



Epidemiologically and Socio-economically Optimal Policies via Bayesian Optimization

Amit Chandak¹ · Debojyoti Dey¹ · Bhaskar Mukhoty¹ · Purushottam Kar¹

Received: 28 April 2020 / Revised: 14 June 2020 / Accepted: 23 June 2020 / Published online: 3 July 2020
© Indian National Academy of Engineering 2020

Abstract

Mass public quarantining, colloquially known as a *lock-down*, is a non-pharmaceutical intervention to check spread of disease. This paper presents ESOP (Epidemiologically and Socio-economically Optimal Policies), a novel application of active machine learning techniques using Bayesian optimization, that interacts with an epidemiological model to arrive at lock-down schedules that optimally balance public health benefits and socio-economic downsides of reduced economic activity during lock-down periods. The utility of ESOP is demonstrated using case studies with VIPER (Virus-Individual-Policy-EnviRonment), a stochastic agent-based simulator that this paper also proposes. However, ESOP is flexible enough to interact with arbitrary epidemiological simulators in a black-box manner, and produce schedules that involve multiple phases of lock-downs.

Keywords Optimal policy · Lock-down · Epidemiology · Bayesian optimization

Introduction

Infectious diseases that are contagious pose a threat to public safety once they attain pandemic status. Several historical instances of such pandemics have taken a heavy toll on human lives. Prominent examples include the H1N1 (Spanish flu) pandemic of 1918 (> 50 million fatalities), the H3N2 (HongKong flu) pandemic of 1968 (\approx 1 million fatalities), the HIV/AIDS pandemic (\approx 32 million fatalities) (Kimball and Bose 2020), the novel influenza-A H1N1 (swine flu) pandemic of 2009 (\approx 0.3 million fatalities) (Roos 2012), and the ongoing CoViD-19 pandemic (\approx 0.45 million fatalities as of writing this document) (WHO 2020).

In such situations, and especially in the absence of vaccines and antiviral treatments, experts often prescribe

guidelines to public such as hand hygiene and respiratory etiquette, as well as two kinds of non-pharmaceutical interventions, namely (1) *mitigation policies* such as human surveillance and contact tracing, and (2) *suppression policies* such as social distancing or its more extreme form colloquially known as a *lock-down* (Aledort et al. 2007). However, their benefits with respect to public health outcomes notwithstanding, severe and extended applications of suppression policies such as lock-downs negatively impact livelihoods and the economy. For instance, Scherbina (2020) estimates the cost of extensive suppression measures to the US economy at \$9 trillion, or about 43% of its annual GDP.

Peak et al. (2017) demonstrate the need for policy decisions to balance suppression and mitigation measures in terms of the epidemiological characteristics of the pandemic, pointing out that suppression measures hold most benefit for fast-course diseases whereas effective mitigation measures may suffice for others at much less socio-economic cost. This points to a need for techniques that can take the disease progression characteristics of a certain outbreak and suggest policies that optimally use suppression and mitigation techniques to offer acceptable health outcomes as well as socio-economic risks within acceptable limits. In fact there is an emerging area variously termed “economic epidemiology” or “epidemiological economics” (Perrings et al. 2014)

Disclaimer This paper makes no recommendation to individuals and its results should not be interpreted by individuals to modulate personal behavior. The authors recommend that individuals continue to follow guidelines offered by local governments with respect to lock-downs and social distancing, and those offered by medical professionals with respect to personal hygiene and treatment.

✉ Purushottam Kar
purushot@cse.iitk.ac.in

¹ Indian Institute of Technology Kanpur, Kanpur, India

that seeks to develop models that can address the interplay of disease and host behavior.

Our Contributions

This paper presents ESOP (Epidemiologically and Socio-economically Optimal Policies), a system that uses Bayesian optimization to automatically suggest suppression policies that optimally balance public health and economic outcomes. ESOP interacts with epidemiological and economic models to automatically suggest policy decisions. The paper also presents VIPER (Virus-Individual-Policy-EnviRonment), an iterative, stochastic agent-based model (ABM) with which case studies are conducted to showcase the utility of ESOP. We note however, that ESOP can readily interact with other epidemiological models, e.g. those that incorporate stratification based on region and age e.g. INDSCI-SIM (Shekatkar et al. 2020), COVision (Nagori et al. 2020), ABCS (Harsha et al. 2020) and IndiaSim (Megiddo et al. 2014). Although machine learning techniques have been used in epidemiological forecasting (Lindström et al. 2015) and estimating model parameters (Dandekar and Barbastathis 2020), we are not aware of prior work using machine learning in epidemiological policy design.

Related Works

Several works exist on modelling epidemic and pandemic progressions using differential equation-based models and using them to make predictions. Some examples include (Efimov and Ushirobira 2020; Lyra et al. 2020; Sardar et al. 2020; Vyasarayani and Chatterjee 2020). Most of these studies utilize differential equation-based models such as SIR and SEIR variants. This paper instead uses a stochastic agent-based model called VIPER that is proposed in Sect. "VIPER: An Iterative Stochastic Agent-based Epidemiological Model". The reason behind this choice was to demonstrate the effectiveness of our proposed techniques when working with epidemiological models that are not described compactly using a few equations and thus, harder to analyze manually.

Micro-simulation studies using UK (Ferguson et al. 2020) and Indian (Singh and Adhikari 2020) data conclude that multiple short-term suppression rounds may offer acceptable health outcomes when a single extended period of suppression is infeasible. However, these works do not offer ways to find either the optimal moment to initiate suppression measures or their duration.

This is important since (Morris et al. 2020; Patterson-Lomba 2020) show that the optimal initiation point and duration of a suppression may depend substantially on the disease and social characteristics themselves (for instance the disease incubation period and the basic reproduction

number R_0). This is understandable since premature suppression would slow down the depletion of the pool of susceptible individuals leaving room open for a second wave of infections whereas delayed suppression may cause the initial wave to be widespread in itself. Scherbina (2020) additionally considers the economic impact of these measures and suggests durations for lock-down periods and their associated economic costs in medical expenses as well as lost value of statistical life.

Prior works offering actual policy advice fall into two categories: (1) those that offer only broad principles on how to target interventions e.g. by identifying simple rules of thumb (Wallinga et al. 2010), and (2) those that offer actionable advice e.g. when to initiate suppression (Klepac et al. 2011; Morris et al. 2020; Torre et al. 2019; Zhao and Feng 2019). However, the latter often do not take the socio-economic impact of these measures into account and moreover, consider only simple theoretical models e.g. SIR that can be manually analyzed but are not very expressive.

A subclass of the latter approaches, for example (Bussell et al. 2019), advocate first fitting an approximate model to the actual simulator (to make it simple enough to enable manual mathematical analyses) before applying optimal control strategies. Such approximations may introduce unmodelled errors into the prediction pipeline and adversely affect their outcome. ESOP instead directly models intervention outcomes in terms of the simulator outputs.

Relevance to CoViD-19 and Future Prospects

Given the evolving nature of the current CoViD-19 pandemic, a technique like ESOP helps in optimally designing multi-phase lock-downs, thus avoiding speculation and human error. As shown in Sect. "Experimental Case Studies", whenever multiple waves of infection are unavoidable due to constraints on lock-downs, ESOP offers lock-down schedules that balance the peaks of these multiple outbreaks, ensuring no peak is too high. To maximize the impact of methods such as ESOP, close interaction and collaboration is needed with experts in the epidemiological and social sciences to better align ESOP with professional epidemiological forecasting and economic forecasting models. ESOP's interaction with these models is of a *black-box* nature which makes integration smoother and simpler.

VIPER: An Iterative Stochastic Agent-based Epidemiological Model

VIPER (Virus-Individual-Policy-EnviRonment) models an *in-silico* population of individuals, supports compartments of the SEIR model (Keeling and Rohani 2008), and allows travel and quarantining of individuals. Being an ABM rather

than an ODE-based model, VIPER can model disease progression within each individual separately and thus, quarantine or expire individuals based on their stage of the disease, something that is difficult to do in ODE-based models. Stochastic ABMs allow diverse socio-medico-economic traits to be modeled at the individual level but cannot be easily represented by a concise system of ODEs. Thus, works such as (Morris et al. 2020) do not apply here.

Details of the VIPER model are described below and succinctly enumerated in Table 1. VIPER consists four *sub-models*, one each devoted to modelling individuals, the virus, environment parameters and policy parameters.

Individual model: An individual is characterized by their susceptibility to infection (SUS), resistance to disease progression (RST), viral load (VLD), recovery load (RLD), current state (STA), quarantine status (QRN) and location (X,Y). SUS, RST, VLD DLD, X and Y are real numbers between 0 and 1, whereas QRN takes Boolean values. The state of an individual STA can be either S (susceptible), E (exposed), I (infectious), R (recovered) or X (expired/

deceased). Individuals are initialized with random values for RST, SUS and their location within the 2-D box $[0, 1]^2$. Individuals progress from state $S \rightarrow E \rightarrow I$, can be optionally quarantined while in state I, and then move to either state R or X.

Viral model: The virus is characterized by its incubation period (INC), the base viral load in an individual at the end of the incubation period (BVL), the disease progression rate (DPR), the viral load over which an individual’s chances of getting expired start increasing (XTH) and the base removal probability of an individual with viral load at XTH (BXP). BVL, DPR, XTH and BXP are real numbers between 0 and 1 whereas INC is a natural number.

Environment model: The environmental factors are modeled using the typical contact radius between individuals (BCR), the probability that a contact between an infectious and susceptible individual will lead to a successful infection (BIP), the fraction of the population that travels at any time instant (BTR), the maximum distance to which they travel (BTD), and the fraction of population that is infected with the virus at start of the simulation (INI). All these values are represented as real numbers between 0 and 1.

Policy model: The policy model comprises the viral load over which an individual’s chances of getting quarantined start increasing (QTH) and the base quarantining probability of an individual with viral load at QTH (BQP). Both are real numbers between 0 and 1. Additionally, the policy prescribes a *lock-down* level which is a real number between 0 and 5. The lock-down level is specified at every time instant t of the simulation. A lock-down level of l at a certain time instant causes BTD as well as BCR to go down by a factor of $\exp(-l)$. Thus, at a high lock-down level, individuals are neither able to travel much, nor interact with other individuals far off from their current location.

Modelling disease-progression dynamics in VIPER : The viral load (VLD) of an individual represents the extent of infection within their system. At the end of the incubation period (INC), an exposed individual always has a “base” viral load of BVL. The virus attempts to increase this viral load according to the disease progression rate (DPR) whereas the individual resists this according to their resistance level (RST) by converting viral load to recovery load (RLD). Disease progression in every infected individual is governed by an SIR-like model with

$$\frac{d \text{VLD}(t)}{dt} = -\text{RST} \cdot \text{VLD}(t) + \text{DPR} \cdot (1 - \text{VLD}(t) - \text{RLD}(t))$$

and

Table 1 VIPER model attributes, valid ranges and default/initial values

Attr	Description	Range	Def
<i>Viral model</i>			
INC	Incubation period	\mathbb{N}	3
BVL	Base viral load	[0, 1]	0.05
DPR	Disease progression rate	[0, 1]	0.1
XTH	VLD threshold for expiry	[0, 1]	0.7
BXP	Expiry probability at XTH	[0, 1]	0.0
<i>Environment model</i>			
BCR	Contact radius b/w individuals	[0, 1]	0.25
BIP	Prob. infection upon contact	[0, 1]	0.5
BTR	Prob. of an individual traveling	[0, 1]	0.01
BTD	Maximum travel distance	[0, 1]	1.0
INI	Initial rate of infection	[0, 1]	0.01
<i>Individual model</i>			
SUS	Susceptibility to infection	[0, 1]	rnd
RST	Resistance to disease progression	[0, 1]	rnd
VLD	Current viral load	[0, 1]	0.0
RLD	Current recovery load	[0, 1]	0.0
STA	Current state	SEIRX	S
QRN	Quarantine status	0 or 1	0
X, Y	Current location	$[0, 1]^2$	rnd
<i>Policy model</i>			
QTH	VLD threshold for quarantine	[0, 1]	0.3
BQP	Quarantine probability at QTH	[0, 1]	0.0
$l(t)$	Lock-down level at time t	[0, 5]	–

See Sect. “VIPER: An Iterative Stochastic Agent-based Epidemiological Model” for details

$$\frac{d \text{RLD}(t)}{dt} = \text{RST} \cdot \text{VLD}(t)$$

Thus, VIPER allows individuals to experience disease progression, as well as associated effects like quarantining or expiry, in a completely individualized manner, something that is readily possible in agent-based models but much more difficult to express in terms of a compact set of differential equations.

An infected individual whose VLD falls below BVL moves on to state R, i.e. recovers. An individual with VLD equal to QTH (resp. XTH) has a probability BQP (resp. BXP) of getting quarantined (resp. expired). An individual with VLD greater than QTH has the following probability of getting quarantined

$$\mathbb{P}[\text{quarantine}] = \text{BQP} + (1 - \text{BQP}) \cdot \frac{\text{VLD} - \text{QTH}}{1 - \text{QTH}},$$

i.e. the probability of getting quarantined increases linearly to 1 as the individual's VLD goes up. At every time step t , a coin is tossed for all individuals with VLD greater than QTH which lands heads with this particular probability. If the coin does indeed land heads, the individual is deemed quarantined. Note that this allows VIPER to model asymptomatic transmission since it allows individuals with low VLD levels (esp. below QTH) to avoid detection with high probability but makes it difficult for those in advanced stages of the disease to avoid quarantine.

Similarly, an individual with VLD greater than XTH has the following probability of getting expired

$$\mathbb{P}[\text{expiry}] = \text{BXP} + (1 - \text{BXP}) \cdot \frac{\text{VLD} - \text{XTH}}{1 - \text{XTH}}.$$

At every time step, a coin is similarly tossed to decide on whether an individual with VLD greater than XTH gets expired or not. The lock-down level needs to be specified at every time instant t of the simulation. A lock-down level of $l(t)$ causes BTD as well as BCR at that time t to go down by a factor of $\exp(-l(t))$. Thus, at high lock-down levels, individuals are neither able to travel much, nor have contact with other individuals far off from their current location. The lock-down level has no effect on the quarantining or expiry processes described above which continue the same way irrespective of the lock-down level.

ESOP: Epidemiologically and Socio-economically Optimal Policies

ESOP encodes interventions as vectors and their health and socio-economic outcomes as functions, e.g., the coordinates of a 2-D vector $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2] \in \mathbb{N}^2$ may encode the starting point (\mathbf{x}_1) and duration (\mathbf{x}_2) of a lock-down. Next, consider

a function $f_{\text{epi}} : \mathbb{N}^2 \rightarrow [0, 1]$ encoding health outcomes with $f_{\text{epi}}(\mathbf{x})$ equal to the peak infection rate (the largest fraction of the total population infected at any point of time) if the intervention \mathbf{x} is applied. Similarly, let $f_{\text{eco}} : \mathbb{N}^2 \rightarrow [0, 1]$ encode economic outcomes with $f_{\text{eco}}(\mathbf{x})$ being the fraction of population that would face unemployment if lock-down were indeed to last \mathbf{x}_2 days. We stress that the functions $f_{\text{epi}}, f_{\text{eco}}$ described here are examples and other measurable outcomes, e.g. cumulative death rate, loss to GDP, can also be used. Predicted estimates for f_{epi} would be obtained from epidemiological simulators such as INDSICI-SIM or IndiaSim (we will use VIPER) and those for f_{eco} would be obtained from economic models. Our goal is to balance health and economic outcomes by solving the following optimization problem:

$$\mathbf{x}^* := \arg \min_{\mathbf{x} \in \mathbb{N}^2} f(\mathbf{x}) \text{ where } f(\mathbf{x}) = f_{\text{epi}}(\mathbf{x}) + f_{\text{eco}}(\mathbf{x})$$

However, it is challenging to perform this optimization using standard descent techniques (Boyd and Vandenberghe 2004) since even obtaining values of the function f at specific query points (let alone gradients) is expensive as it involves querying simulators such as VIPER. An acceptable solution in this case is Bayesian optimization (Jones et al. 1998) which is an *active machine learning* technique used to optimize functions which are expensive to evaluate and to which, moreover we do not have access to gradients.

Figure 1 presents a visual depiction of the key processes in Bayesian optimization. The technique adaptively queries the function at only a few locations to quickly approximate the solution to the optimization problem. The algorithm uses the current observations and Gaussian process regression (Rasmussen and Williams 2006) to obtain a *mean* estimate $f(\hat{\mathbf{x}})$ (dashed line in Fig. 1) of the true function $f(\mathbf{x})$ (bold line in Fig. 1) as well as an estimate $\sigma(\mathbf{x})$ of the *uncertainty* in that estimate (the blue shaded region depicts $f(\hat{\mathbf{x}}) \pm \sigma(\mathbf{x})$ in Fig. 1). Notice that uncertainty drops around observation points since (a good estimate of) the true function value is known there.

Using these, an *acquisition* function is created. Figure 1a uses the simple LCB (lower confidence bound) acquisition function defined as $a(\mathbf{x}) = f(\hat{\mathbf{x}}) - \sigma(\mathbf{x})$. Other possibilities include EI (expected improvement) and KG (knowledge gradient). The (estimated) function value at the point $\hat{\mathbf{x}}_{\text{query}} := \arg \min a(\mathbf{x})$ is now queried. Using $f(\hat{\mathbf{x}}_{\text{query}})$, the mean and uncertainty estimates \hat{f}, σ are updated as shown in Fig. 1b and the process is repeated. At the end, the query point with the lowest function value is returned as the estimated minimum. Note that ESOP can only query VIPER but does not have access to its internal attribute values.

Lack of space does not permit a detailed overview and we refer the reader to (Frazier 2018) for an excellent review. ESOP also employs multi-scale search and caching




Fig. 1 An illustration of the Bayesian optimization process. See Sect. "ESOP: Epidemiologically and Socio-economically Optimal Policies" for details

techniques to accelerate computations. Bayesian optimization routines often enjoy provable convergence bounds which we briefly discuss in Sect. [Convergence Rate of ESOP](#).

Experimental Case Studies

We present case studies with VIPER simulating an in-silico population of 20,000 individuals (ESOP scales to larger populations too). The default attributes settings in VIPER

are given in Table 1. Any modifications to these are mentioned below.

Choice of f_{epi} and f_{eco}

A lock-down is represented as a 3-D vector $\mathbf{x} = (i, p, l)$ of its initiation point (i), period (p) and level (l). Our objective is to minimize $f_{epi} + f_{eco}$. We set $f_{epi}(\mathbf{x})$ to be the peak of the E+I curve (i.e. the peak number of exposed + infectious individuals—please see Fig. 2) if a level- l lock-down is indeed started at $t = i$ for p time steps. We use $f_{eco}(\mathbf{x}) = \frac{l}{5} \cdot p \cdot \frac{N}{1000}$ to estimate job losses due to the lock-down assuming that a

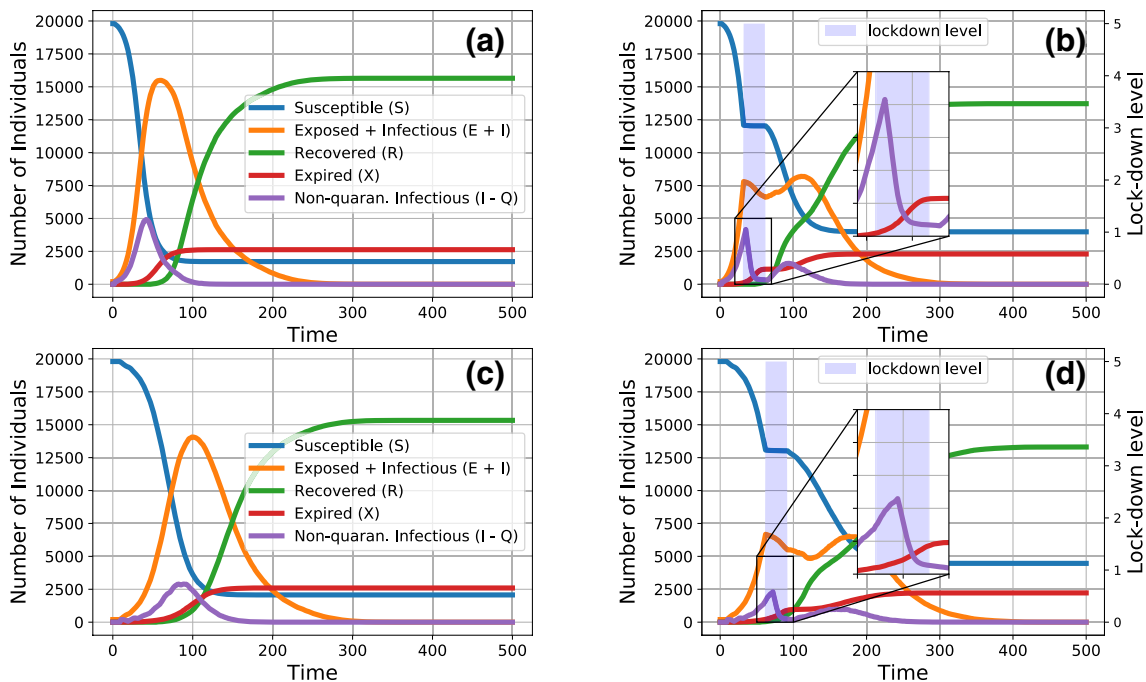


Fig. 2 Using ESOP to discover the optimal point at which to initiate a lock-down. **a, b** Consider a virus strain with an incubation of 3 days whereas **c, d** consider a strain with a 10 day incubation period. The optimal initiation point of a lock-down depends strongly on viral

characteristics e.g. incubation period. ESOP is able to discover a near-optimal initiation point in both cases within very few iterations (see Fig. 3. See Sect. "Using ESOP to Find the Optimal Initiation Point of a Lock-down" for details

level- l lock-down forces an $\frac{l}{50}$ percent of the population of $N = 20000$ into unemployment each day for p days. Note that most natural definitions of f_{epi} and f_{eco} would conflict with each other since f_{epi} would promote aggressive and sustained lock-downs whereas f_{eco} would oppose them.

The notion of $f_{\text{eco}}(\cdot)$ used above is merely demonstrative and ESOP can instead use a more realistic definition of f_{eco} that might be a non-linear function of \mathbf{x} , by asking an economic simulator, just as it asks values of f_{epi} from an epidemiological simulator such as VIPER. Also, defining the objective as an additive sum $f_{\text{epi}} + f_{\text{eco}}$ of the two functions is not essential for ESOP to function, and users may instead prefer other formulations e.g. $f_{\text{eco}}^\alpha \cdot f_{\text{epi}}^\beta$ for some $\alpha, \beta > 0$, etc.

All that ESOP requires is black-box access to values of the objective function, however it may be defined. As Fig. 3 indicates, ESOP is capable of optimizing highly non-linear functions as well. However, the algorithm would naturally require more iterations if the objective becomes highly convoluted and sensitive to changes in the input (see Sect. "Concluding Remarks" for a brief discussion).

Using ESOP to Find the Optimal Initiation Point of a Lock-down

Figure 2 considers a simple case where we have decided to impose a 30 time step lock down at level 5 but are unsure when to initiate the lock down for optimal effect. The objective here is to minimize f_{epi} alone and f_{eco} is not considered since we have already decided the duration and intensity of the lock-down in this case and hence resigned to a predictable economic outcome. For any initiation point $i \in \mathbb{N}$, let $f_{\text{epi}}(i)$ be the largest number of individuals in E and I states at any given point of time (the so called *peak* of the curve) if a level 5 lock-down is initiated at $t = i$.

Figure 2a shows the daily count of individuals in various categories if no suppression is used. Note the large number

of non-quarantined yet infectious individuals (I-Q) who are responsible for disease spread. The number of exposed and infectious (E+I) individuals peaks at around 15,500 at $t = 58$. For Fig. 2b, ESOP was asked to suggest when to start a 30 day lock-down at level 5. It suggested starting at $t = 33$ which brings the peak down to 8200 cases, a reduction of 47%. Figure 2c, d show similar results but for a viral strain that has an incubation of 10 days instead of 3 days.

The results show that the optimal initiation point depends significantly on disease characteristics, e.g., incubation period of the virus. Nevertheless, ESOP discovers a near-optimal solution, offering far superior health outcomes compared to a no-lock-down scenario. Note that the number of non-quarantined infectious individuals continues to rise (see Fig. 2b, d insets) even after imposition of lock-down due to the incubation period of the disease.

Convergence Rate of ESOP

The topic of how fast do Bayesian optimization routines converge to the optimal solution is a subject of intense study but one beyond the scope of this paper. Under appropriate assumptions, Bayesian optimization routines, within T queries to the underlying function (e.g., as ESOP queries f_{epi} via VIPER), are able to offer a solution that is only ϵ_T worse than the optimal solution. The sub-optimality ϵ_T goes down with the number of queries T at a rate $\frac{1}{T^\alpha}$ where $\alpha > 0$ is a constant that depends on the problem setting and the *kernel* used by the Bayesian optimization procedure (all experiments with ESOP use the Matern kernel). We refer the reader to (Srinivas et al. 2010) for technical details of these convergence results.

Figure 3a, b show that ESOP achieved a near-globally optimal initiation point (marked by a yellow diamond) for the two cases considered in Fig. 2 within just 12 and 4 calls to VIPER (see inset figures). To plot the orange curves in Fig. 3, VIPER was queried with all $i \in [0, 100]$ to explicitly




Fig. 3 Convergence rates offered by ESOP. **a, b** Show (see black curve) ESOP's convergence rate in the 3 and 10 day incubation period cases in Fig. 2. The optimal initiation point varies significantly

with the incubation period (33 vs 62 days). See Sect. "Convergence Rate of ESOP" for details

reveal the globally optimum. This was done so that we could verify that ESOP does indeed reach the global optimum. However, ESOP itself requires far fewer, e.g. 12 or 4 calls to VIPER to reach the global minimum.

Using ESOP to Optimize (Multi-phase) Lock-downs under Constraints

Administrative and logistical considerations put constraints on lock-downs, e.g., on when a lock-down can be initiated or, on how long it can last. Another constraint might be taking into account previously applied lock-downs. We demonstrate how ESOP can be used to design lock-down schedules that are optimal among those that satisfy a given set of constraints. Figure 4 considers finding the optimal initiation, period and level of a lock-down, subject to constraints.

Single-phase Lock-downs under Constraints. For Fig. 4a, ESOP was asked to suggest a lock-down (no constraints on duration etc). It suggested one at level 3.5 starting $t = 3$, lasting 69 steps, resulting in a peak of 303 infections (much smaller than the peaks in Fig. 2) and 1092 predicted cases of unemployment. We note that f_{eco} prevents the lock-down from going on indefinitely. The suggested lock-down seems to adopt a containment strategy, initiating a moderate-level suppression very early on to deplete the pool of infectious

individuals. Note that the pool of non-quarantined infectious individuals is indeed exhausted by the time the suggested lock-down is over, preventing a second wave of infections. For Fig. 4b, ESOP was forced to suggest a lock-down starting no earlier than $t = 12$ and lasting no longer than 40 days. Since infections are already rampant by $t = 12$, ESOP instead adopts a mitigation strategy of starting a lock-down at $t = 30$ for 40 steps at level 3.5, causing a peak of 7353 infections and 560 cases of unemployment. Notice that the lock-down is strategically delayed so that the second wave does not have a higher peak, thus balancing the two peaks indeed as dictated by f_{epi} .

Multi-phase Lock-downs. For Fig. 4c, ESOP was given a situation where an earlier lock-down (green shading) had already taken place but was ineffective and left alone, would have caused a massive second wave with a peak of 14384 (dotted orange curve). ESOP was asked to suggest a new lock-down that starts no earlier than 10 days after the previous lock-down ended and lasting no longer than 40 days. ESOP suggested a second lock-down starting at $t = 56$ lasting 35 steps at level 4 which brings the peak down to 8447 (a reduction of 40%) and causing 560 additional cases of unemployment. Figure 4d considers a scenario where the policy makes are still dissatisfied with the outcomes, and request a third lock-down starting no earlier than 10 days

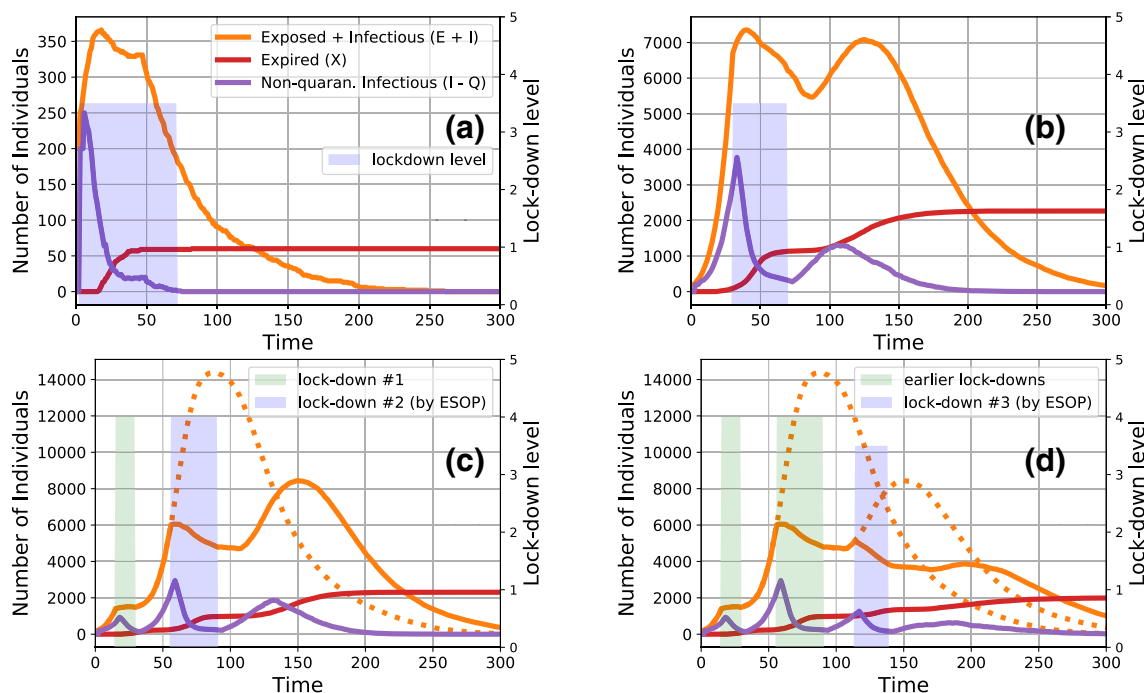


Fig. 4 Optimizing (multi-phase) lock-downs under constraints. See Sect. "Using ESOP to Optimize (Multi-phase) Lock-downs under Constraints" for details. ESOP is able to shift from a containment strategy (e.g. in a) to a mitigation strategy (e.g. in b) depending on

the constraints. When multiple spikes are inevitable due to constraints on the lock-down (e.g. upper limits on the duration of the lock-downs), ESOP balances the heights of the spikes to ensure that infections are evenly distributed among them

after the second lock-down ended and lasting no longer than 40 days. ESOP suggests a third milder lock-down at level 3.5 starting at $t = 114$ and lasting 25 days. This brings the peak to a much lower number (6052 i.e. a further 28% reduction) and 350 additional cases of unemployment.

Can Aggressive Quarantining Permit ESOP to Offer Less Severe Lock-downs?

Intuitively, if aggressive testing and quarantining is applied, then it should be possible to avert an epidemic with milder lock-downs. Figure 5 verifies this claim by considering at the situation in Fig. 4b (i.e. starting a lock-down no earlier than $t = 12$ days and lasting no longer than 40 days), but with varying quarantine aggressiveness. As Table 1 and Sect. 2 explain, individuals with viral loads over a threshold QTH get quarantined with a certain probability profile. Using greater testing and public awareness, this profile can be altered.

Figure 5a considers sluggish quarantining with $QTH = 0.9$ and $BQP = 0.0$ (the quarantining probability profile is shown in Fig. 5a as an inset). Almost no individuals get quarantined and despite its best efforts, ESOP is only able to offer 11338 peak infections and 800 unemployment cases using a lock-down starting at $t = 32$ (level 5, 40 steps). With stronger quarantining at $QTH = 0.4$ and $BQP = 0.3$, the

situation improves in Fig. 5b where ESOP offers fewer infections (7930) and job losses (560) using a lock-down at a reduced level of 3.5 (starting $t = 30$, 40 steps). With still stronger quarantining ($QTH = 0.0$, $BQP = 0.0$) in Fig. 5c, ESOP offers a peak of 871 and 378 job losses using a lock-down lasting fewer (27) steps (starting $t = 15$, level 3.5). Figure 5c reports the outcomes with even stronger quarantining at $QTH = 0.2$ and $BQP = 0.5$ where ESOP offers a peak of just 706 and 192 job losses using a lock-down lasting just 16 steps (level 3 starting $t = 12$).

Variations in Demographics and Geographical Distribution of Population

We also consider location (X, Y), susceptibility (SUS) and resistance (RST) values (see Table 1) which are not set uniformly but fitted to statistical data for India.

Population Clusters. Instead of distributing individuals uniformly in the box $[0, 1]^2$, as we did earlier, we now distribute 34% of the population into 4 Gaussian clusters with standard deviation 0.1 (to simulate crowded urban areas). India does have a similar proportion of urban population (Bank 2018). We distribute the rest 66% of the population uniformly (to simulate scattered non-urban population).

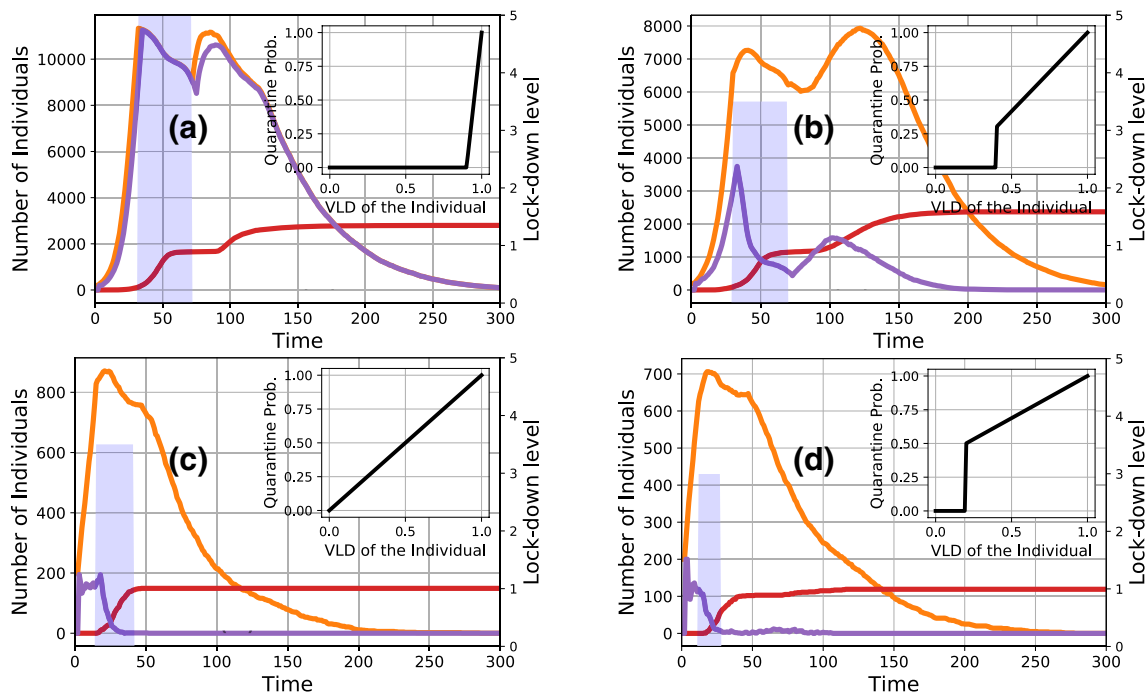


Fig. 5 Effect of quarantine policy on ESOP's ability to offer desirable outcomes. The area under the black curve in the insets roughly measures the aggressiveness of the quarantining policy. See Sect. "Can Aggressive Quarantining Permit ESOP to Offer Less Severe

Lock-downs?" for details. ESOP offers better outcomes, both epidemiologically as well as economically, and that too with milder lock-downs, if strong quarantining is applied

Fitted Demographic Data. Instead of distributing SUS and RST values uniformly in $[0, 1]$ as we did earlier, we now distribute these according to age and co-morbidity statistics for India. (Joshi 2020, [Table 1]) reports that the four factors with highest risk score for CoViD-19 patients are age above 55 years, male gender, hypertension, and diabetes. We collected age-stratified statistics for all these factors from respectively (Home Affairs India 2016, [Detailed Tables]), (Ramakrishnan 2019, [Table 1]) and (Group 2003, [Table 2]). After assigning age, gender, and co-morbidity values to all individuals so as to fit the Indian statistics, individuals were assigned a point each for being male, over 55,

and suffering from each of the two co-morbidities. Thus, an individual could clock up a maximum of 4 points. An individual's SUS value was then decided as $0.1 + 0.2 \times \#points$. Their RST value was set to $RST = 1 - SUS$. Figure 6b shows the resulting distributions.

Results with Fitted Demographic Data and Population Clusters

Figure 7a is identical to Fig. 4b (where ESOP was asked for a lock-down lasting no longer than 40 days and starting no earlier than $t = 12$) except that Fig. 7a also shows




Fig. 6 Demographic data fitted to Indian statistics. The demographic advantage of the Indian population is evident from both curves. **a** Shows that the vast majority of the population do not cross the risk-

factor forming seniority threshold whereas **b** shows that most individuals have high RST values, as calculated in Sect. "Variations in Demographics and Geographical Distribution of Population"

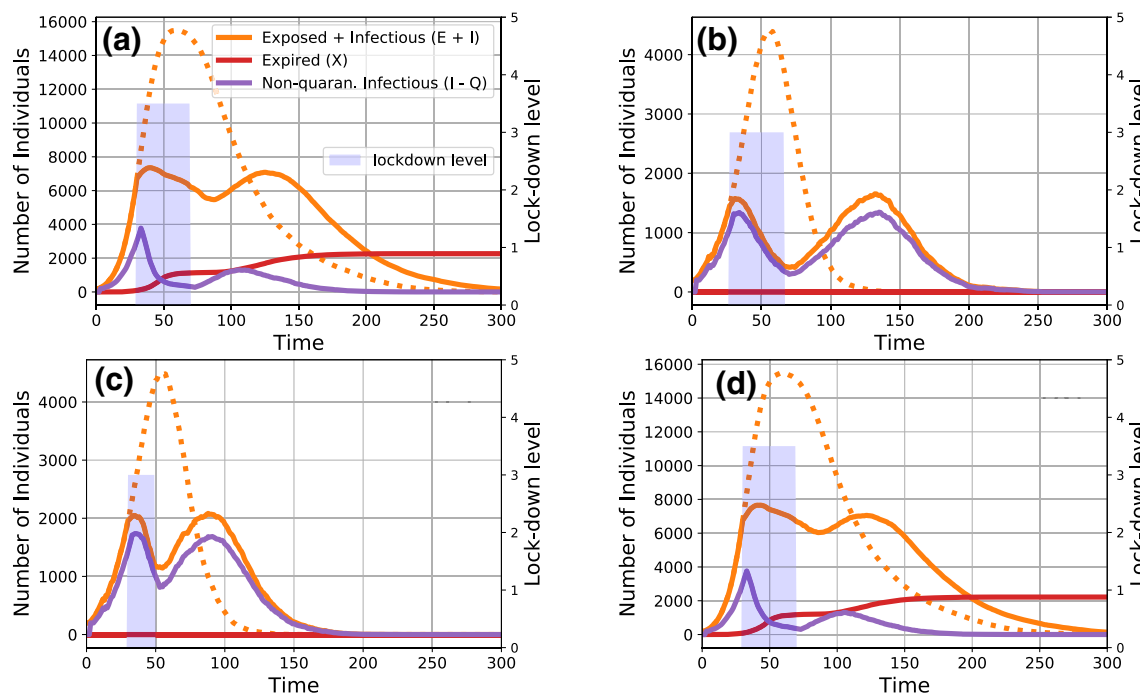


Fig. 7 ESOP with fitted demographic and location data. See Sect. "Results with Fitted Demographic Data and Population Clusters" for details. The dotted lines in all figures gives the infection curve had no lock-down been applied

(using a dotted orange line) what the infection curve would have been, had no lock-down been applied at all. Recall that with the lock-down suggested by ESOP, we had a peak of 7353 infections and 560 cases of unemployment. Figure 7b presents the situation where we use population clusters as well as fitted demographic data. The improvement with ESOP's new suggestion is extremely significant with a peak of just 1653 infections (more than 4× reduction) and 480 cases of unemployment, that too, using a milder lock-down at level 3.0.

To ascertain whether this improvement was due to the demographic change or the population clusters, Fig. 7c, d present ablation studies. In Fig. 7c, only the demographic distributions of RST and SUS are fitted to Indian data but locations are kept uniformly distributed over the 2-D box $[0, 1]^2$ as we were doing earlier. In Fig. 7d, RST and SUS values are uniformly distributed over $[0, 1]$ as we were doing earlier, but 34% of the population is clustered into “cities”.

It is clear that changing the location distribution of the individuals alone worsens the outcome (Fig. 7d reports a peak of 7657 that is higher than Fig. 7a). This is to be expected since people are now crowded into cities. However, changing the demographics of the population to match that of India greatly improves the outcome (Fig. 7c reports a peak of just 2083 and just 240 cases of unemployment). This is also to be expected since India has an extremely favorable age structure. Since applying both changes together, as in Fig. 7b, also presents a significant improvement over Fig. 7a, this seems to suggest that the demographic advantages more than overcome the disadvantage due to crowding in cities.

Concluding Remarks

Incorporating age stratification and climate into VIPER and augmenting ESOP to suggest “personalized” age- and climate-specific policies would be interesting. It would also be interesting to develop epidemiological models that simultaneously track multiple diseases with similar or confounding symptoms, such as CoViD-19 and ILI (Influenza-like illnesses) as it would allow policy models e.g. testing and quarantining models to be checked for false positive and missed detection rates. Lastly, non-linear systems with negative feedback loops such as epidemiological models, are known to exhibit chaotic behavior (Bolker 1993; Eilersen et al. 2020). It is interesting how techniques like ESOP can be adapted to handle such systems.

Acknowledgements D.D. is supported by the Visvesvaraya PhD Scheme for Electronics & IT (FELLOW/2016-17/MLA/194). P.K. thanks Microsoft Research India and Tower Research for research grants.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Code availability All code used for this study is available at the following GitHub Repository <https://github.com/purushottamkar/esop>.

References

- Aledort JE, Lurie N, Wasserman J, Bozzette SA (2007) Non-pharmaceutical public health interventions for pandemic influenza: an evaluation of the evidence base. *BMC Public Health* 7(208):
- Bank TW (2018) Urban population (% of total population) - India. <https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS?locations=IN>
- Bolker B (1993) Chaos and complexity in measles models: a comparative numerical study. *IMA J Math Appl Med Biol* 10:83–95
- Boyd S, Vandenberghe L (2004) *Convex optimization*. Cambridge University Press, Cambridge
- Bussell EH, Dangerfield CE, Gilligan CA, Cunniffe NJ (2019) Applying optimal control theory to complex epidemiological models to inform real-world disease management. *Philosophical Transactions of the Royal Society B* 374(1776):
- Dandekar R, Barbastathis G (2020) Neural Network aided quarantine control model estimation of global Covid-19 spread, [arXiv:2004.02752v1](https://arxiv.org/abs/2004.02752v1) [q-bio.PE]
- Efimov D, Ushirobira R (2020) On an interval prediction of COVID-19 development based on a SEIR epidemic model. *Tech. rep.*, Inria Lille Nord Europe - Laboratoire CRISTAL - Université de Lille, hal-02517866v4
- Eilersen A, Jensen MH, Sneppen K (2020) Chaos in disease outbreaks among prey. *Sci Rep* 10(3907):
- Ferguson NM, Laydon D, Nedjati-Gilani G, Imai N, Ainslie K, Baguelin M, Bhatia S, Boonyasiri A, Cucunubá Z, Cuomo-Dannenburg G, Dighe A, Dorigatti I, Fu H, Gaythorpe K, Green W, Hamlet A, Hinsley W, Okell LC, van Elsland S, Thompson H, Verity R, Volz E, Wang H, Wang Y, Walker PG, Walters C, Winskill P, Whittaker C, Donnelly CA, Riley S, Ghani AC (2020) Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College, London
- Frazier PI (2018) A Tutorial on Bayesian Optimization, [arXiv:1807.02811v1](https://arxiv.org/abs/1807.02811v1) [stat.ML]
- Group TDS (2003) Age- and sex-specific prevalence of diabetes and impaired glucose regulation in 11 Asian cohorts. *Diabetes Care* 26(6):1770–1780
- Harsha P, Juneja S, Patil P, Rathod N, Saptharishi R, Sarath AY, Sriram S, Srivastava P, Sundaresan R, Vaidhiyan NK (2020) COVID-19 Epidemic Study II: Phased Emergence From the Lockdown in Mumbai, [arXiv:2006.03375](https://arxiv.org/abs/2006.03375) [q-bio.PE]
- Home Affairs India OoM (2016) Sample Registration System Statistical Report 2016. https://censusindia.gov.in/vital_statistics/SRS_Reports_2016.html, accessed 14 June, 2020
- Jones DR, Schonlau M, Welch WJ (1998) Efficient global optimization of expensive black-box functions. *J Glob Optim* 13:455–492
- Joshi SR (2020) Indian COVID-19 risk score, comorbidities and mortality. *J Assoc Phys India* 68(5):11–12
- Keeling MJ, Rohani P (2008) *Modeling infectious diseases in humans and animals*. Princeton University Press, Princeton
- Kimball R, Bose D (2020) 11 pandemics that changed the course of human history, from the Black Death to HIV/AIDS - to coronavirus. *Business Insider India*, [https://www.businessinsider.in/science/news/11-pandemics-that-changed-the-course-of-human-histo](https://www.businessinsider.in/science/news/11-pandemics-that-changed-the-course-of-human-history)

- [ry-from-the-black-death-to-hiv/aids-to-coronavirus/articleshow/74695609.cms](#). Accessed 19 April 2020
- Klepac P, Laxminarayan R, Grenfell BT (2011) Synthesizing epidemiological and economic optima for control of immunizing infections. *Proc Natl Acad Sci* 108(34):14366–14370
- Lindström T, Tildesley M, Webb C (2015) A Bayesian ensemble approach for epidemiological projections. *PLOS Comput Biol* 11(4):e1004187
- Lyra W, do Nascimento Jr JD, Belkhiria J, de Almeida L, Chrispim PPM, de Andrade I (2020) COVID-19 pandemics modeling with SEIR(+CAQH), social distancing, and age stratification. The effect of vertical confinement and release in Brazil., medRxiv: 2020.04.09.20060053
- Megiddo I, Colson AR, Nandi A, Chatterjee S, Prinja S, Khera A, Laxminarayan R (2014) Analysis of the universal immunization programme and introduction of a rotavirus vaccine in India with IndiaSim. *Vaccine* 325:A151–A161
- Morris DH, Rossine FW, Plotkin JB, Levin SA (2020) Optimal, near-optimal, and robust epidemic control, [arXiv:2004.02209v1](#) [q-bio.PE]
- Nagori A, Awasthi R, Joshi V, Vyalla SR, Jarodia A, Gupta C, Gulati A, Bandhey H, Kumaraguru P, Sethi T (2020) Less Wrong COVID-19 Projections With Interactive Assumptions, medRxiv: 2020.06.06.20124495
- Patterson-Lomba O (2020) Optimal timing for social distancing during an epidemic, medRxiv: 2020.03.30.20048132
- Peak CM, Childs LM, Grad YH, Buckee CO (2017) Comparing non-pharmaceutical interventions for containing emerging epidemics. *Proceedings of the Nat Acad Sci USA*
- Perrings C, Castillo-Chavez C, Chowell G, Daszak P, Fenichel EP, Finnoff D, Horan RD, Kilpatrick AM, Kinzig AP, Kuminoff NV, Levin S, Morin B, Smith KF, Springborn M (2014) Merging economics and epidemiology to improve the prediction and management of infectious disease. *Ecohealth* 11(4):464–475
- Ramakrishnan S et al (2019) Prevalence of hypertension among Indian adults: results from the great India blood pressure survey. *Indian Heart J* 71(4):309–313
- Rasmussen CE, Williams CKI (2006) *Gaussian processes for machine learning*. MIT Press, Cambridge
- Roos R (2012) CDC estimate of global H1N1 pandemic deaths: 284,000. Newsletter of the Center for Infectious Disease Research and Policy, University of Minnesota, <https://www.cidrap.umn.edu/news-perspective/2012/06/cdc-estimate-global-h1n1-pandemic-deaths-284000>. Accessed 19 April 2020
- Sardar T, Nadimb SS, Chattopadhyay J (2020) Assessment of 21 Days Lockdown Effect in Some States and Overall India: A Predictive Mathematical Study on COVID-19 Outbreak, [arXiv:2004.03487v1](#) [q-bio.PE]
- Scherbina A (2020) Determining the Optimal Duration of the COVID-19 Suppression Policy <https://doi.org/10.2139/ssrn.3562053>. Accessed 19 April 2020
- Shekatkar S, Pujari B, Arjunwadkar M, Hazra DK, Chaudhuri P, Sinha S, Menon GI, Sharma A, Guttal V (2020) INDSCI-SIM: A state-level epidemiological model for India. Ongoing Study at <https://indscicov.in/indscisim>
- Singh R, Adhikari R (2020) Age-structured impact of social distancing on the COVID-19 epidemic in India, [arXiv:2003.12055v1](#) [q-bio.PE]
- Srinivas N, Krause A, Kakade S, Seeger M (2010) Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In: *Proceedings of the 27th International Conference on Machine Learning (ICML)*
- Torre DL, Malik T, Marsiglio S (2019) Optimal Control of Prevention and Treatment in a Basic Macroeconomic-Epidemiological Model, [arXiv:1910.03383](#) [econ.TH]
- Vyasarayani CP, Chatterjee A (2020) New approximations, and policy implications, from a delayed dynamic model of a fast pandemic, [arXiv:2004.03878](#) [q-bio.PE]
- Wallinga J, van Boven M, Lipsitch M (2010) Optimizing infectious disease interventions during an emerging epidemic. *Proceedings of the Nat Acad Sci USA*
- WHO (2020) Coronavirus disease (COVID-19) Pandemic. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019>. Accessed 22 June 2020
- Zhao H, Feng Z (2019) Identifying optimal vaccination strategies via economic and epidemiological modeling. *J Biol Syst* 27(4):423–446

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.