medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

CoSIR: Optimal control of SIR epidemic dynamics by mapping to Lotka-Volterra System

Harsh Maheshwari *
Department of Computer Science
Georgia Insitute of Technology
harshm121@gmail.com

Shreyas Shetty *
Flipkart Internet Pvt. Ltd.
shreyas.shetty@flipkart.com

Nayana Bannur
Wadhwani AI

Srujana Merugu
Independent

Abstract

Multiple macro-phenomena such as disease epidemics, online information propagation, and economic activity can be well-approximated using simple dynamical systems. Shaping these phenomena with adaptive control of key levers has long been the holy grail of policymakers. In this paper, we focus on optimal control of transmission rate in epidemic systems following the widely applicable SIR dynamics. 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 behavior. We exploit this mapping using a control-Lyapunov approach to design a novel adaptive control policy (CoSIR) that nudges the SIR model to the desired equilibrium. Combining CoSIR policy with data-driven estimation of parameters and adjustments for discrete transmission levels yields a control strategy with practical utility. Empirical comparison with periodic lockdowns on simulated and real COVID-19 data demonstrates the efficacy and adaptability of our approach.

1 Introduction

The COVID-19 epidemic, proliferation of fake news, and increasing economic inequality all highlight the necessity and the enormous challenges of controlling social macro systems. Often, these systems can be characterized by simple dynamics parameterized along certain factors. Optimal adaptive control of such factors is a critical problem for policy makers with massive societal consequences.

Management of a pandemic is a prime example of such a control problem with four key levers: (a) contact restrictions, (b) testing, tracing and isolation, (c) provisioning for medical capacity, and (d) vaccinations. Of these, contact restrictions is the most powerful policy instrument, especially in low and middle income countries facing significant resource constraints. However, choosing the optimal restrictions is highly non-trivial not only because of the complex trade-off between the disease impact and socioeconomic disruptions but also due to the rapidly evolving situation on the ground.

Public health interventions related to COVID-19 have largely been driven by scenario-based epidemiological forecasting [Ferguson et al., 2020, Ray et al., 2020] with decision-making based on comparison of a limited number of scenarios over a finite time horizon. Though valuable, this

*Equal Contribution

Preprint. Under review.

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

Table 1: Variants of the relevant SIR models and transmission rate β.

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

approach leans towards a reactive role for the health authorities. In contrast, despite the potentially far-reaching impact, relatively less attention has been devoted to developing an analytic control framework to support proactive decision-making based on the target disease and economic outcomes, and the state of the epidemic. Multiple studies [Chowdhury et al., 2020, Killian et al., 2020, Bin et al., 2021] point to the benefits of periodic lockdowns, but these interventions are based on simulations over restricted scenarios and are not adaptive in nature. Some recent works [Acemoglu et al., 2020] formulate the control problem in terms of net socioeconomic and disease impact but require an expensive numerical solution. Recent reinforcement learning-based formulations [Bhardwaj et al., 2020] are efforts in similar direction but do not fully exploit the mathematical structure of the epidemic dynamics.

In this paper, we explore optimal adaptive control of transmission rate for a desired bound on infectious population. Multiple studies [Chen et al., 2020, Ray et al., 2020] indicate that time-varying SIR and SEIR [Hethcote, 2000] compartmental models permit parsimonious encoding of infection dynamics, data-driven calibration, and accurate forecasts leading to wide adoption of these models. Hence, we focus on the SIR model [Kermack and McKendrick, 1927] for the control analysis.

Contributions:

1. [**Section 4**] We show that the SIR dynamics maps to the well-known Lotka-Volterra (LV) system [Baigent, 2016] on interpreting infectious patients as predators and susceptible contacts (the product of transmission rate and susceptible population) as the prey under a specific transmission rate policy. The resulting system (`LVSIR`) has a stable equilibrium point and an "energy" conservation property. It exhibits a bounded cyclic trend for active infections and a steady decline of susceptibility.

2. [**Section 5**] We exploit the mapping to derive optimal control policy for transmission rate (`CoSIR`) using control-Lyapunov functions [Tsinias, 1988] based on the "LV system energy". The policy is guaranteed to converge to the desired equilibrium (target infection level) from any valid initial state assuming stochastic perturbations. [2]

3. [**Section 6**] We propose a hybrid practical strategy that combines `CoSIR` with data-driven estimation of parameters and restriction to discrete levels (`CoSIR-approx`).

4. [**Section 7**] Empirical results on real and synthetic COVID-19 data demonstrate the efficacy of the CoSIR approach in stabilizing infections and adaptability to perturbations.

Table 1 lists the relevant SIR models for clarity. Our work points to a general control strategy: (a) mapping a real-world phenomenon to a well-behaved Hamiltonian system [Nutku, 1990] amenable for control, (b) designing optimal control mechanisms for levers of interest using control theory, and (c) employing learning methods to estimate the other model parameters from data. Other application domains such as online information propagation could benefit from a similar hybrid approach.

Notation: x_t and $x(t)$ both denote the value of x at time t; $\dot{x} = \frac{dx}{dt}$ denotes the time derivative; $[x_t]$, $[t]_{t_1}^{t_2}$ denotes the series from $t = t_1$ to $t = t_2$.

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 balance the disease and socioeconomic burdens. Strict short-term lockdowns suppress the infection levels but infections tend to rise 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.

Modeling the multi-faceted impact of contact restrictions requires accounting for region-specific cultural and economic constructs as well as the available resources, which is highly complex. For

[2] Extensions to other models with an incubation period (e.g., SEIR) as well as control of other variables such as the infectious period are presented in Appendix A.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

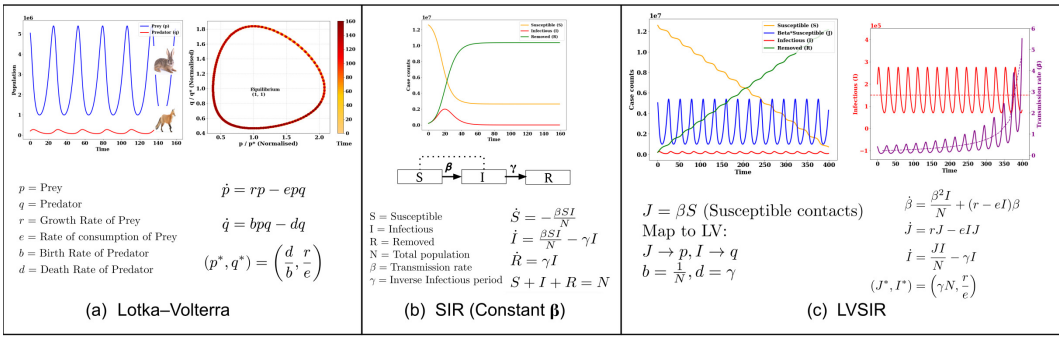

Figure 1: (a) LV system and time evolution of predator and prey populations, (b) SIR dynamics with constant β, (c) SIR model to LV system mapping and the behaviour of case counts S, J, I, R and transmission rate β.[4]

tractability, we assume that the public health goal is to limit active infections to a certain target level determined via an independent impact analysis [BMGF, 2020]. The controls available to the public health authorities can be viewed as configurable knobs (e.g., offices at 20% occupancy). However, the need for clear communication and public compliance entails a simpler strategy centered around a few discrete restriction levels (see Table 6), and a preset schedule for a future time horizon.

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[3] and T, the decision horizon. 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}. In a general case, the target levels could themselves be time-varying functions instead of static values.

Since our goal is to develop an analytical control framework, we make simplifying assumptions on the observability, (i.e., accurate estimation of the infectious population via sero-surveys and diagnostic tests) and the disease dynamics (homogeneous interactions, negligible incubation, constant infectious period). Appendix A describes extensions when some of these assumptions are relaxed.

3 Preliminaries

Compartmental Models. Infectious diseases are often modeled using compartmental models where the population (N) is divided into compartments corresponding to disease stages with transitions governed by the model dynamics. The SIR model [Kermack and McKendrick, 1927] is the simplest and most widely used one comprising three compartments: Susceptible (**S**), Infectious (**I**) and Removed (**R** - includes immune & post-infectious persons) with the dynamics in Figure 1(b). Here β is the rate of disease transmission from infectious to susceptible individuals, which largely depends on the contact restrictions and γ is the inverse of the average infectious period. The effective reproduction number (average number of direct infections from each infection) is $R_{eff} = \frac{\beta S}{\gamma N}$. Existing restriction control approaches [Systrom et al., 2020] are often guided by the principle of ensuring $R_{eff} \simeq 1$.

Lotka-Volterra (LV) Systems. LV systems [Baigent, 2016] model the population dynamics of predator-prey interactions in an ecosystem. In a simple two-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) [5] 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 1(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

[3] Mapping restrictions to transmission is the focus of multiple works [Lehner, 2020] and discussed in Sec 6.
[4] Please see Appendix B and F for proofs and larger figures respectively. Webapp: http://cosir.herokuapp.com/
[5] LV reproduction rate r is different from that of the SIR reproduction number.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

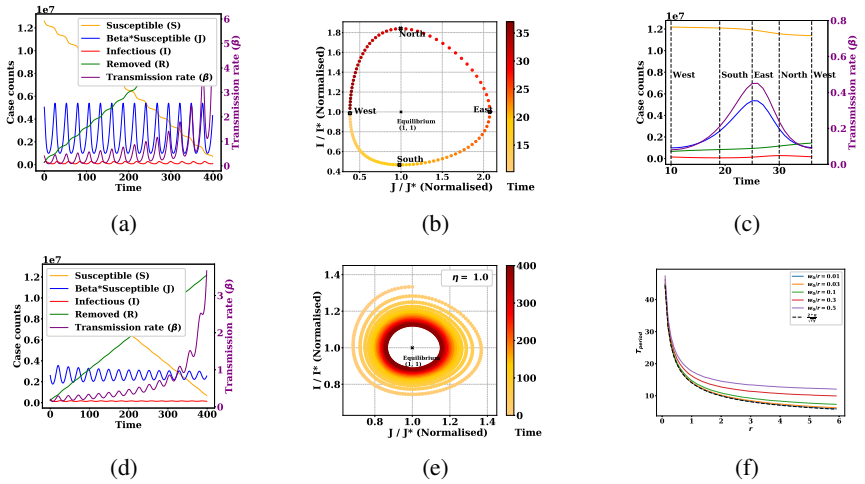

Figure 2: (a-c) System (S, J, I, R, β) evolution when β follows Eqn. 2 [LVSIR] with its phase plot and single cycle annotated with extreme points, (d-e) System evolution when β follows Eqn. 3 [CoSIR] and learning rate $\eta = 1$, (f) Dependence of T_{period} on r for different choices of w_0/r.[2]

considerable research on multiple variants of LV systems [Hening and Nguyen, 2018, Wangersky and Cunningham, 1957] and their Hamiltonian dynamics [Nutku, 1990].

Optimal Control Strategies. Optimal control of dynamical systems Powell [2019] is a well-studied area with connections to reinforcement learning. Given a set of control variables, the optimal control policy describes the time derivatives that minimize the cost function and can be derived using Pontryagin's maximum principle [Powell, 2019] or the Hamilton-Jacobi-Bellman equations [Peng, 1992]. In the case of linear dynamical systems, there are straightforward techniques for optimal control. However, control of non-linear dynamical systems relies heavily on the existence of control-Lyapunov functions often identified using conservation laws of the associated systems, and is non-trivial. Once a suitable Lyapunov function is identified, there exist multiple control design methods such as feedback linearization, backstepping and sliding mode control that are guaranteed to converge using Artstein's theorem [Artstein, 1983]. 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 [Meza et al., 2005, Pchelkina and Fradkov, 2012].

4 Mapping SIR to Lotka-Volterra System (LVSIR)

Solving the restriction policy optimization is our primary goal. We focus on control of SIR dynamics as it captures the core disease spread mechanism of most epidemiological models.

Since existing work [Adda and Bichara, 2012] addresses controllability of SIR dynamics only in the limited context of endemic equilibrium, we adopt a new approach. We first establish a connection between the SIR dynamics and Lotka-Volterra (LV) system and then leverage the LV system properties to propose a strategy for transmission rate control in the SIR system (Section 5).

Stabilizing infection levels 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. Comparing the SIR and LV dynamics in Figure 1, 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

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

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

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

This modified version of SIR model (LVSIR) maps to an LV system. The inverse of the infectious period (γ) corresponds to the predator death rate (d) and the inverse of population $(1/N)$ to the birth rate (b). The LV reproductive rate (r) and prey consumption rate (e) are extra degrees of freedom.

Theorem 1 asserts the existence of an equilibrium for the LVSIR system and the associated behavior.

Theorem 1 [2] *For the LVSIR model in Figure 1(c), the following holds true:*

1. *There exists a stable equilibrium at* $(J^*, I^*) = (\gamma N, r/e)$.
2. *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 susceptible population reaches 0 at* T_{end}.
   (i) $S(t) = S_0 - \gamma I^* t;\ \ R(t) = R_0 + \gamma I^* t$
   (ii) $\beta(t) = \frac{\gamma N}{S_0 - \gamma I^* t};\ \ T_{end} = \frac{S_0}{\gamma I^*}$

Proof sketch. (1) Setting time derivatives to 0 yields the equilibrium points. The stability emerges from the conservation of the "system energy" as defined in Theorem 2(a) and its strict convexity. (2) Invoking the model dynamics (Figure 1(b)) on initial state yields the result. ■

When the LVSIR system is initialized at a non-equilibrium state, it exhibits an oscillatory behavior and a conservation property, which are characterized in Theorem 2 below.

Theorem 2 [2] *For the LVSIR model in Figure 1(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.*

1. *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.*
2. *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* $y - \log(y) = 1 + w_0/r$.
3. *The cyclic period is given by* $T_{period} = \int_{\log(x_{min})}^{\log(x_{max})} (\frac{1}{F_1^{-1}(G(z))} - \frac{1}{F_2^{-1}(G(z))}) dz$,

   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* $T_{period} \simeq \frac{2\pi}{\sqrt{(r\gamma)}}$.
4. *In each cyclic period,* S *reduces by a fixed amount* $\Delta S = \gamma I^* T_{period}$. *When* $S < \Delta S$ *at the start of a cycle, the epidemic terminates during that cycle.*

Proof sketch. (1) Substituting the model dynamics into the energy function yields time-invariance. (2) The closed phase plot emerges from the conservation property with the extreme points determined by the properties of the convex function $f(x) = x - \log(x) - 1$. (3) Derivation of the time period follows a similar analysis for Lotka-Volterra systems [Hsu, 1983]. (4) Fixed drop in S follows from the periodicity and the fact that I^* is the average value of I in each cycle. ■

Figure 2(a) depicts the oscillatory behaviour of the LVSIR model for the hypothetical region in Table 3 . Similar to an LV system, the "energy" which corresponds to a weighted Itakura-Saito distance [Itakura and Saito, 1968] 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, which is desirable from a public health perspective. The transmission rate β also exhibits periodic oscillations but the average steadily goes up to compensate for the reduction in the susceptible population. Figures 2(b-c) show the phase plot and the variation of the key quantities during a single period while Figure 2(f) shows dependence of T_{period} on r and w_0/r.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

Algorithm 1 Practical 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 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} and I_{max}^{target}.
Method:
Data-driven Calibration:
$\gamma \leftarrow \textsc{ComputeGamma}([S_t, I_t, R_t],\ [t]_0^{curr}, N)$
$\rho \leftarrow \textsc{RefineBetaMap}([S_t, I_t, R_t, a_t],\ [t]_0^{curr}, N)$
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.)
$e \leftarrow \frac{r}{I^*}$
$\eta \leftarrow$ DESIRED CONVERGENCE RATE
Determine Optimal Restrictions:
$[\beta_t^{ideal}], [t]_{curr+1}^{curr+T} \leftarrow$
$\textsc{CoSIR}(N, S_{curr}, I_{curr}, R_{curr}, \hat{\beta}_{curr}, r, e, \gamma, \eta, T)$
$a_t = \textsc{argmin}_{a \in \mathcal{A}} (\beta_t^{ideal} - \rho(a))^2,\ [t]_{curr+1}^{curr+T}$
return $\mathbf{a}^* = [a_t],\ [t]_{curr+1}^{curr+T}$

Table 2: Transmission control policies used for simulation.

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 3: State of the epidemic and public health requirements used for the simulations in Section 7 for a hypothetical region (HYP), India (IN) and Netherlands (NL).

Factor	Symbol	HYP	IN	NL
Population	N	13M	1.36B	17.1M
Susceptible Population	S_{curr}	12.6M	1.36B	17.09M
Infectious Population	I_{curr}	0.2M	279k	6.7k
Hospitalization ratio	h	-	2%	6.6%
Target manageable I	I_{avg}^{target}	0.15M	6.8M	0.16M
Target max to avg ratio	$\frac{I_{max}^{target}}{I_{avg}^{target}}$	1.3	1.13	1.13
Target H inflow	ΔH_{avg}^{target}	-	27.5k	2.1k
Periodicity	T_{period}	7 days	7 days	7 days
Inverse Infectious period	γ	0.2	0.2	0.2
LV Reproduction rate	r	4.2	4.2	4.2
Consumption rate	e	r/I^*	r/I^*	r/I^*
Learning rate	η	5	5	5

5 Transmission Rate Control (CoSIR)

Consider an epidemic system from a control theory perspective with the compartmental populations as the system state, the transmission rate as the control input, and the external flows as perturbations. The SIR system has a natural positive feedback loop because of the infection mechanism (since $\dot{S} = \frac{\beta S I}{N}$) which leads to an exponential-like behavior in the early phase of the epidemic. The LVSIR dynamics neutralizes this feedback loop to create oscillatory behavior. We now explore the problem of controlling the transmission rate β 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) defined below. Hence, our immediate goal is to exploit the SIR-LV mapping to derive CLFs based on the "Lotka-Volterra energy" function $w(J, I)$ in Theorem 2.

Definition 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 [1983] 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 [1983]) *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 iff 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 [Artstein, 1983, Tsinias, 1989]. For the current scenario, we consider appropriately defined functions of the LV energy. Let $\mathbf{z} = (\frac{J}{J^*} - 1, \frac{I}{I^*} - 1)$ so that the equilibrium $\mathbf{z}^* = (0, 0)$ and $L(w) : \mathbb{R}_+ \mapsto \mathbb{R}$ be a continuously differentiable function such that $|\frac{dL}{dw}| > 0$ and $L(w) \geq 0,\ \forall w$ and $L(w) = 0 \iff w = 0$. Then, the function $V(\mathbf{z}) = L(w(J, I))$ can be shown to be a CLF for the SIR model and used to construct an affine control policy (CoSIR) as described below.

Theorem 4 (CoSIR) [2] *For the SIR model, a proportional additive control on β of the form,*

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

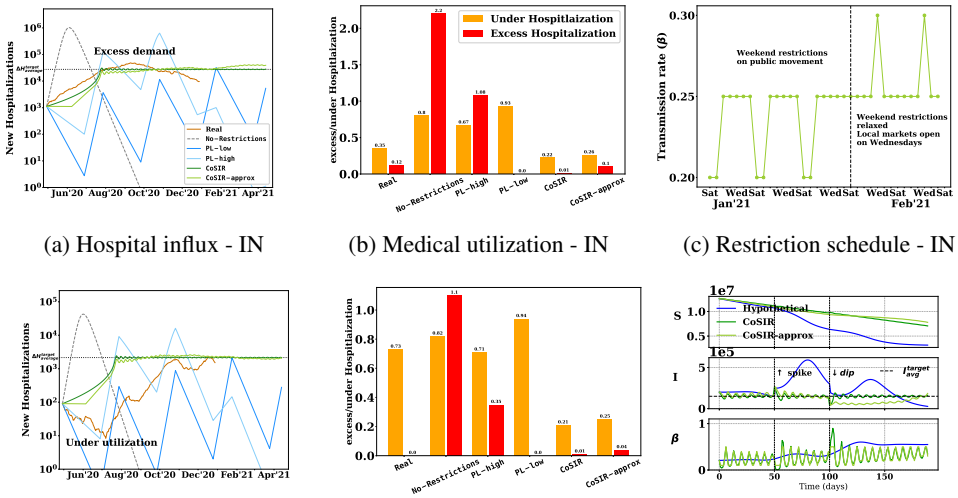

(a) Hospital influx - IN (b) Medical utilization - IN (c) Restriction schedule - IN

(d) Hospital influx - NL (e) Medical utilization - NL (f) Adaptability

Figure 3: (a-b, d-e) New hospitalization influx and excess/under hospitalization with various transmission control policies for two regions: India-IN and Netherlands-NL, (c) `CoSIR-approx` schedule (IN) for a 5-week period (see Appendix F for full range) with a possible mapping to real restrictions, (f) Disease progression indicates that `CoSIR` variants adapt to target levels and also to a sudden increase ($t = 50$) or decrease ($t = 100$). [2]

converges to the equilibrium (J^*, I^*) *when* $u = -\eta(t)\frac{dL}{dw}\left(\frac{J}{J^*} - 1\right)$ *with learning rate* $\eta(t) > 0,\ \forall t$.

Proof sketch. Follows from Definition 1 and Theorem 3 using $|\frac{dL}{dw}| > 0$. ■

Pchelkina and Fradkov [2012] proposed a special case of the above construction with $L(w) = \frac{1}{2}w^2$, which is referred to as the speed-gradient method. Figures 2(d-e) show the damping oscillatory behavior of the CoSIR model for this special case, for the hypothetical region in Table 3. 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. Even with external perturbations, the system can recalibrate β, but the convergence properties depend on the nature of perturbations and need to be explored.

6 Practical Transmission Rate Control

Algorithm 1 outlines a practical approach to solve the restriction policy optimization problem (Section 2) using the optimal β-control in Theorem 4. There are four key steps.

Input Collection. Infection level targets (I^{target}_{avg}, I^{target}_{max}), periodicity of the restriction schedule (T_{period}), decision horizon (T) need to be determined based on 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 calibration and policy optimization.

Data-driven Calibration. We then employ SIR calibration methods to estimate the parameter γ (COMPUTEGAMMA) and the state of the epidemic (S_{curr}, I_{curr} R_{curr}) from historical data. There is a huge body of literature in this space [Ray et al., 2020, Bettencourt and Ribeiro, 2008, Baek et al., 2020] based on likelihood maximization, Bayesian optimization and MCMC methods that can be readily employed. Similarly, the restriction level to transmission map (ρ) can be initialized from public health guidelines [COVID-AMP, 2020] and refined using the observed β for past restrictions (REFINEBETAMAP) via online learning or contextual bandit methods [Beygelzimer et al., 2011]. For the sake of brevity, we skip a detailed description of these techniques (more details in Appendix D).

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

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. Equilibrium I^* is chosen to be the same as the average target infectious level, LV reproductive rate (r) is based on the desired cyclic period T_{period} and the ratio of the maximum and average target levels, and lastly the consumption rate (e) is determined by r and I^*. 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 (`CoSIR`) from Eqn. 3 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 (`CoSIR-approx`).

The main utility of the above approach is that it combines estimation of model parameters (ComputeGamma, RefineBetaMap) with the underlying system dynamics to arrive at a practical strategy. Additional aspects relevant for a real-world implementation are discussed in Section 9.

7 Empirical Results

We present simulation results on the COVID-19 pandemic to demonstrate the efficacy and adaptability of the proposed `CoSIR` approaches relative to other widely cited approaches [Chowdhury et al., 2020].

Experimental Setup.
Data. Simulated data (Apr 2020-Apr 2021) from two regions (India, Netherlands) was used along with real COVID-19 data [Dong et al., 2020] from the initial phase of pandemic (till end of Dec 2020) where the SIR dynamics are more applicable due to negligible reinfection and vaccination rates.
Algorithms. We consider five transmission control policies (see Table 2) assuming SIR dynamics. These include a completely non-restrictive policy (`No-Restrictions`), two CoSIR-based policies (`CoSIR`, `CoSIR-approx`), and two baseline policies (`PL-low`, `PL-high`) based on dynamic periodic interventions [Chowdhury et al., 2020]. We also include the real outcome as indicative of the actual public health transmission control that was adopted.
Parameter Choices. Table 3 enumerates the parameters chosen based on COVID-19 characteristics and practical considerations (details in Appendix C). Two critical choices in our simulations are the restriction periodicity which was set to a weekly cadence and the target infectious level I_{avg}^{target} which was chosen based purely on medical capacity constraints due to the absence of other information. We estimate a manageable hospitalization inflow ΔH_{avg}^{target} based on factors such as population and per-capita bed capacity and set $I_{avg}^{target} = \Delta H_{avg}^{target} \gamma h^{-1}$ noting that in the case of SIR dynamics, the new hospitalizations equal $\gamma h I$ where h denotes the fraction of cases requiring hospitalization.

Hospitalization Influx. Figures 3(a,d) show the simulated daily new hospitalizations ($\gamma h I$) while Figures 3(b,e) depict the average excess and under hospital inflow relative to the target rate as follows: *Excess demand* $:= \underset{t}{\text{mean}}\big[\min\big(\frac{\gamma h I(t)}{\Delta H_{avg}^{target}} - 1, 0\big)\big]$, *Under utilization* $:= \underset{t}{\text{mean}}\big[\min\big(1 - \frac{\gamma h I(t)}{\Delta H_{avg}^{target}}, 0\big)\big]$. Real hospitalizations are based on appropriate scaling of the reported active cases accounting for under-reporting (Appendix C) and point to periods where the healthcare system was overwhelmed (India) or the restrictions were more stringent than necessary (Netherlands). The `CoSIR` variants result in nearly steady hospitalization inflow with the smallest deviation from the target levels relative to all other choices including the actual implemented policy. The key takeaway is that optimizing resource utilization requires choosing β based on the constraints and the varying susceptibility.

Feasibility of CoSIR-approx Restriction Policy. Figure 3(c) shows the `CoSIR-approx` β schedule for a 5-week duration comprising two phases each of which has a weekly periodicity. This schedule could potentially be implemented by having tighter weekend restrictions in the initial phase which are then eased in the later phase along with permitting high contact activities (e.g., local market) on a single weekday. Public communication for these multi-week phases is quite feasible.

Adaptability of CoSIR. Figure 3(f) shows the epidemic evolution of the two CoSIR variants along with a hypothetical but a realistic premature relaxation policy. The CoSIR variants not only stabilize infections, but are able to adapt to sudden upward (t = 50) or downward perturbations (t = 100) as in

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

the case of super-spreader events or sudden quarantine restrictions respectively and continue pushing towards the equilibrium. In practice, the adaptation is not seamless and depends on the nature of perturbations and the lags in the data collection and public communication processes.

8 Related work

Most of the existing literature on epidemiological modeling focuses on the following aspects: (a) Design of models to capture the disease dynamics [Harko et al., 2014], (b) Accurate forecasting of future case counts [Ray et al., 2020], (c) Estimation of model parameters corresponding to non-pharmaceutical interventions (NPIs) [Davies et al., 2020] (d) Optimization of public-health policy (especially NPIs such as quarantine and lockdown policy) based on economic impact and disease burden [Acemoglu et al., 2020, Petrie and Masel, 2020]. There is also work [Naji and Mustafa, 2012] in the space of eco-epidemiology that interleaves SIR and LV dynamics to model prey-predator populations when there is an infection spreading among the prey but it is substantially different from the current effort. Our current work is aligned with the last area, often referred to as economic epidemiology, since the objective is identification of optimal NPIs. However, most of the research in this space is focused on the variation of disease dynamics across subgroups using compartmental models and the explicit modeling of the economic and disease impact along with the trade-offs. The formulated optimization problems are often not readily tractable. Due to the subjective nature of the socioeconomic modeling assumptions and the computationally expensive solutions, this work is more suited for static recommendations. Our work, on the other hand, focuses on dynamic adaptive control of the transmission rate assuming simple SIR dynamics with the goal of maintaining a specific (possibly varying) target level of infection. The resulting analytic control policy adjusted for practical considerations can be useful for real-time decision support.

9 Assumptions & Limitations

Real-world utility of the proposed control framework relies on a few key assumptions. The first pertains to the applicability of the SIR dynamics (i.e., ODEs in Figure 1 with β possibly varying with time) and the associated homogeneous mixing property for a given epidemic and region. This has been validated for MMR and Ebola epidemics for certain regions [Yuan et al., 2015] with SEIR model being preferable for COVID-19. The second assumption relates to the existence of a 1-1 mapping between contact restrictions and the transmission, which is somewhat validated by multiple studies [Liu et al., 2021] though the relationship can include delays and stochasticity due to compliance issues. The third assumption is on the nature of perturbations or the external flows since the optimality of control and convergence results are guaranteed only for the ideal β control with normal stochastic perturbations. Adaptability under arbitrary/adversarial perturbations remains to be explored. Effective implementation in a practical public health setting similar to the dynamic protocol deployed in California [CA, 2020] also depends heavily on ensuring (a) appropriateness and availability of a target (possible time-varying) infection level through an independent impact analysis, (b) robust data-driven estimation of model parameters (e.g., γ) as well as the restriction to β mapping, (c) adjusting for the lags and under-reporting in the data collection process, and (d) accounting for delays in public communication systems and varying compliance. An important limitation of our current work is that the efficacy of control strategies could only be demonstrated via generic SIR-based simulations since real experiments are impractical due to the associated human impact. As more case and restriction data becomes available, we plan to perform off-policy evaluation.

10 Conclusion & Future Work

We propose an analytical framework for epidemic control with the intent of supporting an active goal-oriented public health response. This framework relies on a mapping between SIR dynamics to Lotka-Volterra system under a specific transmission rate policy (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 compartmental models and additional control variables (Appendix A). Results in diverse settings point to the utility of this approach. Qualitative insights such as weekly restriction patterns were employed in a regional COVID-19 response. Regret bounds for a non-stochastic control setting [Hazan et al., 2020] with arbitrary bounded perturbations is a promising direction to explore. Additionally, we could investigate control strategies for applications such as online information propagation and regulation of occupational groups that exhibit Hamiltonian dynamics [Nutku, 1990].

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

References

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

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.

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

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. *arXiv preprint arXiv:2006.06373*, 2020.

Stephen A Baigent. Lotka-Volterra dynamics: an introduction. World Scientific, 2016.

Yinon M. Bar-On, Ron Sender, Avi I. Flamholz, Rob Phillips, et al. A quantitative compendium of COVID-19 epidemiology, 2020.

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.

Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, et al. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 19–26, 2011.

Ankit Bhardwaj, Han Ching Ou, Haipeng Chen, Shahin Jabbari, et al. Robust lock-down optimization for COVID-19 policy guidance. In *AAAI Fall Symposium*, 2020.

Michelangelo Bin, Peter Y. K. Cheung, Emanuele Crisostomi, Pietro Ferraro, Hugo Lhachemi, Roderick Murray-Smith, Connor Myant, Thomas Parisini, Robert Shorten, Sebastian Stein, and Lewi Stone. Post-lockdown abatement of COVID-19 by fast periodic switching. *PLOS Computational Biology*, 17:1–34, 01 2021. doi: 10.1371/journal.pcbi.1008604.

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

W.E. Boyce, R.C. DiPrima, and D.B. Meade. *Elementary Differential Equations and Boundary Value Problems*. Wiley, 2017. ISBN 9781119443766.

COVID-19 CA. Tracking COVID-19 in CA, 2020. URL `https://covid19.ca.gov/state-dashboard/`.

Medical capacity. List of countries by hospital beds, Dec 2020. URL `https://en.wikipedia.org/wiki/List_of_countries_by_hospital_beds`.

Yi-Cheng Chen, Ping-En Lu, Cheng-Shang Chang, and Tzu-Hsuan Liu. A time-dependent SIR model for COVID-19 with undetectable infected persons. *IEEE Transactions on Network Science and Engineering*, 7(4):3279–3294, Oct 2020. ISSN 2334-329X. doi: 10.1109/tnse.2020.3024723. URL `http://dx.doi.org/10.1109/TNSE.2020.3024723`.

Rajiv Chowdhury, Kevin Heng, Md Shajedur Rahman Shawon, 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.

COVID-AMP. Analysis and mapping of policies - COVID-AMP, 2020. URL `https://covidamp.org/`.

COVID19India. Coronavirus in India: Latest Map and Case Count, 2020. URL `https://www.covid19india.org/`.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

Nicholas G Davies, Adam J Kucharski, Rosalind M Eggo, Amy Gimma, 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*, 5(7):e375–e385, 2020. ISSN 2468-2667. doi: https://doi.org/10.1016/S2468-2667(20)30133-X. URL `https://www.sciencedirect.com/science/article/pii/S246826672030133X`.

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. License CC by 4.0.

Neil Ferguson, Daniel Laydon, Gilani G Nedjati, N Imai, K Ainslie, et al. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Technical report, Imperial College, London, 2020.

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. ISSN 0096-3003. doi: 10.1016/j.amc.2014.03.030".

Elad Hazan, Sham M. Kakade, and Karan Singh. The Nonstochastic Control Problem. *arXiv preprint arXiv:1911.12178*, 2020.

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.

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

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.

Fumitada Itakura and Shuzo Saito. Analysis synthesis telephony based upon the maximum likelihood method. In *International Congress on Acoustics*, pages 17–20, 1968.

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.

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.

Jackson A. Killian, Marie Charpignon, Bryan Wilder, Andrew Perrault, et al. 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.

Daniel B Larremore, Bryan Wilder, Evan Lester, Soraya Shehata, et al. Test sensitivity is secondary to frequency and turnaround time for COVID-19 surveillance. *MedRxiv*, 2020. doi: 10.1101/2020.06.22.20136309.

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

Yang Liu et al. The impact of non-pharmaceutical interventions on SARS-CoV-2 transmission across 130 countries and territories. *BMC Medicine*, 19(1):40, 2021. doi: 10.1186/s12916-020-01872-8. URL `https://doi.org/10.1186/s12916-020-01872-8`.

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.

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

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(2):124–135, 2005.

Siuli Mukhopadhyay and Debraj Chakraborty. Estimation of undetected COVID-19 infections in India. *medRxiv*, 2020. doi: 10.1101/2020.04.20.20072892.

Manoj V Murhekar, Tarun Bhatnagar, Sriram Selvaraju, Kiran Rade, 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.

Raid Kamel Naji and Arkan N Mustafa. The Dynamics of an Eco-Epidemiological Model with Nonlinear Incidence Rate. *J. Appl. Math.*, 2012:852631, 2012. ISSN 1110-757X. doi: 10.1155/2012/852631.

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

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.

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

James Petrie and Joanna Masel. Marginal value of quarantine. *medRxiv*, 2020. doi: 10.1101/2020.11.24.20238204.

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

Warren B Powell. From reinforcement learning to optimal control: A unified framework for sequential decisions, 2019.

Evan L Ray, Nutcha Wattanachit, Jarad Niemi, Abdul Hannan Kanji, et al. Ensemble forecasts of Coronavirus disease 2019 (COVID-19) in the U.S. *medRxiv*, 2020. doi: 10.1101/2020.08.19.20177493.

Shagi-Di Shih. The period of a Lotka-Volterra system. *Taiwanese Journal of Mathematics*, 1(4):451–470, 1997. ISSN 10275487, 22246851.

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

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

John Tsinias. Sufficient Lyapunov-like conditions for stabilization. *Mathematics of Control, Signals and Systems*, 2(4):343–357, 1989. ISSN 1435-568X. doi: 10.1007/BF02551276.

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

Worldometer. Worldometer, 2021. URL `https://www.worldometers.info/`.

Sanling Yuan, Amira Rachah, and Delfim F. M. Torres. Mathematical Modelling, Simulation, and Optimal Control of the 2014 Ebola Outbreak in West Africa. *Discrete Dynamics in Nature and Society*, 2015:842792, 2015. doi: 10.1155/2015/842792. URL `https://doi.org/10.1155/2015/842792`.

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.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

Appendix

A Extensions

Delayed SIR & SEIR Models. Certain infectious diseases have a significant incubation period when the individuals are infected but are not spreading the disease (non-infectious). The SEIR model [Hethcote, 2000] includes an additional **E** (exposed) compartment to model the incubation phase. This model is known to closely mimic the behaviour of the delayed SIR model [Kaddar et al., 2011]. 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 [Martin and Ruan, 2001, Wangersky and Cunningham, 1957]. 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

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

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 [Bar-On et al., 2020, Larremore et al., 2020]. This is analogous to the culling of predators (infectious population) by increasing the death rate for which there already exists multiple control mechanisms [Meza et al., 2005]. In particular, choosing $V(\mathbf{z}) = L(w(J, I))$ as the CLF of interest, we obtain the control, $\gamma(t) = \gamma_0 + \zeta(t)\frac{dL}{dw}(\frac{I}{I^*} - 1)$, with $\zeta(t) > 0$.

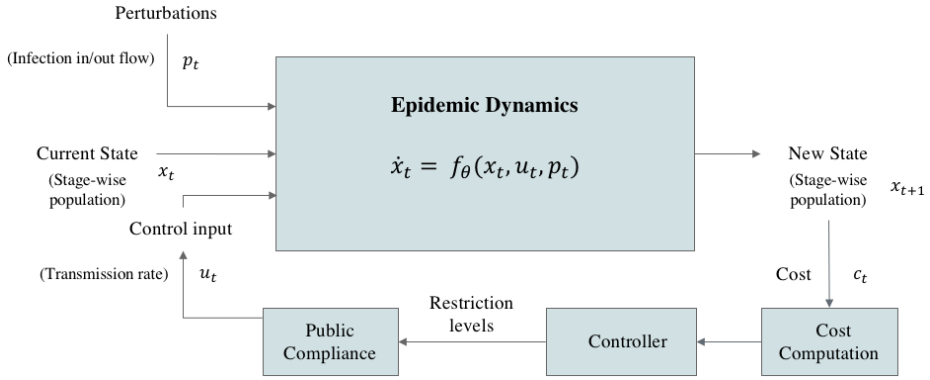

Figure 4: The control theory perspective of the epidemic dynamics.

Non-stochastic Control. Figure 4 depicts the epidemic system from a control perspective. Theoretical guarantees of optimal control based on control Lyapunov functions rely on the assumptions of (a) deterministic transition dynamics, (b) stochastic perturbations $p(t)$, and (c) exact control. In practice, however, the parameters of the SIR model such as the inverse infectious period (γ) vary with time, the perturbations due to in/out flow might be arbitrary, and the control is approximate because of public communication and compliance issues. To address these constraints, we need to continuously re-estimate the model parameters as well as the mapping between the restriction levels and β from historical data. In this online learning setting, the restriction levels can be viewed as bandit arms each associated with a specific transmission rate (or probability distribution) and compartmental populations correspond to the state. The arm-specific transmission rates are estimated in each step

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

using observations based on an appropriate loss function. When the loss function is defined in terms of the Lotka-Volterra energy over a time horizon, this formulation can be viewed as a non-linear contextual bandit [Beygelzimer et al., 2011] and the properties of the LV energy permit derivation of regret bounds similar to the non-stochastic control framework for linear dynamical systems [Hazan et al., 2020].

B Proofs

Definition 2 *[Boyce et al., 2017] 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 1.1
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 2.1, we note that

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

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 2, $\mathbf{z}^* = (1, 1)$ (or equivalently (J^*, I^*)) is a stable equilibrium. ■

Proof of Theorem 1.2

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,

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

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 2.1

The energy function of the LVSIR system in Figure 2(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

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

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

Proof of Theorem 2.2

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

$$\begin{aligned}
w(J, I) &= \gamma \left(\frac{J}{J^*} - \log\left(\frac{J}{J^*}\right) - 1 \right) \\
&\quad + r \left(\frac{I}{I^*} - \log\left(\frac{I}{I^*}\right) - 1 \right) \\
&= w_0
\end{aligned}$$

If J, I functions are continuous[6], 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

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

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

.

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 $f(z) \leq c \Rightarrow z_{min} \leq z \leq z_{max}$ where (z_{min}, z_{max}) correspond to the roots of $f(z) = c$.

To identify the extreme x values, we observe that

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

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 2.3

The period of a Lotka-Volterra system has been derived in multiple works Shih [1997]. We include the below proof based on Hsu's method Hsu [1983] for completeness.

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

$$\dot{x} = -rx(y - 1) \tag{4}$$

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

From 4, we have

$$\begin{aligned} \ddot{x} &= -rx\dot{y} - r\dot{x}(y - 1) \\ &= -rx\gamma y(x - 1) - r(y - 1)\dot{x} && \text{(substituting } \dot{y} \text{ from 5)} \\ &= -r\gamma x(-\frac{\dot{x}}{rx} + 1)(x - 1) + \frac{r(\dot{x})^2}{rx} && \text{(substituting y from 4)} \\ &= -\gamma(rx - \dot{x})(x - 1) + \frac{\dot{x}^2}{x} \end{aligned}$$

Thus,

$$\ddot{x} - \frac{\dot{x}^2}{x} - \gamma(x - 1)(\dot{x} - rx) = 0. \tag{6}$$

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

From 6, we have

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

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 $\log(1-a) = a - \frac{a^2}{2}$, 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 2.4
Assuming a continuous form for J, we observe that

$$\begin{aligned}
\dot{J} &= (r - eI)J \\
\Rightarrow \frac{1}{J}\frac{dJ}{dt} &= r - eI \\
\frac{d}{dt}(\log(J)) &= r - eI \\
\int_{t=t_0}^{t=t_0+T_{period}} d(\log(J)) &= \int_{t=t_0}^{t=t_0+T_{period}} (r - eI)dt \\
&\text{(since } J \text{ is periodic)} \\
\int_{t=t_0}^{t=t_0+T_{period}} d(\log(J)) &= 0 \text{ (for any } t_0) \\
\Rightarrow \int_{t=t_0}^{t=t_0+T_{period}} I dt &= \frac{r}{e} T_{period} \\
&= I^* T_{period}
\end{aligned}$$

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

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

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

■

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 $V(\mathbf{z}) = L(w(J, I))$ to be a control-Lyapunov function, we require $\dot{V}(\mathbf{z}, \mathbf{u}) < 0$.

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

Factor	**Notation**	**Hypothetical**	**India**	**Netherlands**
Population	N	13M	1.366B	17.1M
Stage of pandemic	(S_{curr}, I_{curr})	(12.6M, 0.2M)	(1.36B, 279k)	(17.09M, 6.7k)
Bed Occupancy level (Phillip et al. [1984])	κ	-	0.76	0.76
Medical Capacity (per capita) (capacity [2020])	p	-	5.3×10^{-4}	3.32×10^{-3}
Hospitalization ratio	h	-	2%	6.6%
Time spent in hospital	T_{hosp}	-	20 days	20 days
Target manageable infections	I_{avg}^{target}	0.15M	6.88M	0.16M
Target max to average ratio	$I_{max}^{target}/I_{avg}^{target}$	1.3	1.13	1.13
Infection-Fatality ratio (Meyerowitz-Katz and Merone [2020], Mukhopadhyay and Chakraborty [2020], Murhekar et al. [2020])	IFR	-	0.002 [7]	0.0068
Death detection ratio	DDR	-	0.3	1
Reported Infections on 1st Dec 2020 (Worldometer [2021], COVID19India [2020])		-	9.5M	527k
Reported Fatalities on 15th Dec 2020 (Worldometer [2021], COVID19India [2020])		-	144k	10k
Under reporting factor	URF	-	25.28	2.83
Periodicity	T_{period}	7 days	7 days	7 days
Reciprocal of infectious period	γ	0.2	0.2	0.2
LV Reproduction rate	r	4.2	4.2	4.2
Consumption rate	e	2.8×10^{-5}	6.1×10^{-7}	2.6×10^{-5}
Initial transmission rate	β_0	0.2	0.2	0.2
Learning rate	η	5	5	5

Table 4: State of the epidemic, public health requirements, and CoSIR parameters used for the simulations in Section 7.

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

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

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

C Parameter Choices

Table 4 enumerates the detailed parameter choices and the relevant assumptions for the different regions that can be broadly divided into three groups. We now discuss the rationale for the various parameter choices.

(a) Target Hospitalization & Infectious levels. There are multiple ways to arrive at an acceptable level for infectious population I_{avg}^{target} based on socioeconomic impact considerations [BMGF, 2020]. When the focus is on medical capacity constraints, this target level can be set chosen so 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

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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

inflow that will not overwhelm the medical infrastructure is given by $\Delta H_{avg}^{target} = \frac{Np\kappa}{T_{hosp}}$. As described earlier in Section 7, assuming SIR dynamics, the new hospitalizations would be given by γhI. 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 capacity [2020] assuming an average hospitalization period $T_{hosp} = 20$ days and hospitalization ratio between $2 - 6.6\%$ based on literature. 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 weeks between infection and the associated fatalities. In particular, we estimate the under reporting factor (URF) as

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

The purpose of computing URF is to figure out the appropriate scaling of real cases for a fair comparison with the simulated results.

(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 $\eta = 5$.

D Additional Experiments

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.

Figures 5(a-b) show the variation of the transmission rate (β) with time for four intervention policies (`CoSIR`, `CoSIR-approx`, `PL-high`, `PL-low`). 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 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(c-d) 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 [Zachary and Mobarak, 2020]. 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

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

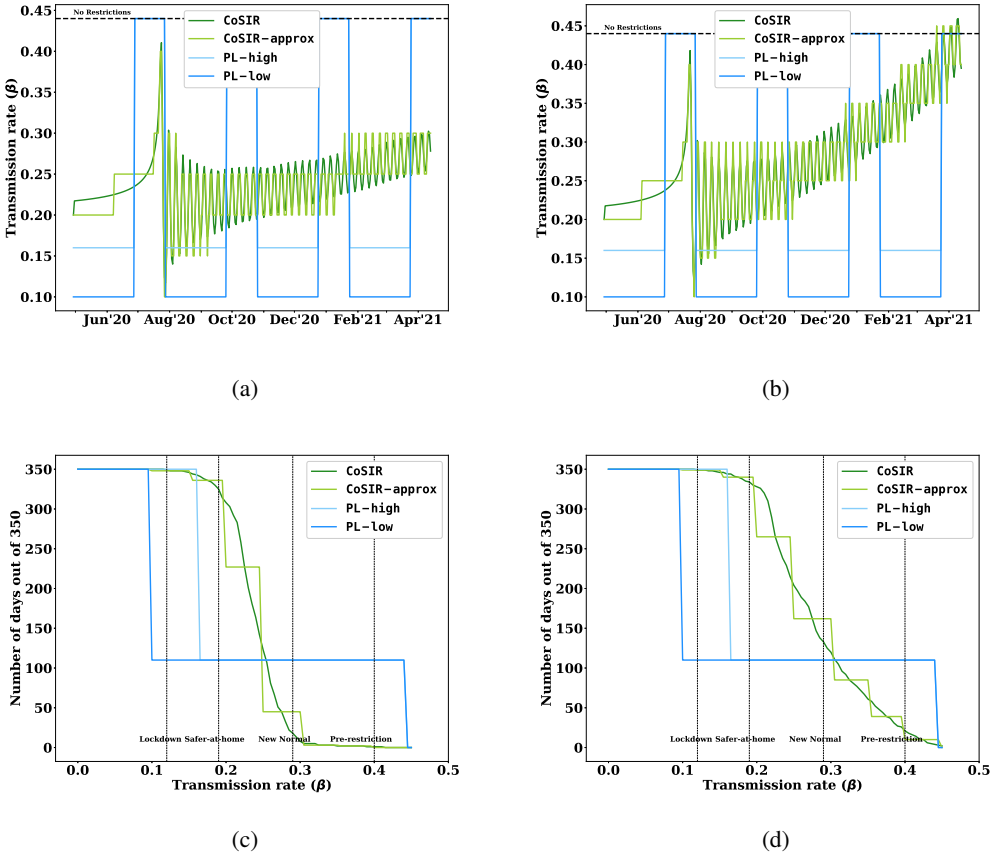

Figure 5: Plots (a-b) depict the variation of policy transmission rate with time. Plots (c-d) show the restrictiveness profile (number of days where the policy is less restrictive than a specified β) for the periodic lockdown and CoSIR-based policies for India and Netherlands

Policy	India	Netherlands
CoSIR-ideal	2026.6	2009.7
CoSIR-approx	2184.6	2133.2
PL-high	6561.9	5416.7
PL-low	6744.9	6826.6

Table 5: Total LV Energy for the entire duration of 350 days corresponding to all policies for India and Netherlands respectively.

β profile are likely to be amplified when the socioeconomic impact is considered. Note that if the simulation were to cover the entire duration of the pandemic, the CoSIR based approaches would turn out to be more relaxed than the fixed periodic lockdowns.

Figure 6 show the LV Energy curve with respect to time for the different policies. The energy, which denotes how far the systtem is from the target values, quickly comes down for the CoSIR variants. Table 5 shows the total energy of each of the policy over the entire duration of 350 days. The total energy of the CoSIR variants is significantly low which signifies that the system remained close to the target levels.

D.1 Empirical results till April 2021

Given the heavy interest in the subsequent waves of infections in India and Europe, in the interest of transparency, we also have included plots showing the real outcome till April (Figure 7). However, as

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

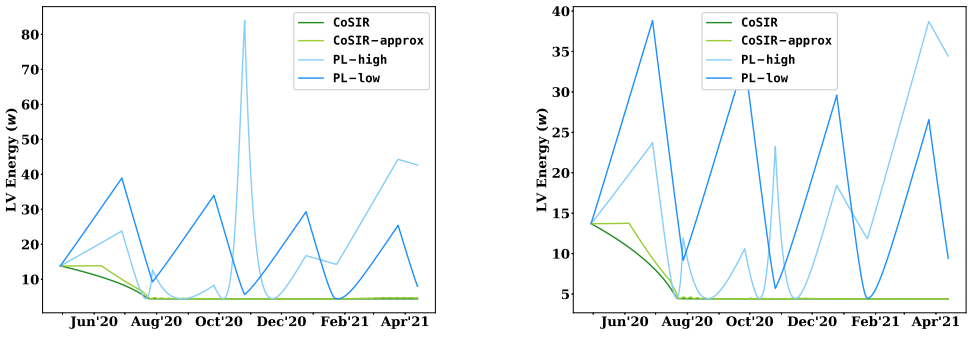

(a) (b)

Figure 6: Plots (a-b) depict the variation of LV Energy with time corresponding to all policies for India and Netherlands respectively.

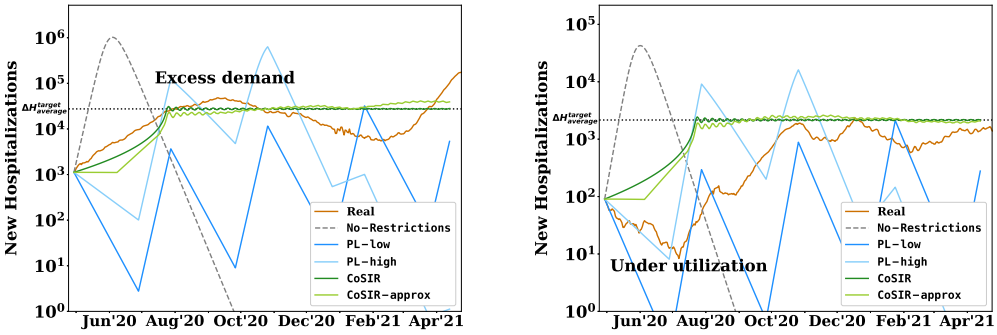

(a) Hospital influx - India (b) Hospital influx - Netherlands

Figure 7: New hospitalization influx with various transmission control policies for two regions: India-IN and Netherlands-NL with real data till April 2021

per the discussion in Section 9, one needs to be mindful that the proposed control policy (CoSIR) applies only when SIR dynamics are valid. In the first phase of the pandemic (most of 2020) when reinfection and vaccination rate was negligible, this is largely true but it not the case for the most recent phase. The right approach is to consider a new compartmental model that incorporates the reinfection and vaccination dynamics and derive the corresponding optimal control, which is a direction that we are exploring.

E Societal Impact

Epidemic and population control have applications in multiple domains such as public health, public communications, viral marketing, regulation of economic activity in multi-sided market places, all of which have an immediate and significant societal impact. Better understanding and controllability of these macro-systems that impact human lives definitely holds the promise of positive outcomes . However, as in the case of generic AI techniques, the proposed control methods create an empowered artificial agency that does entail multiple risks for the society including but not limited to:

- Inappropriate and unsafe use of control techniques due to poor understanding of the limitations and underlying assumptions of the proposed methods (see Section 9) as well as inadequate safeguards to address these constraints, especially in areas related to goal/input specification, data observability, and model dynamics. Such misuse (e.g., an inappropri-

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

ate target infectious level) can happen even with benevolent controllers and the best of intentions.

- Use of control techniques for objectives contrary to societal good such as propagating misinformation on social media, marketing harmful products, and bio-warfare.
- Increase automation of higher order decision that could lead to more efficiency on one end, but also loss of employment and further empowerment of capital relative to human resources.

Both the benefits and risks associated with the proposed techniques are amplified by the fact that the immediate applications such as control of pandemics and online discourse pertain to regulation of macro-systems affecting large human populations.

F Miscellaneous

This section includes miscellaneous content and larger versions of figures for the readers benefit. Table 6 contains an example of restriction levels and the corresponding effective transmission rates (β).

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 6: Example of restriction levels recommended by public health authorities [COVID-AMP, 2020] and the corresponding effective R value and transmission rate (β) for COVID-19.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

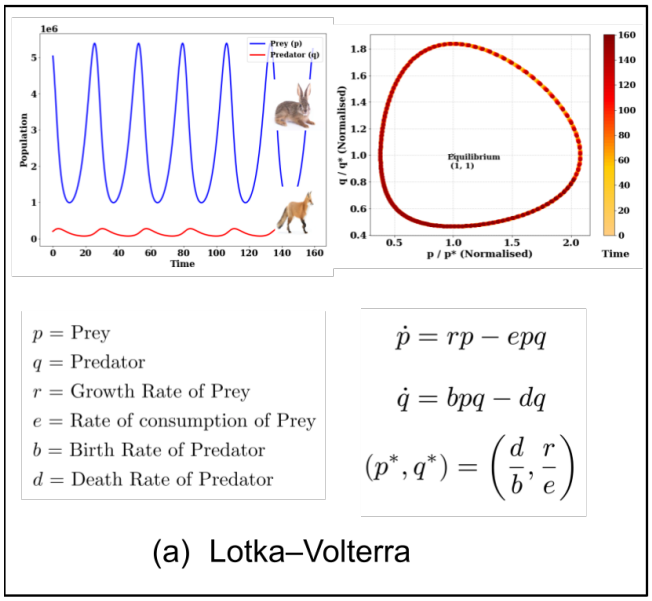

Figure 8: (a) Lotka-Volterra system and evolution of predator (fox) and prey (hare) population over time.

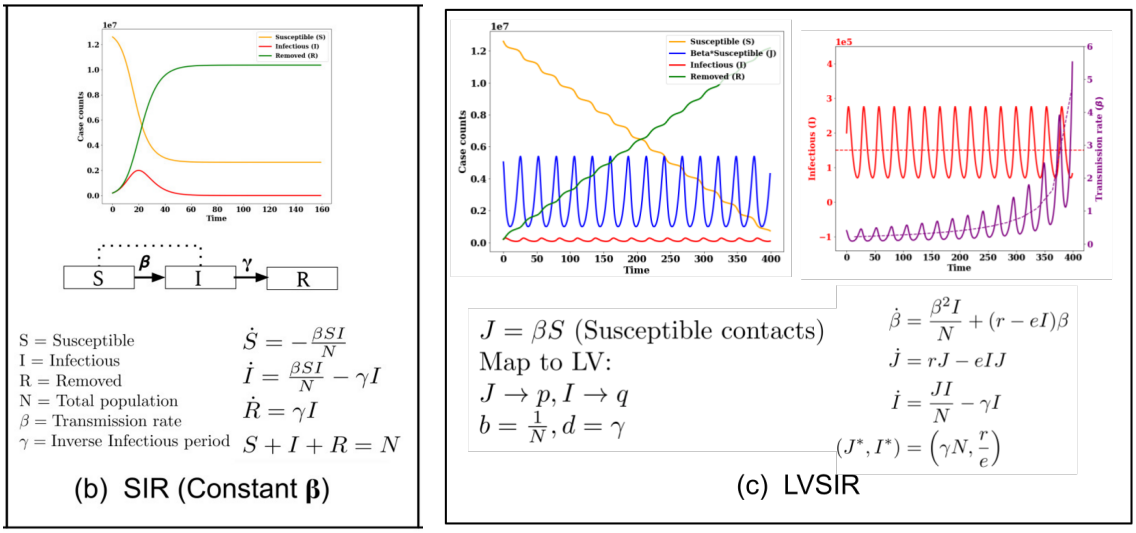

Figure 9: (b) SIR compartmental dynamics for epidemiological modeling. (c) Mapping of SIR model to LV system and the behaviour of case counts S, J, I, R and transmission rate β.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

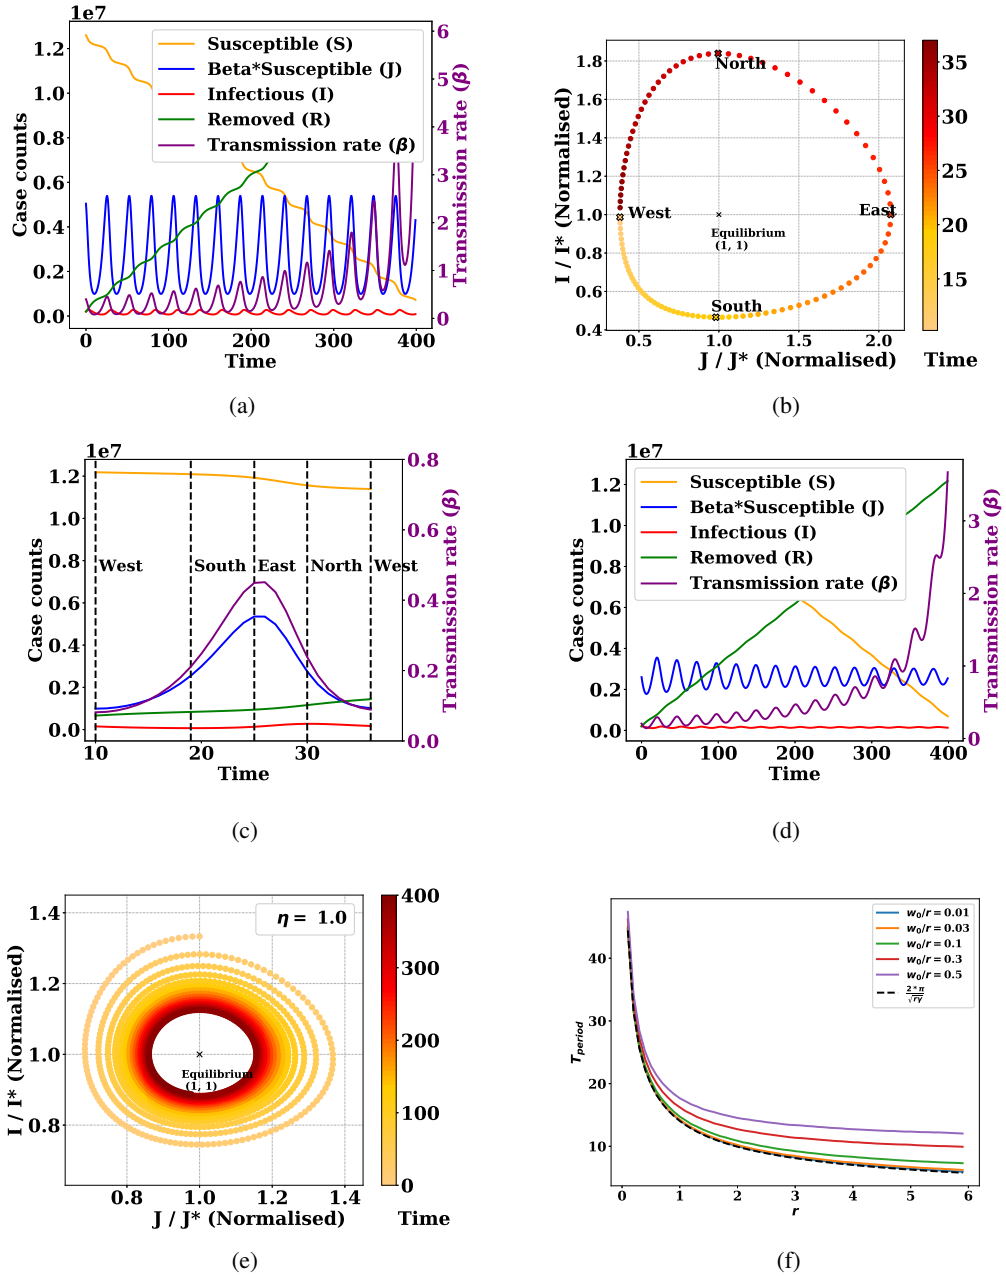

Figure 10: (a-c) System (S, J, I, R, β) evolution when β follows Eqn. 2 [LVSIR] with its phase plot and single cycle annotated with extreme points, (d-e) System evolution when β follows Eqn. 3 [CoSIR] and learning rate $\eta = 1$, (f) Dependence of T_{period} on r for different choices of w_0/r.

medRxiv preprint doi: https://doi.org/10.1101/2020.11.10.20211995; this version posted October 6, 2021. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license .

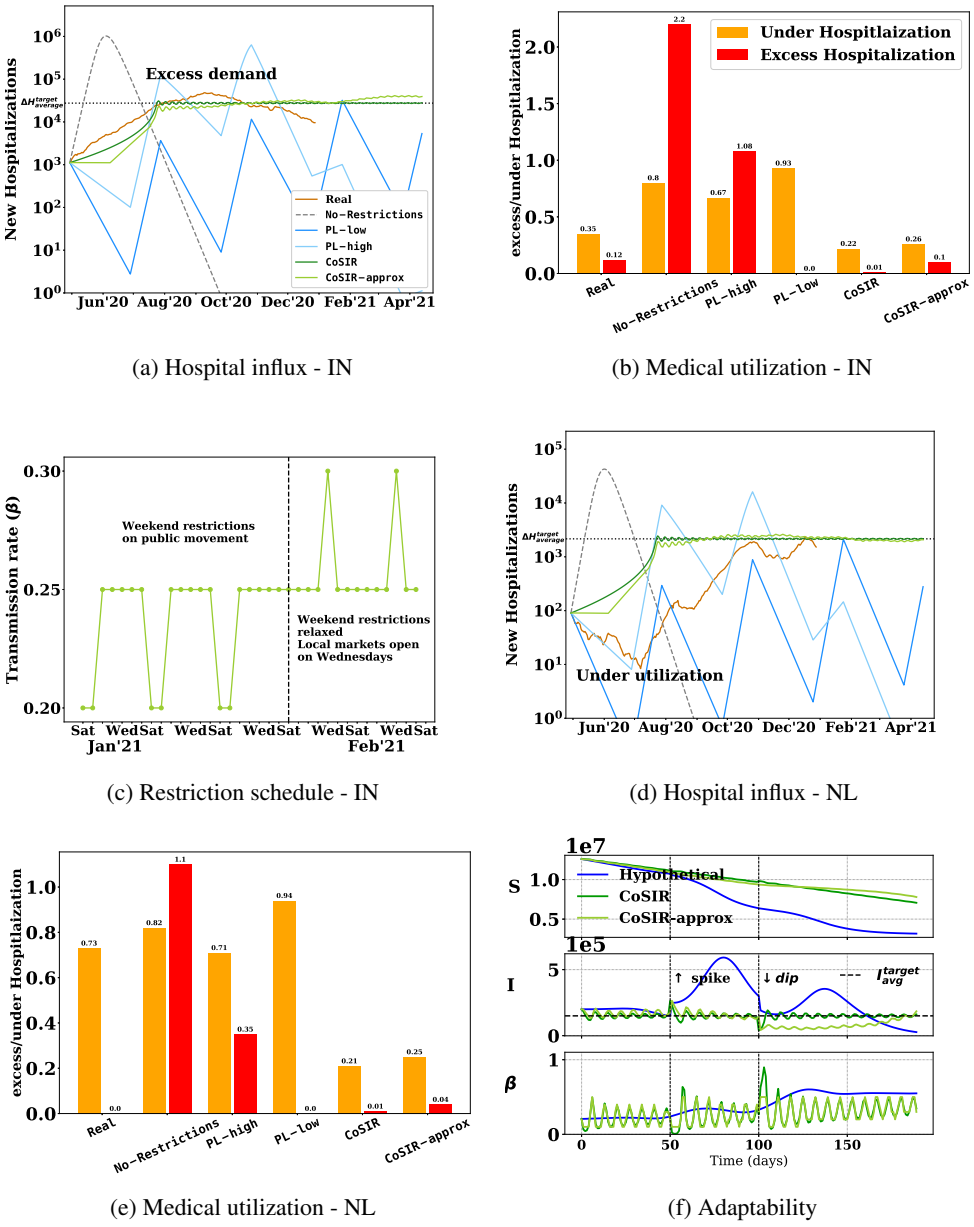

(a) Hospital influx - IN

(b) Medical utilization - IN

(c) Restriction schedule - IN

(d) Hospital influx - NL

(e) Medical utilization - NL

(f) Adaptability

Figure 11: (a-b, d-e) New hospitalization influx and excess/under hospitalization with various transmission control policies for two regions: India-IN and Netherlands-NL, (c) `CoSIR-approx` schedule (IN) for a 5-week period (see Appendix F for full range) with a possible mapping to real restrictions, (f) Disease progression indicates that `CoSIR` variants adapt to target levels and also to a sudden increase ($t = 50$) or decrease ($t = 100$).