet the hospitalization requirements at a steady state. Let p denote the per-capita hospital bed capacity of a region with population N, κ the bed occupancy level for normal functioning, and T_{hosp} the average duration of hospitalization. Then, a manageable hospitalization inflow that will not overwhelm the medical infrastructure is given by $\Delta H_{avg}^{target} = \frac{Np\kappa}{T_{hosp}}$. On the other hand, following the dynamics of the SIR model, we know that new cases being identified (i.e., outflow from the I bucket) is given by γI. Let h denote the fraction of active cases requiring hospitalization. Then, the new hospitalizations would be given by $\gamma h I$. Hence, we choose the target infectious level as $I_{avg}^{target} = \frac{Np\kappa}{\gamma h T_{hosp}}$.To concretely instantiate these choices for various regions, we use the populations of the regions N, the per-capita medical capacity [5] assuming an average hospitalization period $T_{hosp} = 20$ days and hospitalization ratio between $2 - 6.6\%$. Note that these parameters are likely to vary across regions and there is a significant uncertainty in the estimates across multiple studies. However, the key insights on the relative behavior of the transmission control policies hold true regardless of the specific choices of the parameter values.

(b) Under Reporting Factor. While comparing real observed case counts with simulations, a critical factor to consider is the level of under reporting. This factor was computed based on the assumption that the infection fatality rate (IFR) and death detection rate (DDR) are largely variant within a region and there is a steady lag of 2

[4]Since these estimations vary, we choose an intermediate value.

Factor	Notation	Hypothetical	India	Netherlands	Mexico
Population	N	13M	1.366B	17.1M	127.5M
Stage of pandemic	(S_{curr}, I_{curr})	(12.6M, 0.2M)	(1.36B, 279k)	(17.09M, 6.7k)	(127.2M, 30k)
Bed Occupancy level [50]	κ	-	0.76	0.76	0.76
Medical Capacity (per capita) [5]	p	-	5.3×10^{-4}	3.32×10^{-3}	1.38×10^{-3}
Hospitalization ratio	h	-	2%	6.6%	2%
Time spent in hospital	T_{hosp}	-	20 days	20 days	20 days
Target manageable infections	I_{avg}^{target}	0.15M	6.88M	0.16M	1.67M
Target max to average ratio	$I_{max}^{target}/I_{avg}^{target}$	1.3	1.13	1.13	1.13
Infection-Fatality ratio [41, 43–45]	IFR	-	0.002 [4]	0.0068	0.0068
Death detection ratio	DDR	-	0.3	1	0.35
Reported Infections on 1st Dec 2020 [2, 64]		-	9.5M	527k	1.11M
Reported Fatalities on 15th Dec 2020 [2, 64]		-	144k	10k	114k
Under reporting factor	URF	-	25.28	2.83	43.127
Periodicity	T_{period}	7 days	7 days	7 days	7 days
Reciprocal of infectious period	γ	0.2	0.2	0.2	0.2
LV Reproduction rate	r	4.2	4.2	4.2	4.2
Consumption rate	e	2.8×10^{-5}	6.1×10^{-7}	2.6×10^{-5}	2.5×10^{-6}
Initial transmission rate	β_0	0.2	0.2	0.2	0.2
Learning rate	η	5	5	5	5

Table 3: State of the epidemic, public health requirements, and CoSIR parameters used for the simulations in Section 7.

weeks between infection and the associated fatalities. In particular, we estimate the under reporting factor (URF) as

$$URF = \frac{\text{Reported Infections}}{\text{Total Infections}} = \frac{(\text{Reported Infections on Dec 1st}) \times DDR \times IFR}{\text{Reported Fatalities on Dec 15th}}$$

(c) SIR & CoSIR Parameters. The primary parameter of interest in the SIR model apart from the transmission rate is γ, i.e., reciprocal of the infectious period, which was chosen to be 0.2 across all simulations since it is a disease-specific factor. For practical implementation of a control policy, a periodicity aligned with typical economic activity is preferable and hence, we choose a periodicity of $T_{period} = 7$ days for the CoSIR model. Assuming the maximum value of infections to be 13% higher than the average and following Algorithm 1, we obtain the LV reproductive rate $r = 4.2$. We choose the steady state infectious level to be the target infectious level, i.e., $I^* = I_{avg}^{target}$ and the consumption rate is given by $e = r/I^*$. To ensure effective control, we pick an aggressive learning rate $\eta = 5$.

Evaluation Metrics. In order to assess the effectiveness of the different control policies, we focus on two aspects.

(a) **Hospitalization Influx,** which is defined as the number of new cases being hospitalized every day, which equals γhI for a SIR-based model. We compare this inflow rate with the manageable target influx based on the available capacity, i.e., $\Delta H_{avg}^{target} = \frac{Np\kappa}{T_{hosp}}$. We also estimate the relative demand surplus and gap relative to this target level and average it across the relevant duration.

- Excess hospitalization $= \min\left(\frac{\gamma hI}{\Delta H_{avg}^{target}} - 1, 0\right)$
- Under hospitalization $= \min\left(1 - \frac{\gamma hI}{\Delta H_{avg}^{target}}, 0\right)$

In addition to influx analysis, one can also evaluate the utilization rate which depends on both inflow and outflow. Consider an SIHR compartmental model with an additional H compartment for post-infectious cases with an average duration of T_{hosp} days. Assessing the H compartment size adjusted for hospitalization ratio relative to the available capacity gives an estimate of utilization, but we omit these results in the interest of brevity since the insights are similar.

(b) **Mobility Restriction Distribution.** To assess the impact of the different transmission control policies on socioeconomic activity, we consider the variation of the restriction level, i.e., β over time. We also compute for each policy, the total number of days where the transmission rate associated with the policy was higher (i.e., less restrictive) than a particular β and discuss the associated implications on socioeconomic impact.

7.2 Medical Capacity Utilization.

Figures 5[a-c] show results from simulating the five transmission control policies over the three regions of interest (India, Mexico, Netherlands) from April 2020 to April 2021 along with the real case counts. The curves depict the simulated daily new hospitalizations ($= \gamma hI$). Real hospitalization counts are obtained by appropriate scaling of the reported active case counts using the under reporting factor URF and the hospitalization fraction h. From the simulations, it is evident that the `No-Restrictions` policy leads to a rapid escalation of cases resulting in a hospitalization demand significantly higher than the capacity. The periodic lockdown approach `PL-high` with $\beta_{relax} = 0.16$ is also not restrictive enough to limit the hospitalization needs, while the periodic lockdown `PL-low` turns out to be excessively conservative especially in Netherlands. The `CoSIR` approach, on the other hand, results in a nearly steady (periodic with small oscillations) hospitalization inflow that matches the specified target value resulting in a near-optimal utilization of the medical capacity. The approximate (`CoSIR-approx`) approach with ten restriction levels results in a similar outcome, but begins to diverge when the susceptible population reduces so much that

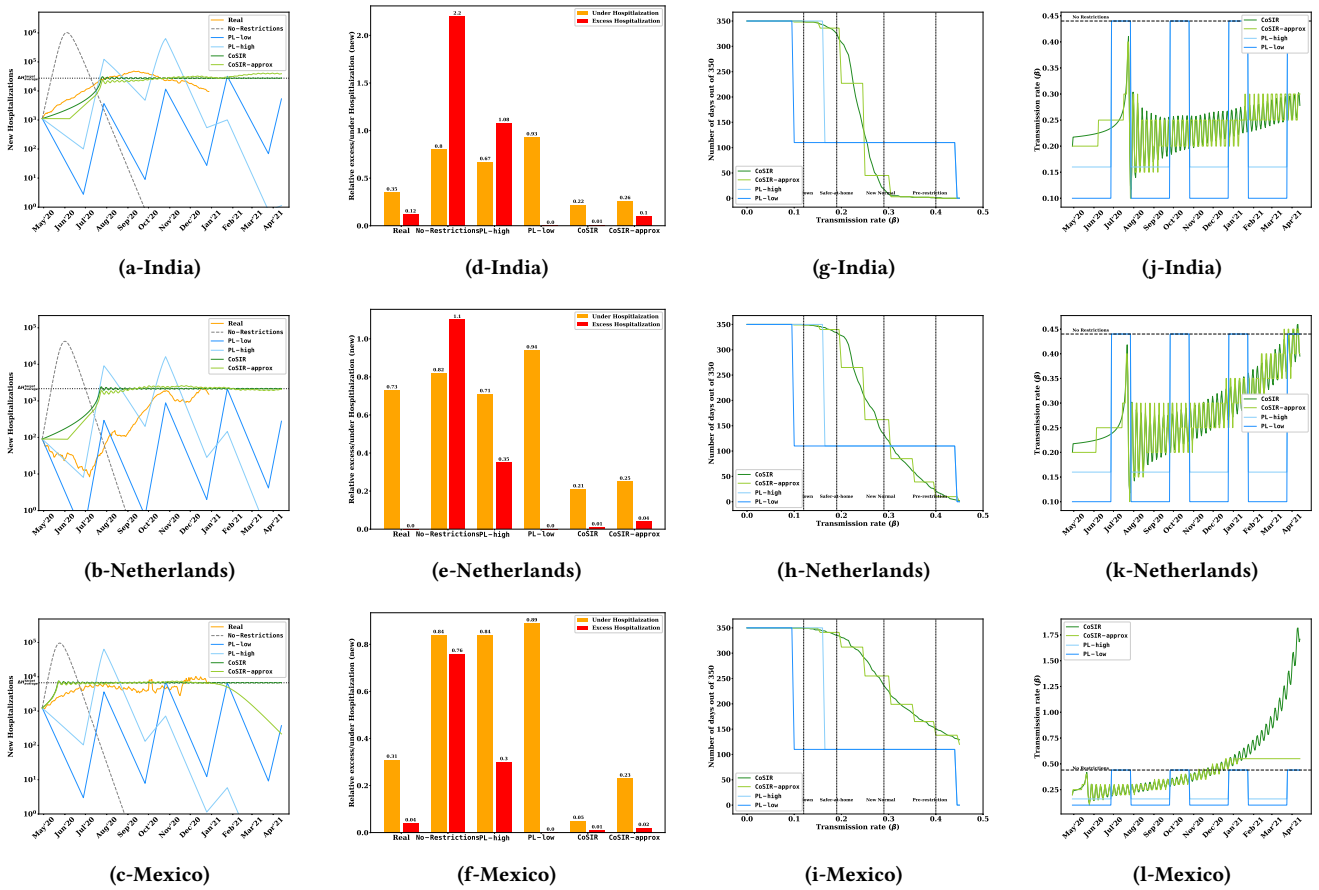

Figure 5: Simulation results using transmission control policies for three regions (India, Mexico, Netherlands). Plots [a-f] show the hospitalization inflow for various policies along with the real observations as well as the average excess and under hospitalization levels. Plots [g-l] depict the variation of policy transmission rate with time and also the restrictiveness profile (number of days where the policy is less restrictive than a specified β) for the periodic lockdown and CoSIR-based policies.

the `CoSIR` optimal β is higher than the no restrictions level. This is especially evident in the case of Mexico (Figure 5-l). The real hospitalizations corresponding to the actual imposed restrictions point to periods where the healthcare system in some regions (India, Mexico) was overwhelmed. In the case of Netherlands, the early policy was possibly more restrictive than necessary.

Figures 5[d-f] depict the average excess and under hospitalizations relative to the specified target hospital inflow rate. The `CoSIR`-based approaches result in smaller overall deviations from the target levels relative to other policies including the actual public health policy that was adopted. There is a possibility that the respective health authorities preferred a different target than the one in our simulations. The main takeaway is that periodic lockdowns can be beneficial in containing the infection spread, but the actual transmission rate β needs to be determined based on the available capacity and the varying susceptible population, as is being done in the CoSIR approach.

7.3 Mobility Restriction Impact.

Figures 5[j-l] show the variation of the transmission rate (β) with time for four intervention policies (`CoSIR`, `CoSIR-approx`, `PL-high`, `PL-low`) for the same four regions. All the approaches involve alternating between different levels of restrictions but the CoSIR based approaches adapt these levels as the susceptible population decreases over time. In the case of Mexico and Netherlands, this approach nearly permits a return to no restrictions by April 2021 even without other interventions such as vaccination (conditioned on the assumptions on the acceptable hospitalization levels).

Figures 5[g-i] show the number of days where a control policy allows a transmission rate higher than a particular β, i.e., in other words allows all socioeconomic activities that necessitate a transmission rate of β. The periodic lockdown policies tend to have a discontinuous profile while the CoSIR based policies allow a more continuous transition. In particular, a large number of vital socioeconomic activities correspond to β between "safer-at-home" and "new normal". The CoSIR based policies clearly dominate in this region. Given a socioeconomic model that maps β to socioeconomic

costs and the distribution of the transmission rate β, we could quantitatively evaluate the socioeconomic impact of the various policies. We do not provide such an impact analysis due to the complexity of socioeconomic modeling in the context of significant variations between high income and low-middle income regions [57, 66]. However, since small relaxations in restriction levels in the range between "stay/safer at home" policies and "new normal" do seem to allow a substantial increase in socioeconomic activity with diminishing returns beyond that stage, the differences in the β profile are likely to be amplified when the socioeconomic impact is considered. In the case of Mexico, where the simulation duration covers most of the pandemic, the CoSIR based approaches turn out to be more relaxed than the fixed periodic lockdowns and that would be the case even for the other regions if the entire duration of the pandemic is considered.

7.4 Adaptability of CoSIR

An important characteristic of the CoSIR approach is its adaptability to sudden changes in the infection levels. Figure 6 shows the epidemic evolution with three different control policies for a hypothetical city with parameters listed in Table 3. The first policy is a `hypothetical`, but realistic one similar to the one adopted in a large number of cities that experienced multiple waves of infections while the other two are the `CoSIR` and `CoSIR-approx` based on the acceptable target hospitalization levels. As in the case Section 7.2, the CoSIR-based approaches lead to a relatively steady rate of infections while the hypothetical premature relaxation policy tends to result in infections peaking in narrow time intervals and burdening the healthcare system. More importantly, when the infection levels are subject to sudden upward ($t = 50$) or downward perturbations ($t = 100$) as in the case of super-spreader events or sudden quarantine restrictions respectively, the `CoSIR` β-control mechanism is able to adapt to these changes and continue pushing towards the equilibrium. In practice, the adaptation is not seamless and depends on the data collection frequency and the lag required for public communications to ensure compliance. Incorporation of mobility data, which is a leading indicator of actual infections, can aid with faster adaptation.

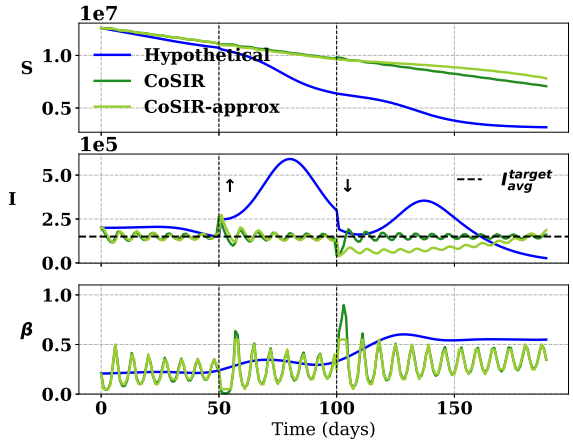

Figure 6: Disease progression (S, I, β) with three restriction policies: `hypothetical`, `CoSIR` and `CoSIR-approx` (using specified levels). The `CoSIR` variants maintain infectious counts close to target levels and also adapt to a sudden increase ($t = 50$) or decrease ($t = 100$).

8 EXTENSIONS

We now briefly describe some key extensions.

Delayed SIR & SEIR Models. The SEIR model allows explicit modelling of the incubation period and is known to closely mimic the behaviour of the delayed SIR model [34]. When β follows Eqn. 2, the delayed SIR model readily maps to a delayed LV system with a non-preying growth period for the predators, which is a special case of the well-studied Wangersky–Cunningham systems [40, 63]. It can be shown that the modified delayed SIR system with a delay τ has the same equilibrium $(J^*, I^*) = (\gamma N, r/e)$, exhibits (unbounded) oscillations and permits control of the form Eqn. 3, where

$$u(t) = -\eta(t)\frac{dL}{dw}\left(\frac{J(t)}{J^*} - 1\right) - r\left(\frac{(I(t) - I^*)}{(J(t) - J^*)}\right)\left(\frac{J(t-\tau)I(t-\tau) - J(t)I(t)}{I(t)I^*}\right)$$

with $\eta(t) > 0$. There is a need for special handling when J approaches J^* with the behavior depending on τ.

Testing & Isolation Policy. Testing, tracing and isolation also play a critical role in regulating the epidemic. In terms of SIR and SEIR dynamics, the net effect of aggressive testing is minimizing the infectious period [13, 37]. This is analogous to the culling of predators (infectious population) by increasing the death rate for which there already exists multiple control mechanisms [42]. In particular, choosing $V(\mathbf{z}) = L(w(J, I), w^*)$ as the CLF of interest, we obtain the control, $\gamma(t) = \gamma_0 + \zeta(t)\frac{dL}{dw}(\frac{I}{I^*} - 1)$, with $\zeta(t) > 0$.

Online Learning. The restriction control problem can also be posed as a non-linear contextual bandit [16] formulation with the cumulative LV energy $w(J, I)$ of the CoSIR model over a future horizon interpreted as the (negative) "reward". Here, the discrete restriction levels can be viewed as the multiple arms of a bandit, the context includes the state of the epidemic, and the "reward" distribution is computed using the context and the observed transmission rate for the arms.

9 CONCLUSION & FUTURE DIRECTIONS

Our current work proposes an analytical framework for epidemic control with the intent of supporting an active goal-oriented public health response. The proposed framework relies on a mapping between SIR dynamics to Lotka-Volterra systems when the transmission rate varies with time following a certain form (LVSIR) and an additional feedback control mechanism (CoSIR). Given the vast literature on control of LV systems, this mapping can be leveraged to design new epidemic control techniques as well as extend current results to richer heterogeneous compartmental models and additional control variables (e.g., testing levels). Simulation results on a diverse set of regions point to the potential efficacy and utility of this approach. Effective practical implementation as part of public health systems requires robust estimation of parameters associated with NPIs, addressing the limitations of SIR dynamics, the lags in observation and communication systems, and incorporation of additional signals such as mobility [4]. It might also be beneficial to explore online and reinforcement learning variants [17, 52].

REFERENCES

[1] Analysis and mapping of policies - COVID-AMP, 2020. Accessed October 2, 2020. URL: https://covidamp.org/.

[2] Coronavirus in India: Latest Map and Case Count, 2020. Accessed January 1, 2021. URL: https://www.covid19india.org/.

[3] COVID-19: A Global Perspective, 2020. URL: https://www.gatesfoundation.org/goalkeepers/report/2020-report/#GlobalPerspective.

[4] Google mobility data, 2020. Accessed October 2, 2020. URL: https://www.google.com/covid19/mobility/.

[5] List of countries by hospital beds, Dec 2020. URL: https://en.wikipedia.org/wiki/List_of_countries_by_hospital_beds.

[6] New Zealand COVID-19 alert levels, 2020. Accessed October 2, 2020. URL: https://covid19.govt.nz/assets/resources/tables/COVID-19-alert-levels-detailed.pdf.

[7] Daron Acemoglu, Victor Chernozhukov, Iván Werning, and Michael D Whinston. Optimal targeted lockdowns in a multi-group SIR model. *NBER Working Paper*, 27102, 2020.

[8] Phillipe Adda and Derdei Bichara. Global stability for SIR and SIRS models with differential mortality. *International Journal of Pure and Applied Mathematics*, 80(3):425–433, 2012.

[9] Fernando E Alvarez, David Argente, and Francesco Lippi. A simple planning problem for COVID-19 lockdown. Technical report, National Bureau of Economic Research, 2020.

[10] Zvi Artstein. Stabilization with relaxed controls. *Nonlinear Analysis: Theory, Methods & Applications*, 7(11):1163–1173, 1983. doi:10.1016/0362-546X(83)90049-4.

[11] Jackie Baek, Vivek F. Farias, Andreea Georgescu, Retsef Levi, Tianyi Peng, Deeksha Sinha, Joshua Wilde, and Andrew Zheng. The limits to learning an SIR process: Granular forecasting for COVID-19, 2020. arXiv:2006.06373.

[12] Stephen Baigent. Lotka-Volterra dynamics: An introduction. 2016. URL: https://www.ucl.ac.uk/~ucess29/LTCC_LV2010.pdf.

[13] Yinon M. Bar-On, Ron Sender, Avi I. Flamholz, Rob Phillips, and Ron Milo. A quantitative compendium of COVID-19 epidemiology, 2020. arXiv:2006.01283.

[14] David W Berger, Kyle F Herkenhoff, and Simon Mongey. An SEIR infectious disease model with testing and conditional quarantine. Technical report, National Bureau of Economic Research, 2020. URL: https://www.nber.org/papers/w26901.

[15] Luís M. A. Bettencourt and Ruy M. Ribeiro. Real time Bayesian estimation of the epidemic potential of emerging infectious diseases. *PLOS ONE*, 3(5):1–9, 2008. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0002185.

[16] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26, 2011. URL: http://proceedings.mlr.press/v15/beygelzimer11a.html.

[17] Ankit Bhardwaj*, Han Ching Ou*, Haipeng Chen, Shahin Jabbari, Milind Tambe, Rahul Panicker, and Alpan Raval. Robust lock-down optimization for COVID-19 policy guidance. In *AAAI Fall Symposium*, 2020.

[18] Michelangelo Bin, Peter Cheung, Emanuele Crisostomi, Pietro Ferraro, H. Lhachemi, R. Murray-Smith, Connor Myant, Thomas Parisini, Robert Shorten, S. Stein, and L. Stone. On fast multi-shot COVID-19 interventions for post lock-down mitigation, 2020. arXiv:2003.09930.

[19] W.E. Boyce, R.C. DiPrima, and D.B. Meade. *Elementary Differential Equations and Boundary Value Problems*. Wiley, 2017. URL: https://books.google.co.in/books?id=SyaVDwAAQBAJ.

[20] Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A Lyapunov-based approach to safe reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 8092–8101, 2018.

[21] Rajiv Chowdhury, Kevin Heng, Md Shajedur Rahman Shawon, Gabriel Goh, Daisy Okonofua, Carolina Ochoa-Rosales, Valentina Gonzalez-Jaramillo, Abbas Bhuiya, Daniel Reidpath, Shamini Prathapan, et al. Dynamic interventions to control COVID-19 pandemic: a multivariate prediction modelling study comparing 16 worldwide countries. *European Journal of Epidemiology*, 35(5):389–399, 2020. doi:10.1007/s10654-020-00649-w.

[22] IHME COVID, Christopher JL Murray, et al. Forecasting the impact of the first wave of the COVID-19 pandemic on hospital demand and deaths for the USA and European economic area countries. *medRxiv*, 2020. doi:10.1101/2020.04.21.20074732.

[23] Nicholas G Davies, Adam J Kucharski, Rosalind M Eggo, Amy Gimma, W John Edmunds, Thibaut Jombart, Kathleen O'Reilly, Akira Endo, Joel Hellewell, Emily S Nightingale, et al. Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study. *The Lancet Public Health*, 2020.

[24] Amélie Desvars-Larrive, Elma Dervic, Nils Haug, Thomas Niederkrotenthaler, Jiaying Chen, Anna Di Natale, Jana Lasser, Diana S. Gliga, Alexandra Roux, Johannes Sorger, Abhijit Chakraborty, Alexandr Ten, et al. A structured open dataset of government interventions in response to COVID-19. *Scientific Data*, 7(1):285, 2020. doi:10.1038/s41597-020-00609-9.

[25] Ensheng Dong, Hongru Du, and Lauren Gardner. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet infectious diseases*, 20(5):533–534, 2020.

[26] Awad El-Gohary and M.T. Yassen. Optimal control and synchronization of Lotka–Volterra model. *Chaos, Solitons & Fractals*, 12(11):2087 – 2093, 2001. doi:10.1016/S0960-0779(00)00023-0.

[27] Neil Ferguson, Daniel Laydon, Gilani G Nedjati, N Imai, K Ainslie, M Baguelin, S Bhatia, A Boonyasiri, Perez Z Cucunuba, G Cuomo-Dannenburg, A Dighe, et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Technical report, Imperial College, London, 2020. doi:10.25561/77482.

[28] Tiberiu Harko, Francisco S.N. Lobo, and M.K. Mak. Exact analytical solutions of the susceptible-infected-recovered (SIR) epidemic model and of the SIR model with equal death and birth rates. *Applied Mathematics and Computation*, 236:184 – 194, 2014. doi:10.1016/j.amc.2014.03.030".

[29] Alexandru Hening and Dang H. Nguyen. Coexistence and extinction for stochastic Kolmogorov systems. *Annals Applied Probability*, 28(3):1893–1942, 06 2018. doi:10.1214/17-AAP1347.

[30] Herbert W Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4):599–653, 2000. doi:10.1137/S0036144500371907.

[31] Sze-Bi Hsu. A remark on the period of the periodic solution in the Lotka-Volterra system. *Journal of Mathematical Analysis and Applications*, 95(2):428–436, 1983. doi:10.1016/0022-247X(83)90117-8.

[32] Timon Idema. The behaviour and attractiveness of the Lotka-Volterra equations. Master's thesis, Universiteit Leiden, 2005. URL: https://www.sas.upenn.edu/~tidema/publications/maththesis.pdf.

[33] Fumitada Itakura and Shuzo Saito. Analysis synthesis telephony based upon the maximum likelihood method. pages 17–20, 1968.

[34] Abdelilah Kaddar, Abdelhadi Abta, and Talibi Hamad. A comparison of delayed SIR and SEIR epidemic models. *Nonlinear Analysis. Modelling and Control*, 2:181–190, 2011. doi:10.15388/NA.16.2.14104.

[35] William Ogilvy Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London*, 115(772):700–721, 1927. doi:10.1098/rspa.1927.0118.

[36] Jackson A. Killian, Marie Charpignon, Bryan Wilder, Andrew Perrault, Milind Tambe, and Maimuna S. Majumder. Evaluating COVID-19 lockdown and business-sector-specific reopening policies for three US states. In *KDD 2020 Workshop on Humanitarian Mapping*, 2020. doi:10.2139/ssrn.3598744.

[37] Daniel B Larremore, Bryan Wilder, Evan Lester, Soraya Shehata, James M Burke, James A Hay, Milind Tambe, Michael J Mina, and Roy Parker. Test sensitivity is secondary to frequency and turnaround time for COVID-19 surveillance. *MedRxiv*, 2020. doi:10.1101/2020.06.22.20136309.

[38] Lukas Lehner. COVID19 policy trackers, 2020. URL: https://lukaslehner.github.io/covid19policytrackers/.

[39] A. Magyar, G. Szederkényi, and K.M. Hangos. Globally stabilizing feedback control of process systems in generalized Lotka–Volterra form. *Journal of Process Control*, 18(1):80 – 91, 2008. doi:10.1016/j.jprocont.2007.05.003.

[40] Annik Martin and Shigui Ruan. Predator-prey models with delay and prey harvesting. *Journal of Mathematical Biology*, 43:247–267, 2001. doi:10.1007/s002850100095.

[41] Gideon Meyerowitz-Katz and Lea Merone. A systematic review and meta-analysis of published research data on COVID-19 infection-fatality rates. *medRxiv*, 2020.

[42] Magno Enrique Mendoza Meza, Amit Bhaya, and Eugenius Kaszkurewicz. Controller design techniques for the Lotka-Volterra nonlinear system. *SBA: Controle & Automaccão Sociedade Brasileira de Automatica*, 16:124 – 135, 2005. doi:10.1590/S0103-17592005000200002.

[43] Siuli Mukhopadhyay and Debraj Chakraborty. Estimation of undetected COVID-19 infections in India. *medRxiv*, 2020. URL: https://www.medrxiv.org/content/early/2020/05/03/2020.04.20.20072892, arXiv:https://www.medrxiv.org/content/early/2020/05/03/2020.04.20.20072892.full.pdf, doi:10.1101/2020.04.20.20072892.

[44] Manoj Murhekar, Tarun Bhatnagar, Sriram Selvaraju, V Saravana Kumar, Jeromie Wesley Vivian Thangaraj, Naman Shah, Muthusamy Santhosh Kumar, Kiran Rade, R Sabarinathan, Smita Asthana, et al. SARS-CoV-2 antibody prevalence in India: Findings from the second nationwide household serosurvey, August-September 2020. 2020.

[45] Manoj V Murhekar, Tarun Bhatnagar, Sriram Selvaraju, Kiran Rade, V Saravanakumar, Jeromie Wesley Vivian Thangaraj, Muthusamy Santhosh Kumar, Naman Shah, R Sabarinathan, Alka Turuk, et al. Prevalence of SARS-CoV-2 infection in India: Findings from the national serosurvey, May-June 2020. *Indian Journal of Medical Research*, 152(1):48, 2020.

[46] Yavuz Nutku. Hamiltonian structure of the Lotka-Volterra equations. *Physics Letters A*, 145(1):27 – 28, 1990. doi:10.1016/0375-9601(90)90270-X.

[47] Irina Pchelkina and Alexander L. Fradkov. Control of oscillatory behavior of multispecies populations. *Ecological Modelling - ECOL MODEL*, 227:1–6, 2012. doi:10.1016/j.ecolmodel.2011.10.022.

[48] Shige Peng. Stochastic Hamilton–Jacobi–Bellman equations. *SIAM Journal on Control and Optimization*, 30(2):284–304, 1992. doi:10.1137/0330018.

[49] James Petrie and Joanna Masel. Marginal value of quarantine. *medRxiv*, 2020. URL: https://www.medrxiv.org/content/early/2020/11/28/2020.11.24.20238204, arXiv:https://www.medrxiv.org/content/early/2020/11/28/2020.11.24.20238204.full.pdf, doi:10.1101/2020.11.24.20238204.

[50] P Joseph Phillip, Ross Mullner, and Steven Andes. Toward a better understanding of hospital occupancy rates. *Health Care Financing Review*, 5(4):53, 1984.

[51] Facundo Piguillem and Liyan Shi. Optimal COVID-19 quarantine and testing policies. 2020.

[52] Warren B Powell. From reinforcement learning to optimal control: A unified framework for sequential decisions, 2019. arXiv:1912.03513.

[53] Yuri Aleksandrovich Pykh. Lyapunov functions for Lotka-Volterra systems: An overview and problems. *IFAC Proceedings Volumes*, 34(6):1549 – 1554, 2001. doi:10.1016/S1474-6670(17)35410-1.

[54] Evan L Ray, Nutcha Wattanachit, Jarad Niemi, Abdul Hannan Kanji, Katie House, Estee Y Cramer, Johannes Bracher, Andrew Zheng, Teresa K Yamana, Xinyue Xiong, Spencer Woody, Yuanjia Wang, et al. Ensemble forecasts of Coronavirus disease 2019 (COVID-19) in the U.S. *medRxiv*, 2020. doi:10.1101/2020.08.19.20177493.

[55] Anna Scherbina. Determining the optimal duration of the COVID-19 suppression policy: A cost-benefit analysis. *Economics Working Paper*, 3, 2020.

[56] Snehal Shekatkar, Bhalchandra Pujari, Mihir Arjunwadkar, Dhiraj Kumar Hazra, Pinaki Chaudhuri, Sitabhra Sinha, Gautam Menon, Anupama Sharma, and Vishwesha Guttal. INDSCI-SIM A state-level epidemiological model for India, 2020. Ongoing Study at https://indscicov.in/indscisim.

[57] Adam Sheridan, Asger Lau Andersen, Emil Toft Hansen, and Niels Johannesen. Social distancing laws cause only small losses of economic activity during the COVID-19 pandemic in Scandinavia. *Proceedings of the National Academy of Sciences*, 117(34):20468–20473, 2020. URL: https://www.pnas.org/content/117/34/20468, arXiv:https://www.pnas.org/content/117/34/20468.full.pdf, doi:10.1073/pnas.2010068117.

[58] Shagi-Di Shih. The period of a Lotka-Volterra system. *Taiwanese Journal of Mathematics*, 1(4):451–470, 1997.

[59] Eduardo D. Sontag. A 'universal' construction of Artstein's theorem on nonlinear stabilization. *Systems & Control Letters*, 13(2):117 – 123, 1989. doi:10.1016/0167-6911(89)90028-5.

[60] Kevin Systrom, Thomas Vladek, and Mike Krieger. Rt. live. https://github.com/rtcovidlive/covid-model, 2020.

[61] John Tsinias. Stabilization of affine in control nonlinear systems. *Nonlinear Analysis: Theory, Methods & Applications*, 12(11):1283 – 1296, 1988. doi:10.1016/0362-546X(88)90060-0.

[62] John Tsinias. Sufficient Lyapunov-like conditions for stabilization. *Mathematics of Control, Signals and Systems*, 2(4):343–357, 1989. doi:10.1007/BF02551276.

[63] Peter J Wangersky and WJ Cunningham. Time lag in prey-predator population models. *Ecology*, 38(1):136–139, 1957. doi:10.2307/1932137.

[64] worldometer. worldometer, 2021. URL: https://www.worldometers.info/.

[65] Xiang-Ping Yan and Yan-Dong Chu. Stability and bifurcation analysis for a delayed Lotka–Volterra predator–prey system. *Journal of Computational and Applied Mathematics*, 196(1):198 – 210, 2006. doi:10.1016/j.cam.2005.09.001.

[66] Barnett-Howell Zachary and Ahmed Mushfiq Mobarak. Should low-income countries impose the same social distancing guidelines as Europe and North America to halt the spread of COVID-19? *Yale Research Initiative on Innovation and Scale (Y-RISE)*, 2020.

A PROOFS

Definition 1. *[19] Let $\mathbf{z}^* \in \mathbb{R}^n$ be a critical point of a system of ODEs. The critical point $\mathbf{z}^*$ is stable if, for any $\epsilon > 0\ \exists\ \delta > 0$ such that if $\mathbf{z} = \phi(t)$ satisfies $||\phi(0) - \mathbf{z}^*|| < \delta$ then $||\phi(t) - \mathbf{z}^*|| < \epsilon,\ \forall\ t > 0$.*

Proof of Theorem 1(a).
At equilibrium (J^*, I^*), we have $\dot{J} = 0$ and $\dot{I} = 0$. From Eqn. 1, it follows that

$$\dot{J}|_{(J,I)=(J^*,I^*)} = rJ^* - eI^*J^* = 0 \Rightarrow I^* = \frac{r}{e},$$

$$\dot{I}|_{(J,I)=(J^*,I^*)} = \frac{J^*I^*}{N} - \gamma I^* = 0 \Rightarrow J^* = \gamma N.$$

To prove the stability of the critical point at (J^*, I^*), let us consider the normalized variables $\phi(t) = \big(x(t), y(t)\big)$ where $x(t) = \frac{J(t)}{J^*}$ and $y(t) = \frac{I(t)}{I^*}$. $\mathbf{z}^* = (1, 1)$ is the corresponding critical point.

$$||\phi(0) - \mathbf{z}^*||^2 = (x_0 - 1)^2 + (y_0 - 1)^2 < \delta^2$$
$$\Rightarrow 1 - \delta < x_0 < 1 + \delta,\ 1 - \delta < y_0 < 1 + \delta.$$

Let $f(s) = s - \log(s) - 1$. Since $f(s)$ is a convex function, $1 - \delta < s < 1 + \delta$, implies $f(s) < \max\{f(1 + \delta), f(1 - \delta)\}$. Denoting this bound by D_{max} implies $f(x_0) < D_{max}$ and $f(y_0) < D_{max}$
From Theorem 2(a), we note that

$$\begin{aligned} w(J_0, I_0) &= \gamma f(x_0) + r f(y_0) \\ &< (r + \gamma) D_{max} \end{aligned}$$

Denoting $w_b = (r + \gamma)D_{max}$, from Theorem 2(a), we note that

$$w(J(t), I(t)) = w(J_0, I_0) < w_b$$
$$\Rightarrow f\big(x(t)\big) < \frac{w_b}{\gamma},\ f\big(y(t)\big) < \frac{w_b}{r}.$$

Given the nature of $f(s)$, $f(s) < c \Rightarrow s_{min} < s < s_{max}$ where (s_{min}, s_{max}) are the finite-valued roots of $f(s) = c$. Hence $x(t)$ and $y(t)$ are both bounded on either side. Consequently, $(x(t) - 1, y(t) - 1)$ is confined to a bounded rectangle and thus,

$$\begin{aligned} \Rightarrow ||\phi\big(x(t), y(t)\big) - \mathbf{z}^*||^2 &= \big(x(t) - 1\big)^2 \\ &\quad + \big(y(t) - 1\big)^2 \\ &< \epsilon^2 \end{aligned}$$

where ϵ can be directly expressed in terms of δ and vice versa. Hence, from Definition 1, $\mathbf{z}^* = (1, 1)$ (or equivalently (J^*, I^*)) is a stable equilibrium. □

Proof of Theorem 1(b).
When initial state (J_0, I_0) is at equilibrium $(J^*, I^*) = (\gamma N, r/e)$, we have $(J(t), I(t)) = (\gamma N, r/e),\ \forall t$. Hence,

$$\dot{S} = -\frac{\beta S I^*}{N} = -\frac{J^* I^*}{N} = -\gamma I^*$$
$$\Rightarrow S(t) = S_0 - \gamma I^* t.$$

Similarly,

$$\dot{R} = \gamma I^* \Rightarrow R(t) = R_0 + \gamma I^* t,$$
$$\beta(t) = \frac{J(t)}{S(t)} = \frac{\gamma N}{S_0 - \gamma I^* t}.$$

At $t = T_{end}$, the susceptible population $S(t) = 0$. Hence,

$$S(T_{end}) = 0 \Rightarrow T_{end} = \frac{S_0}{\gamma I^*}.$$

□

Proof of Theorem 2(a).
The energy function of the LVSIR system in Figure 4(c) corresponds to

$$w(J, I) = \gamma\left(\frac{J}{J^*} - \log\left(\frac{J}{J^*}\right) - 1\right) + r\left(\frac{I}{I^*} - \log\left(\frac{I}{I^*}\right) - 1\right),$$

and the dynamics of I, J are given by

$$\dot{J} = (r - eI)J = r\left(1 - \frac{I}{I^*}\right)J$$

$$\dot{I} = \left(\frac{J}{N} - \gamma\right)I = \gamma\left(\frac{J}{J^*} - 1\right)I.$$

Considering the time derivative of $w(J,I)$, we have

$$\begin{aligned}\dot{w}(J,I) &= \gamma\left(\frac{\dot{J}}{J^*} - \frac{\dot{J}}{J}\right) + r\left(\frac{\dot{I}}{I^*} - \frac{\dot{I}}{I}\right)\\ &= \gamma\dot{J}\left(\frac{1}{J^*} - \frac{1}{J}\right) + r\dot{I}\left(\frac{1}{I^*} - \frac{1}{I}\right)\\ &= \gamma r J\left(1 - \frac{I}{I^*}\right)\left(\frac{1}{J^*} - \frac{1}{J}\right)\\ &\quad + r\gamma I\left(\frac{J}{J^*} - 1\right)\left(\frac{1}{I^*} - \frac{1}{I}\right)\\ &\quad \text{(substituting for } \dot{I},\ \dot{J})\\ &= r\gamma\frac{(I^* - I)(J - J^*)}{I^*J^*}\\ &\quad + r\gamma\frac{(I - I^*)(J - J^*)}{I^*J^*}\\ &= 0\end{aligned}$$

Hence, $w(J,I)$ remains invariant throughout and is equal to $w(J_0, I_0) = w_0$. □

Proof of Theorem 2(b):
Let $w_0 = w(J_0, I_0)$ be the energy associated with the modified SIR system in Figure 4(c). The conservation law implies that every valid state (J,I) corresponds to a point on the level curve given by

$$w(J,I) = \gamma\left(\frac{J}{J^*} - \log\left(\frac{J}{J^*}\right) - 1\right) + r\left(\frac{I}{I^*} - \log\left(\frac{I}{I^*}\right) - 1\right) = w_0.$$

If J, I functions are continuous[5], then these would be periodic functions. In terms of normalized variables, $x = \frac{J}{J^*}$ and $y = \frac{I}{I^*}$, the phase plot reduces to

$$\gamma(x - \log(x) - 1) + r(y - \log(y) - 1) = w_0$$

.

Consider the continuously differentiable function $f(z) = z - \log(z) - 1$ defined on $\mathbb{R}_{++}$. Since $\frac{df}{dz} = 1 - \frac{1}{z}$ and $\frac{d^2f}{dz^2} = \frac{1}{z^2} > 0$, $f(z)$ is a convex function with a single global minimum at $z = 1$ corresponding to $f(1) = 0$. Hence $f(z) \le c \Rightarrow z_{min} \le z \le z_{max}$ where (z_{min}, z_{max}) correspond to the roots of $f(z) = c$.

To identify the extreme x values, we observe that

$$\begin{aligned}\gamma(x - \log(x) - 1) + r(y - \log(y) - 1) &= w_0\\ \Rightarrow \gamma(x - \log(x) - 1) &\le w_0\\ (\text{ since } f(y) > 0)&\\ \Rightarrow x_{min} \le x &\le x_{max}\end{aligned}$$

where (x_{min}, x_{max}) are roots of $f(x) = \frac{w_0}{\gamma}$. Both the extreme values of x are realized for $y = 1$. Similarly, the extreme values of y are attained for $x = 1$ and given by (y_{min}, y_{max}) which correspond to the roots of $f(x) = \frac{w_0}{r}$.

□

Proof of Theorem 2(c).
The period of a Lotka-Volterra system has been derived in multiple works [58]. We include the below proof based on Hsu's method [31] for completeness.

Let $x = \frac{J}{J^*}$ and $y = \frac{I}{I^*}$. Then we have

$$\dot{x} = -rx(y-1) \tag{5}$$

$$\dot{y} = \gamma y(x-1). \tag{6}$$

From 5, we have

$$\begin{aligned}\ddot{x} &= -rx\dot{y} - r\dot{x}(y-1)\\ &= -rx\gamma y(x-1) - r(y-1)\dot{x} && \text{(substituting } \dot{y} \text{ from 6)}\\ &= -r\gamma x(-\frac{\dot{x}}{rx} + 1)(x-1) + \frac{r(\dot{x})^2}{rx} && \text{(substituting y from 5)}\\ &= -\gamma(rx - \dot{x})(x-1) + \frac{\dot{x}^2}{x}\end{aligned}$$

Thus,

$$\ddot{x} - \frac{\dot{x}^2}{x} - \gamma(x-1)(\dot{x} - rx) = 0. \tag{7}$$

Let $z = \log(x)$. Then, $\dot{x} = e^z\dot{z}$ and $\ddot{x} = e^z(\ddot{z} + \dot{z}^2)$.

From 7, we have

$$\begin{aligned}e^z(\ddot{z} + \dot{z}^2) - \frac{e^{2z}(\dot{z})^2}{e^z}&\\ -\gamma(e^z - 1)(e^z\dot{z} - re^z) &= 0\\ \Rightarrow \ddot{z} - \gamma(e^z - 1)(\dot{z} - r) &= 0.\end{aligned}$$

Choosing $s = \dot{z} \Rightarrow s = \frac{\dot{x}}{x} = -r(y-1)$.

Let $w_0 = w(J_0, I_0)$. Then, the trajectory corresponds to

$$\begin{aligned}&\gamma(x - \log x - 1) + r(y - \log y - 1) = w_0\\ &\Rightarrow \gamma(e^z - z - 1) + r(y - \log y - 1) = w_0\\ &\Rightarrow \gamma(e^z - z - 1) + r(-\frac{s}{r} - \log(1 - \frac{s}{r})) = w_0\\ &\Rightarrow \gamma(e^z - z - 1) - w_0 = s + r\log(1 - \frac{s}{r})\\ &\Rightarrow G(z) = F(s)\end{aligned}$$

where $G(z) = \gamma(e^z - z - 1) - w_0$ and $F(s) = s + r\log(1 - \frac{s}{r})$.

Let $F_1(s), F_2(s)$ be the restrictions of $F(s)$ for the lower and upper parts of the phase plot. Then the time period for the lower section is given by

$$\int_{z_{min}}^{z_{max}} \frac{dz}{s} = \int_{z_{min}}^{z_{max}} \frac{dz}{F_1^{-1}(G(z))}.$$

The total time for both the lower and upper section is given by

$$\int_{\log(x_{min})}^{\log(x_{max})} \left(\frac{1}{F_1^{-1}(G(z))} - \frac{1}{F_2^{-1}(G(z))}\right) dz.$$

When $w_0 \simeq 0$, linearization is possible. Simplifying the trajectory $F(s) = G(z)$ using the approximations $e^a = 1 + a + \frac{a^2}{2}$ and $\log(1-a) =$

[5]Note that J, I are actually discrete population counts and not continuous functions.

$a - \frac{a^2}{2}$, we have

$$\begin{aligned}
&\gamma(e^z - z - 1) - w_0 = s + r\log(1 - \frac{s}{r}) \\
&\Rightarrow \frac{\gamma z^2}{2} - w_0 = s + r(-\frac{s}{r} - \frac{s^2}{2r^2}) \\
&\Rightarrow \frac{\gamma z^2}{2} + \frac{s^2}{2r} = w_0
\end{aligned}$$

Essentially, we have an elliptical curve with x, y following sinusoidal behavior with a period $\frac{2\pi}{\sqrt{r\gamma}}$. □

Proof of Theorem 2(d).
Assuming a continuous form for J, we observe that

$$\begin{aligned}
\dot{J} &= (r - eI)J \\
\Rightarrow \frac{1}{J}\frac{dJ}{dt} &= r - eI \\
\frac{d}{dt}(\log(J)) &= r - eI \\
\int_{t=t_0}^{t=t_0+T_{period}} d(\log(J)) &= \int_{t=t_0}^{t=t_0+T_{period}} (r - eI)dt
\end{aligned}$$

(since J is periodic)

$$\int_{t=t_0}^{t=t_0+T_{period}} d(\log(J)) = 0 \text{ (for any } t_0)$$

$$\begin{aligned}
\Rightarrow \int_{t=t_0}^{t=t_0+T_{period}} Idt &= \frac{r}{e}T_{period} \\
&= I^* T_{period}
\end{aligned}$$

In other words, I^* is also the average value of I in each cycle.

Considering the time derivatives of S and I, we have $\dot{I} = \frac{\beta SI}{N} + \gamma I$ and $\dot{S} = -\frac{\beta SI}{N}$.

Let ΔS be the drop in S during a single cycle starting at any t_0, then

$$\begin{aligned}
\Delta S &= \int_{t=t_0}^{t=t_0+T_{period}} \dot{S}dt \\
&= \int_{t=t_0}^{t=t_0+T_{period}} (-\dot{I} + \gamma I)dt \\
&\quad \text{(since } \dot{I} = -\dot{S} + \gamma I) \\
&= 0 + \gamma \int_{t=t_0}^{t=t_0+T_{period}} Idt \\
&\quad \text{(since } I \text{ is periodic)} \\
&= \gamma I^* T_{period} \text{ (from above)}
\end{aligned}$$

□

Proof of Theorem 4.
Assuming a proportional additive control on β of the form $\dot{\beta} = \frac{\beta^2 SI}{N} + (r - eI)\beta + u\beta$, the variation of the susceptible contacts J is given by

$$\begin{aligned}
\dot{J} &= \dot{\beta}S + \beta\dot{S} \\
&= \left(\frac{\beta^2 SI}{N} + (r - eI)\beta + u\beta\right)S \\
&\quad + \beta\left(-\frac{\beta SI}{N}\right) \\
&= (r - eI + u)J
\end{aligned}$$

Let $\mathbf{z} = (J/J^* - 1, I/I^* - 1) = (z_1, z_2)$ so that $\mathbf{z} = (0, 0)$ corresponds to the equilibrium state. Then $w(J, I) = \gamma(z_1 - \log(1 + z_1)) + r(z_2 - \log(1 + z_2))$. For $V(\mathbf{z}) = L(w(J, I), w^*)$ to be a control-Lyapunov function, we require $\dot{V}(\mathbf{z}, \mathbf{u}) < 0$.

$$\begin{aligned}
\dot{V}(\mathbf{z}, \mathbf{u}) &= \langle \nabla V(\mathbf{z}), \dot{\mathbf{z}} \rangle \\
&= \frac{dL}{dw}\langle \nabla w(\mathbf{z}), \dot{\mathbf{z}} \rangle \\
&= \frac{dL}{dw}\left(\frac{dw}{dz_1}\dot{z}_1 + \frac{dw}{dz_2}\dot{z}_2\right) \\
&= \frac{dL}{dw}\left(\gamma\left(1 - \frac{1}{z_1 + 1}\right)\dot{z}_1\right. \\
&\quad \left. + r\left(1 - \frac{1}{z_2 + 1}\right)\dot{z}_2\right) \\
&= \frac{dL}{dw}\left(\gamma\left(\frac{J}{J^*} - 1\right)u\right).
\end{aligned}$$

When the control is chosen as $u = -\eta(t)\frac{dL}{dw}(\frac{J}{J^*} - 1)$ and $\eta(t) > 0,\ \forall t$,

$$\begin{aligned}
\dot{V}(\mathbf{z}, \mathbf{u}) &= \frac{dL}{dw}\left(\gamma\left(\frac{J}{J^*} - 1\right)\right) \\
&\quad \times \left(-\eta(t)\frac{dL}{dw}\left(\frac{J}{J^*} - 1\right)\right) \\
&= -\eta(t)\gamma\left(\frac{dL}{dw}\left(\frac{J}{J^*} - 1\right)\right)^2 \\
&< 0 \text{ (unless } w = w^*)
\end{aligned}$$

Hence, Artstein's theorem guarantees convergence to the equilibrium. □