

MODELING AND SIMULATION FOR THE SPREAD OF COVID-19 IN AN INDIAN CITY: A CASE STUDY

Aditya A. Paranjape
Souvik Barat
Anwasha Basu

Rohan Salvi

Software Systems & Services, TCS Research
Tata Research Development and Design Centre
Hadapsar Industrial Estate
Pune 411013, INDIA

School of Information Sciences
University of Illinois at Urbana-Champaign
501 E. Daniel Street
Champaign, IL 61801, USA

Supratim Ghosh

Vinay Kulkarni

Decision & Data Sciences, TCS Research
Tata Research Development and Design Centre
Hadapsar Industrial Estate
Pune 411013, INDIA

Software Systems & Services, TCS Research
Tata Research Development and Design Centre
Hadapsar Industrial Estate
Pune 411013, INDIA

ABSTRACT

We present a case study on modeling and predicting the course of Covid-19 in the Indian city of Pune. The results presented in this paper are concerned primarily with the wave of infections triggered by the Delta variant during the period between February and June 2021. Our work demonstrates the necessity for bringing together compartmental stock-and-flow and agent-based models and the limitations of each approach when used individually. Some of the work presented here was carried out in the process of advising the local city administration and reflects the challenges associated with employing these models in a real-world environment with its uncertainties and time pressures. Our experience, described in the paper, also highlights the risks associated with forecasting the course of an epidemic with evolving variants.

1 INTRODUCTION

Ever since Covid-19 emerged as a pandemic, a broad range of modeling and simulation tools have been employed to predict its course, examine candidate intervention strategies, predict the impact of vaccination campaigns, and inform government policy. Our prior work (Barat et al. 2021) describes an agent-based model developed for the city of Pune in western India. The model was used, notably, to inform the relaxation of restrictions on people's movement right after the nationwide lockdown in India ended in June 2020 and the prerogative for imposing restrictive measures was passed on to the local administrations.

The results reported in (Barat et al. 2021) included a prediction for a second wave of Covid-19 infections in Pune starting approximately March 2021. This prediction was made under the assumption that there would be no major changes in the infectivity or virulence of the Covid-19 variants in circulation when compared to the Alpha variant. Thus, when Covid-19 cases began to rise towards the end of February, we presented our predictions to the Pune administration with the objective of guiding their decision-making. Over a period of two weeks, during which we engaged with and guided the administration, we realized

that our assumptions about the variant were inaccurate as the number of cases vastly exceeded our initial expectations. In this paper, we present that engagement as a case study, with emphasis on modeling and simulation.

We pay particular attention to three points which may be of interest for general epidemiological modeling. The first point concerns the limitations of ABMs at the onset of a wave and how SEIR models can help tune the internal features of the ABM. The second point, which is related to the first, concerns the impact of delays in imposing lockdowns or similar restrictions. The third point concerns the necessity for modeling immune escape and loss of immunity.

1.1 Literature Review

Epidemiological modeling has a long history dating back to the seminal work by Kermack and McKendrick (Kermack and McKendrick 1927) which led to the well-known SIR (susceptible-infected-recovered) model and its variants (Anderson and May 1992). Broadly speaking, there are two families of techniques which employ the SIR approach for modeling epidemics. The first family of techniques is based on modeling the movement of *individuals* from susceptible to infected and onward to recovered (or dead). These agent-based models (ABMs) permit highly realistic and faithful modeling of social behavior and geographical dispersal of populations. (Eubank et al. 2004; Kerr et al. 2021; Ferguson et al. 2020; Wolfram 2020; Silva et al. 2020; Cuevas 2020). In contrast to ABMs, compartmental SIR models work with *cohorts* of individuals such as those that are susceptible to a disease, infected and those that have recovered or died (Anderson and May 1992). Although computationally efficient, compartmental SIR models do not work well when the microscopic behavior of individuals is highly stochastic or their geographic dispersal is substantial (Keeling and Grenfell 2000).

The key objective of ABMs is to capture the behaviour of micro-elements such as people, households and places (*e.g.*, workplace, school, shops) to predict macroscopic indicators such as the number of infected cases and cases that need medical attention. Several ABMs (Silva et al. 2020; Cuevas 2020) consider high-level classify individuals and places into cohorts, where each cohort is internally represented using aggregated equations, and represent the whole system as a connected network of a limited number of cohorts. These models address scalability by aggregation and can estimate the impact of fine-grained interventions. However, like coarse-grained models, their ability to capture emergent behaviour and micro-causalities of pandemic dynamics is limited. The ABM underlying Covasim (Kerr et al. 2021) captures the individualistic behaviour of a wide range of micro-elements and their interactions. It linearly scales down the population (on the order of 10^4) to make the simulation manageable, and captures demographic variations in the population together with a wide range of places and interventions. Although agent-based models can afford a high degree of fidelity, they can be significantly large in size. For instance, the ABM in our prior work (Barat et al. 2021) involved 50,000 agents for prototypical city wards. A high-fidelity model for the predicting the course of Covid in Chicago has been developed with over 2 million agents (Ozik et al. 2021).

The state space can be expanded to accommodate specific characteristics of a disease. For instance, epidemic models for TB typically include multiple infected states, covering latent and active infection (Brandeau 2005). Likewise, models for Covid-19 have put emphasis on asymptomatic infected individuals. One reason for the popularity of SIR models and its variants is that they are relatively low-order models (compared to agent-based or graph-based models) and can be used effectively to design control strategies (Brandeau 2005; Enns and Brandeau 2015) and perform extensive “what-if” analyses (Ray et al. 2020) without being restricted by the heavy computational demands of agent or graph-based methods. SIR models have been used in the recent past for analysing the SARS epidemic from 2003 (Riley et al. 2003) as well as to inform Government policy for responding to Covid-19: see, for instance, the SUTRA model (Agrawal et al. 2021) used to advise the Government of India. A hierarchical model with a high-level compartmental model and a low-level ABM has been presented for the United States in (Mokhtari et al. 2021).

SIR models, being nonlinear, can yield qualitatively dissimilar solutions (trajectories) as functions of system parameters. This is evident in the form of bifurcations in systems that present equilibrium or

periodic solutions (Kuznetsov and Piccardi 1994). Since the susceptible population for Covid-19 reduces monotonically over short time scales (which are of interest here), the only steady state solutions are those that correspond to the attainment of “herd immunity.” However, as demonstrated in the paper, the short-term behavior of these models can be qualitatively sensitive to system parameters.

1.2 Contribution and Organization

As noted earlier, this paper is presented as a case study of Covid-19 modeling and simulation aimed particularly at informing government decisions. In Section 2, we develop a compartmental SEIR (susceptible - exposed - infected - recovered/dead) model which incorporates time delays to account for the time spend by individuals in a given state. This SEIR model forms the basis of a more detailed stock and flow model that we incorporate into a Vensim simulator for parameter identification. We also describe, briefly, an agent based model (ABM) developed using the Enterprise Simulation Language (ESL).

In Section 3, we present the key results of the case study. We start by presenting predictions made using an agent-based model prior to the wave of infections triggered by the Delta variant. Next, we present *post facto* analysis based on the ABM as well as stock-and-flow models to understand the factors which may have influenced case numbers in the Delta-driven wave. Finally, we present novel results that show how the short-term behavior of compartmental SIR models may be qualitatively sensitive to system parameters. We use these results to identify conditions under which restricting social mixing and intermingling at the onset of an epidemic can make a significant difference at least in the short term.

2 MODELING

2.1 Compartmental SEIR Model with Distributed Time Delay

The course of an epidemic in a large, homogeneous population can be described by a set of differential equations featuring five compartments or stocks: susceptible (S), exposed (E), infected (I), recovered (R) and dead (D). This SEIRD model is used to compute the time history of the *number of people* in the respective bins. For a disease like Covid-19, the movement between E, I and R (or D) bins is typically a well-characterized (in a probabilistic sense) time-delayed phenomenon and best described using a distributed time delay. This contrasts with models such as (Kermack and McKendrick 1927; Ma et al. 2002) where a distributed delay is used to model time-varying infectiousness characteristics. Instead, we assume that the infectiousness of individuals is uniform across their period of exposure/infection. We start by defining a finite history of the number of exposed and infected people as follows:

$$E_{\tau}[k] = [\delta E[k], \delta E[k-1], \dots, \delta E[k-\tau]]^{\top}$$

where $\delta E[\cdot]$ denotes the number of new cases that enter the bin E at time k . We clarify that this is not the *net* number of new cases at time k (which would be obtained by subtracting those cases that move out of the bin at time k). For consistency, we set $\delta E[0] = E[0]$. Along similar lines, we also define $I_{\tau}[k]$. The value of τ depends on how long the individuals spend in the exposed/infected stage, and is known to be around 12 days for Covid-19. We rewrite the dynamics as follows:

$$\begin{aligned} S[k+1] &= S[k] - \mu_1 S[k] E[k] - \mu_2 S[k] I[k] + m_S[k] \\ E[k+1] &= E[k] + \mu_1 S[k] E[k] + \mu_2 S[k] I[k] - (\mu_3 + \mu_4) p_{E,\tau}^{\top} E_{\tau}[k] + m_E[k] \\ I[k+1] &= I[k] + \mu_3 p_{E,\tau}^{\top} E_{\tau}[k] - (\mu_5 + \mu_6) p_{I,\tau}^{\top} I_{\tau}[k] + m_I[k] \\ R[k+1] &= R[k] + \mu_4 p_{E,\tau}^{\top} E_{\tau}[k] + \mu_5 p_{I,\tau}^{\top} I_{\tau}[k] + m_R[k] \\ D[k+1] &= D[k] + \mu_6 p_{I,\tau}^{\top} I_{\tau}[k] \end{aligned} \tag{1}$$

where $p_{E,\tau}^{\top} \mathbf{1}_{\tau+1} = p_{I,\tau}^{\top} \mathbf{1}_{\tau+1} = 1$. The two vectors $p_{E,\tau}$ and $p_{I,\tau}$ capture the probability that individuals who have spent a certain of amount of time in a given bin will progress to a subsequent bin. The terms $m_{(\cdot)}[k]$

capture the net in-bound migration of people from other geographical areas within the respective bin. These terms have been included for completeness although we do not consider migration in this paper.

2.2 Parameter Estimation

The parameters μ_1 and μ_2 are difficult to model and need to be predicted using the available data. There are several well-known techniques for this purpose; see (Ozik et al. 2021; Agrawal et al. 2021; Canto et al. 2017) for a few examples. Most of these methods rely on sum-of-squares or equivalent error minimization. We use a similar approach, but assume that a human provides a reasonably accurate initial estimate for the parameters. Thereafter, a gradient descent approach is used to obtain a more precise estimate of the system parameters using data.

Consider a general system of difference equations

$$x[k+1] = f(x[k], \rho[k]), \quad x[0] = x_0 \quad y[k] = h(x[k])$$

where x denotes the state of the system, y are the measured outputs, and ρ are parameters whose value needs to be estimated. We will assume that $x[0]$ is known and $\rho_j[k] > 0$ for all j and k . Suppose a reference trajectory $r[k]$ is known for the outputs $y[k]$. Then, $\rho[k]$ is found by minimizing a suitable cost-to-go:

$$\rho[k] = \operatorname{argmin} J[k], \quad J[k] = \sum_{j=k}^T \|y[j] - r[j]\|_Q$$

where T denotes the terminal time and $\|\cdot\|_Q$ denotes the weighted l_2 norm with $Q > 0$. We use the following update law to solve for $\rho[k]$, given an initial estimate:

$$\rho[k] \leftarrow \rho[k] - \sigma \nabla_{\rho[k]} J[k]$$

where the step size σ is chosen to ensure that $\rho_j[k] > 0$ for all j after the update. The gradient $\nabla_{\rho[k]} J[k]$ is computed using simulation-aided regression. We reduce the step-size until convergence; i.e., if, for a given step size, there is no improvement in the cost function. The parameter update stops until either the step size becomes smaller than a prescribed threshold or a certain maximum number of updates is reached.

2.3 Implementation

We implement a stock and flow model with the aforementioned time-delay logic in Vensim. Since the effect of Covid-19 correlates strongly with age, we divide the total population into three homogeneous groups: children (0-20 yrs old), young and middle-age adults (21-60 yrs old) and senior citizens (60+). This classification has the additional advantage that the roll out of vaccines has been carried out with in age-wise manner. Our compartment model has been illustrated schematically in Figure 1. In addition to the usual S, E, I, R and D bins, we have introduced additional bins for individuals who have received one or two doses of vaccines (V_{D1} and V_{D2} , respectively). We assume that vaccinated individuals experience a different trajectory upon exposure, which is why two separate 'S' stocks are introduced for those individuals together with the following paths (in green and blue, respectively) We assume that only unvaccinated people or those who have received a single dose become critically ill ('C' stock) and proceed to die.

2.4 Agent-Based Model

We virtually represent individual localities in Pune using a parameterized agent-based digital twin for simulation-based experimentation. The digital twin is built using the Enterprise Simulation Language (ESL) (Clark et al. 2017) and models five key aspects: individual citizens; localities and places; movements and contact; virus characteristics, and interventions. Each of these elements is modeled as an *agent* in our digital twin and endowed with city-specific characteristics and with the known characteristics of vaccines and variants (where applicable).

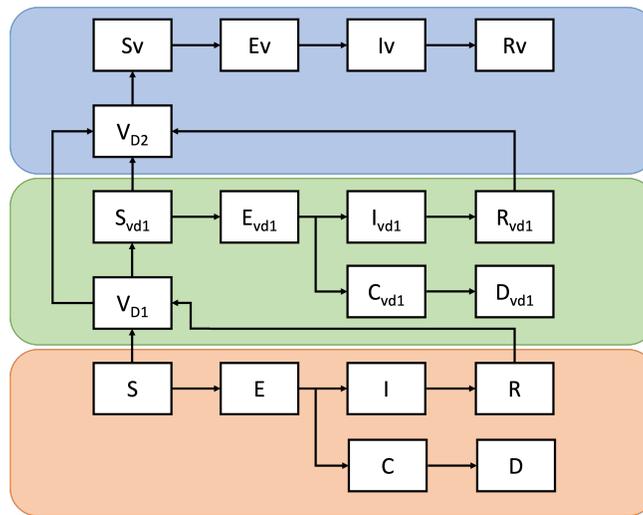


Figure 1: A schematic of the compartments for each age group in our stock and flow model. Stocks in orange, green and blue zones are for unvaccinated individuals and those with one and two doses, respectively.

- *Citizen model* captures the relevant individual characteristics and behavioural patterns of individuals in a locality. It includes age, gender, comorbidity, profession, household structure, vaccination status, and infection status of the individuals. To represent the city of Pune, 21 types of households and 15 professions archetypes are considered. The household structure ranges from a two-member family (1 male adult and 1 female adult) to a twelve-member family (3 male and 3 female adults, 4 children and 2 senior citizens). Citizen archetypes cover a wide range of professions with different behavioural patterns, such as back-office and customer-facing employees at an office, cab drivers, students, etc.
- *Place model* represents relevant places of a city where people frequently visit, spend time and come into contact with other individuals. We visualize a city as a set of prototypical districts (called wards in Pune), where each ward is a combination of well-to-do and slum localities with representative sets of houses and households, citizens (with different age, gender, comorbidities and archetype), and commercial places. We consider 20 different place archetypes including various modes of transport (i.e., own car, bus or shared cab) to precisely represent well-to-do and slum areas.
- *Movement model* captures the movement of individuals *within* a given place (e.g., in an office or school) as well as *between* places (e.g., home to car or office to shop) to understand how and when people come into contact with each other. Movements between the places are modeled as functions of an individual's profession, with certain random movements for living, socializing and entertainment. Movements *within* a given place are assumed to be inherently random.
- *Virus characteristics* define two aspects of Covid-19 variants: transmission dynamics and transition dynamics. Transmission is the propensity of a person to get exposed to a variant from an infectious person – it is modeled as a function of the individuals, the duration of contact and the infectivity of the variant. Exposed individuals go through a number of stages (i.e., asymptomatic, mildly infected, critical) before they recover or die. These transitions are modeled as functions of three factors – a) virus characteristics, i.e., severity and mortality rates of the variants, b) citizen characteristics, i.e., age, gender and comorbidity, and c) vaccination details, i.e., last vaccine dose and its efficacy.
- *Interventions* are considered along four dimensions: administrative, health-care related, social, and vaccination. Administrative interventions are related to citizen movements, (partial) closure of places and capacity constraints in cabs and public transport. Interventions from a health care standpoint include testing of mildly infected citizens (in addition to severely infected citizen), contact tracing and isolation of detected mildly infected citizens. Social interventions include mask use and social

distancing. Vaccines are modelled using two aspects a) vaccine administration policy and rate, and b) vaccine efficacy in terms of reducing infection, severity and fatality probability.

We capture the aforementioned elements as *agents* in our digital twin and simulate it by contextualizing all model elements with city-specific information and known facts about vaccines, variants, and interventions. Situations *emerge* in our digital twin as citizens with specific demographic characteristics, comorbidity, vaccine doses and infection history move around and interact with each other. They get exposed to a specific variant when they interact with an infectious person for a specific duration. Thereafter, they progress along one of the paths in Figure 1 depending on their individual characteristics, infection and vaccine histories, and the characteristics of the infecting variant. Each experiment using the digital twin involves simulation with the desired changes (i.e., a hypothesis) incorporated into the model. Further details on our ABM, including initialization, can be found in (Barat et al. 2021).

3 SIMULATION RESULTS

3.1 Prediction Using the ABM

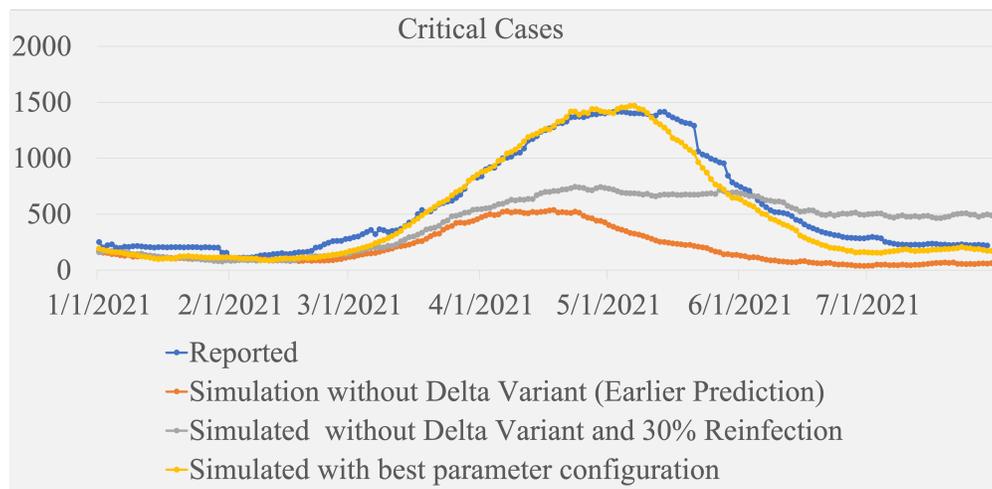


Figure 2: Summary plot showing the critical cases in Pune from Jan 1 2021 to July 30 2021. We plot the actual number of cases (blue), the number predicted by our prior model (orange), one illustrative case that was investigated in real-time at the onset of the Delta-driven wave when the characteristics of that variant were not understood (grey) and the the best fit (yellow) obtained *post facto*.

Figure 2 shows the simulated trends of critical cases in Pune for four different experimental settings (hypotheses). The actual reported critical cases are shown by the blue line, while the yellow line is the best-fit hypothesis. Each curve represents the average taken over 10 simulations. Chronologically, the first hypothesis that we made is shown in orange in Figure 2. This curve was obtained in the latter half of 2020 (Barat et al. 2021), at a time when the Alpha variant was dominant in India and the Delta variant was unknown (the first case was reported on Oct 5 2020). We will refer to that model as our 2020 model.

The 2020 model assumed (correctly) that the administrative restrictions on people’s movement would be lifted more or less entirely by Jan 31 2021. Until mid-February 2021, we note that the actual number of critical cases and those predicted by our 2020 model (orange) appear to follow a similar trend, as indicated by their slopes. The actual values differ by no more than 100 until about Feb 15. When the case numbers started increasing around Feb 20, we brought our predictions to the attention of the Pune city administration: while a wave appeared imminent, its peak amplitude was unlikely to exceed that of the first wave (not plotted here, see (Barat et al. 2021)) by a substantial amount. The orange peak in Figure 2

is approximately 70% as that of the first wave. This assertion appeared to be supported until early March: notice the flattening part of the blue segment between March 1 and 10 in Figure 2. Starting mid-March, the slopes of the 2020 model (orange) and the actual trajectory (blue) started diverging, by which point it was too late to avoid a sharp increase in the number of cases. The subsequent six-week lockdown which started on 4 April 2021 helped mitigate its amplitude.

In order to understand what may have led to such a sharp increase, we simulated a wide range of scenarios including (i) greater violation of Covid-appropriate behavior (CAB); (ii) greater mixing than anticipated; (iii) greater infectivity compared to the original Alpha variant; and (iv) reinfection and loss of immunity. The results of one such simulation, with 30% reinfection after 6 months and without accounting for a new variant, are shown in Figure 2 as the grey line. This curve severely underestimates the case load.

It is worth noting that the first cases of the Delta variant were detected outside India only in mid-February 2021, and it was designated as a variant of concern by the WHO only on May 11 2021. By this time, the second wave was receding and even the critical cases (whose trend lags that of new cases substantially) had started to decline. Thus, most of our hypotheses had to be assessed *post facto*. The best match from that analysis is shown as the yellow curve in Figure 2. It was obtained after including: (i) higher infectivity consistent with the Delta variant; (ii) 10% higher non-compliance with CAB compared to the pre-Delta baseline; and (iii) violation of home-quarantine in nearly 70% of the cases. Although institutional quarantine facilities were set up in mid-2020, these had been dismantled prior to the Delta-driven wave and not available for a large part of the Delta wave. *Post facto* analysis also suggested that the peak number of cases could have been closer to that of the first wave had the lockdown been imposed 10 days before it was eventually imposed on 4 April. The necessity for early imposition of lockdowns at the onset of a wave triggered by a new variant is examined further in Section 3.3.

3.2 Post Facto Analysis Using Stock and Flow Model

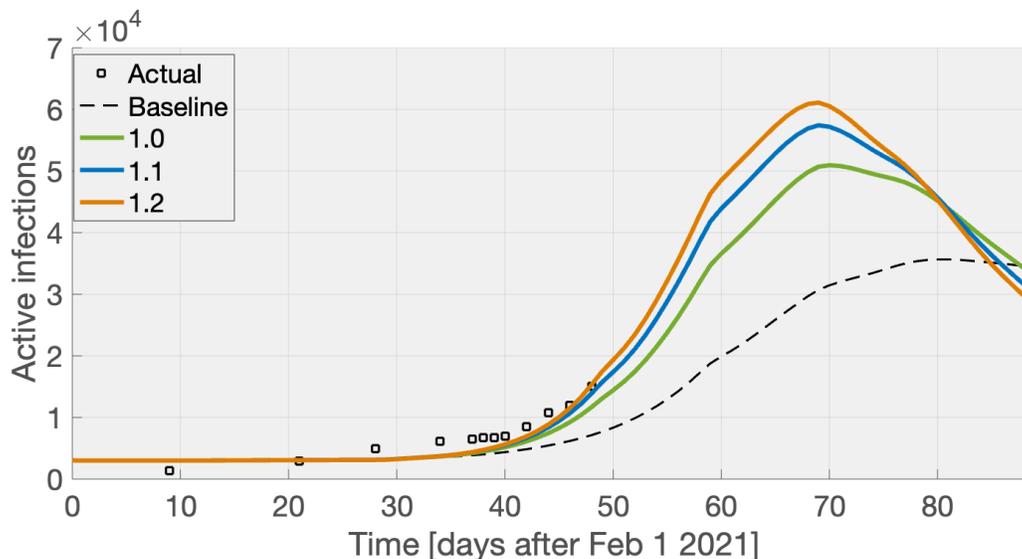


Figure 3: Expected track of active detected infections (the I compartment) starting 1 Feb 2021 and ending 30 Apr 2021, with actual data available only until 15 March. Baseline denotes a track which would yield a peak case load close to the first wave.

As described in the previous section, it became clear by mid-March that our 2020 model was likely to underestimate the trajectory of the second wave. Since an ABM takes a considerable amount of time to run and we had no more than a few days to make actionable recommendations, we attempted to derive

worst-case estimates using our SEIR model, based on (1). This was done by conducting a parametric study by varying the values of μ_1 and μ_2 in (1) after Jan 31 2021.

Figure 3 is a plot of the expected number of active cases for a range of values of mixing constants μ_1 (with $\mu_2 = 0.2\mu_1$ to account for better isolation of symptomatic patients). The baseline case yields a peak load comparable to the first wave. This particular plot has been plotted post facto and uses data available only until March 15 to illustrate the complexity of the situation. Even if the baseline case is ignored, the three cases (green, blue and red) yield peak case loads that differ from each other by up to 10,000 cases. Importantly, the three larger values of μ_1 considered in the plot (around 1.1) were over 50% higher than the baseline and difficult to justify in physical terms without including a much higher degree of infectiousness - recall that the Delta variant was labeled as a variant of concern by the WHO only in May.

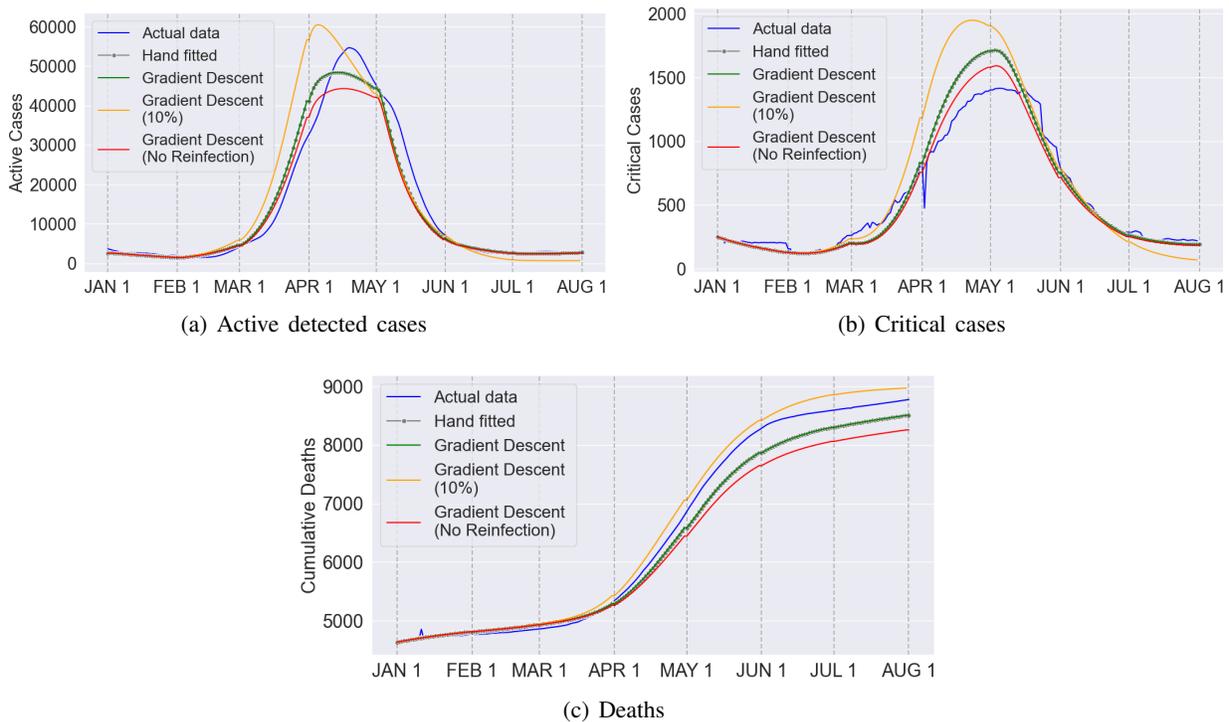


Figure 4: Trajectories obtained by fitting the rate parameters to the stock and flow Vensim simulator and the actual trajectory (blue). The parameters behind the yellow curve were obtained using gradient descent after a 10% initial perturbation to the manually fit parameters.

Other than an increased infectiousness, it was not clear how far the second wave was driven by loss of immunity in those who had previously contracted Covid, especially the earlier variants. In order to study the impact of reinfection, we use the parameter fitting approach in Section 2.2. More specifically, we fit the parameters μ_1 and μ_2 to the available data from 1 Jan 2021 - 31 July 2021 for two different cases: no reinfection and 30% risk of reinfection 6 months after contracting any variant of Covid. Figure 4 shows the resulting trajectories, including a trajectory where the parameters were fitted manually. In all of the cases in Figure 4, we find that the best parameter fit for any given variable does worse than its competitor for a different variable. For instance, while reinfection appears to better explain the trajectory of active infections, it does worse for that of critical cases. All of the parameter fits produce large errors around the peak, but generally do well on either side of the peak. This suggests that the impact of reinfection was low at least for the second wave - this assertion matches with the best fit obtained using the ABM (see Figure 2) which, notably, did not include a reinfection model.

3.3 State-Feedback Approach and Impact of Early Lockdowns

Although lockdowns or similar restrictions of a suitable degree are generally considered essential for controlling an epidemic like Covid-19, a prolonged lockdown can be economically unviable. A balanced approach to imposing lockdowns involves the use of *state feedback*, wherein a short-lived lockdown is imposed or relaxed in response to data to ensure that the healthcare system is not overwhelmed and the economic costs are also kept to a minimum.

Although one can determine the conditions for imposing or relaxing lockdowns by solving an optimal control problem, this approach may not always be practically feasible. For instance, if the optimal control law is not intuitively predictable, it may lead to popular discontentment and risks for the local economic productivity. Instead, lockdowns imposed on the basis of a handful of variables crossing pre-defined thresholds can be easier to justify in a real-world setting. One such feedback law that uses the rate of growth of cases takes the form

$$\text{Lockdown} = \begin{cases} \text{on,} & \langle I[k+1] - I[k] \rangle > -\varepsilon \\ \text{off,} & \text{otherwise} \end{cases} \quad (2)$$

where ε is a positive number and $\langle \cdot \rangle$ denotes a moving average computed over a suitable time window. Its obvious drawback is the possibility of incurring a higher loss (economic as well as lives lost), but it offers better predictability for the administrative authorities as well as the general public. This feedback law *does not* apply as a condition for imposing the very first lockdown: notice that the growth rate of cases is zero even before the epidemic, which is not grounds for imposing a lockdown. In order to determine the number of cases that must be reached before the first lockdown becomes necessary, we can use a simulation-based approach to compute the relevant metrics (such as the number of patients needing hospitalization) as functions of time.

In order to illustrate this approach, we conduct simulations on a cluster of 22000 people for two different values of the initial threshold for the number of infections. For simplicity, we model a lockdown by setting $\mu_1 = 5\mu_2 = \mu_1^{\text{lock}} = 0.14$, and the absence of lockdown by setting $\mu_1 = 5\mu_2 = 1$. The time histories of the instantaneous and the cumulative number of infected individuals (the latter almost equal to the number of recovered cases) is plotted in Figure 5 for two cases wherein the the number of infected individuals required to trigger *the first* lockdown is changed “slightly” (in relative terms).

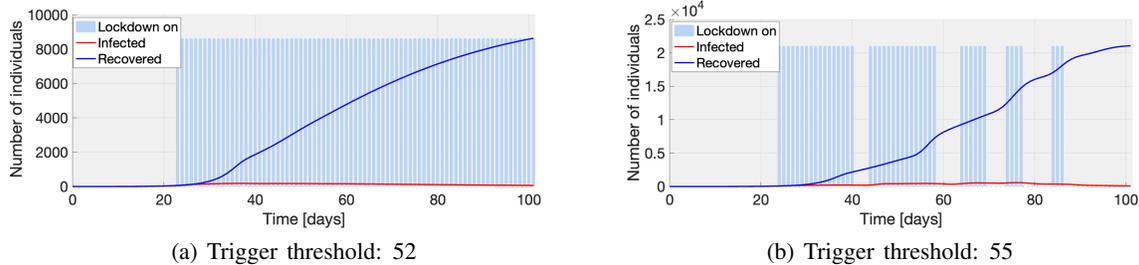


Figure 5: Time history for two different values of the number of infections needed to trigger the first lockdown. Blue regions indicate that a lockdown is on.

The plots in Figure 5 point to a qualitative change in the evolution of the epidemic around a trigger threshold of 53 cases. On the lower side of this number, the cases rise slowly and the maximum number of infected cases at any given time is 185. Notice that the lockdown persists for the entire duration of the simulation. When the trigger threshold exceeds 53, the lockdowns tend to be short-lived but the entire population is exposed within 100 days. The maximum number of infected cases is 550. This qualitative change is not a bifurcation in the traditional sense of the word since there is no underlying steady state of

interest. In both cases, the cumulative number of infected individuals eventually (but not necessarily over the 100-day simulation window) turns out to be identical, but the near-term behavior is vastly different. The qualitative variation comes across clearly in Figure 6, where a jump is discernible at a trigger threshold of 53 cases. The fraction of locked-down days reduces substantially thereafter, accompanied by an increase in the peak of infected cases. This plot suggests that an early lockdown may significantly temper the growth rate of the pandemic in the early stages.

The critical value of the number of infections in this example depends on a number of system parameters and it might thus be difficult to predict at the onset of a pandemic. Figure 7 shows, for instance, the peak active infections as functions of the trigger threshold and the values of μ_1^{lock} for four different values of ϵ (see (2)). The critical value of the trigger threshold can be discerned from the plots, informally speaking, as the cliff from the blue to yellow regions. The qualitative change is prominent when ϵ and $\mu_{1,2}^{\text{lock}}$ are small; i.e., when the lockdowns are highly effective, and the threshold for easing them is lenient.

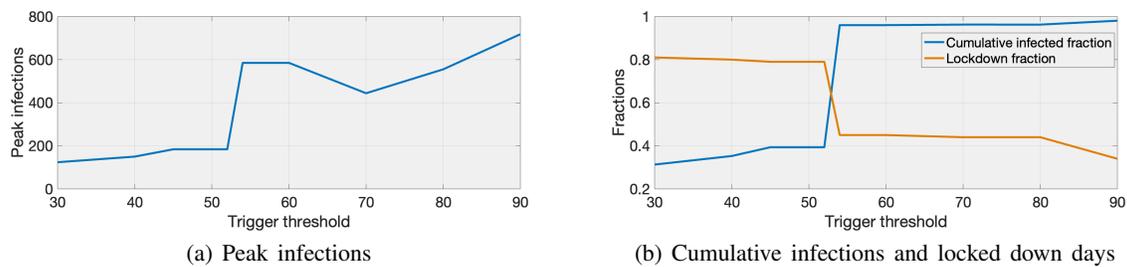


Figure 6: The peak number of infected cases, the cumulative infected population (as a fraction of the total), and the number of days spent in lockdown (as a fraction of 100) as functions of the trigger threshold.

4 CONCLUDING DISCUSSION

In this paper, we presented our work on modeling and simulation for the spread of Covid-19 in the Indian city of Pune during the first half of 2021. Some of the work reported here was carried out in “real-time” with significant time constraints as we attempted to help the Pune city administration assess the requirement for a lockdown in the month of March. We had developed several techniques which allowed us to support this assessment, including an SEIR model with time-delayed transitions and a high-fidelity agent-based model to assess the impact of microscopic behavior and interventions. Although our models indicated the presence of a highly infectious variant in March 2021, we did not have the requisite information to advise as such definitively. This demonstrates the risk associated with the use of any model or simulator in real-life, time-critical environments. Post-facto analysis using our models helped identify sociological factors which may have contributed to the rapid spread of the Delta variant in Pune, and it also suggested that reinfection made a minor contribution to the overall case load. Our results suggest the importance of early lockdowns, by identifying a bifurcation-like phenomenon in the growth rate of cases in the early stages of an epidemic. This might be an interesting topic for theoretical investigation into SEIR models.

REFERENCES

- Agrawal, M., M. Kanitkar, and M. Vidyasagar. 2021. “SUTRA: A Novel Approach to Modelling Pandemics with Undetected (Asymptomatic) Patients, and Applications to COVID-19”. In *60th IEEE Conference on Decision and Control (CDC)*. December 13th-15th, Austin, Texas, United States, 3531.
- Anderson, R. M., and R. M. May. 1992. *Infectious Diseases of Humans: Dynamics and Control*. New York: Oxford University Press, Inc.
- Barat, S., R. Parchure, S. Darak, V. Kulkarni, A. Paranjape, M. Gajrani, and A. Yadav. 2021. “An Agent-Based Digital Twin for Exploring Localized Non-pharmaceutical Interventions to Control COVID-19 Pandemic”. *Transactions of the Indian National Academy of Engineering* 6(2):323–353.

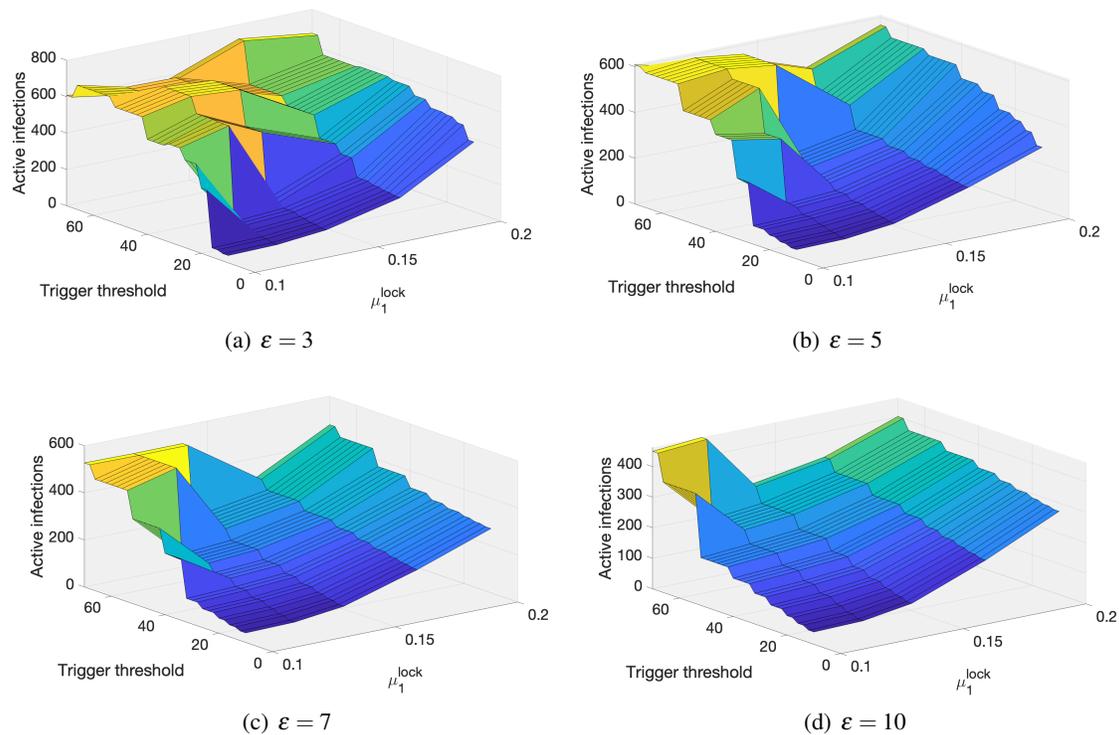


Figure 7: The peak infection load as a function of μ_1^{lock} and the trigger threshold for various values of ε .

- Brandeau, M. L. 2005. “Allocating Resources to Control Infectious Diseases”. In *Operations Research and Health Care*, edited by M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, 443–464. Boston, MA: Springer.
- Canto, B., C. Coll, and E. Sanchez. 2017. “Estimation of Parameters in a Structured SIR Model”. *Advances in Differential Equations* 33(1):1–13.
- Clark, T., V. Kulkarni, S. Barat, and B. Barn. 2017. “ESL: an Actor-Based Platform for Developing Emergent Behaviour Organisation Simulations”. In *Advances in Practical Applications of Agents and Multi-Agent Systems: The PAAMS Collection*, edited by Y. Demazeau, P. Davidsson, J. Bajo, and Z. Vale, 311–315. Cham: Springer.
- Cuevas, E. 2020. “An Agent-Based Model to Evaluate the COVID-19 Transmission Risks in Facilities”. *Computers in Biology and Medicine* 121:103827.
- Enns, E. A., and M. L. Brandeau. 2015. “Link removal for the Control of Stochastically Evolving Epidemics Over Networks: A Comparison of Approaches”. *Journal of Theoretical Biology* 371:154–165.
- Eubank, S., H. Guclu, V. Anil Kumar, M. V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. 2004. “Modelling Disease Outbreaks in Realistic Urban Social Networks”. *Nature* 429(6988):180–184.
- Ferguson, N. M., D. Laydon, G. Nedjati-Gilani et al. 2020. “Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID19 Mortality and Healthcare Demand”. Technical report, Imperial College London. <https://www.imperial.ac.uk/media/imperial-college/medicine/mrc-gida/2020-03-16-COVID19-Report-9.pdf>, accessed 20.09.2022.
- Keeling, M. J., and B. T. Grenfell. 2000. “Individual-Based Perspectives on R_0 ”. *Journal of Theoretical Biology* 203(1):51–61.
- Kermack, W. O., and A. G. McKendrick. 1927. “A Contribution to the Mathematical Theory of Epidemics”. *Proceedings of the Royal Society of London. Series A* 115(772):700–721.
- Kerr, C. C., R. M. Stuart, D. Mistry, R. G. Abeysuriya, K. Rosenfeld, G. R. Hart, R. C. Núñez, J. A. Cohen, P. Selvaraj, B. Hagedorn et al. 2021. “Covasim: An Agent-Based Model of COVID-19 dynamics and interventions”. *PLOS Computational Biology* 17(7):e1009149.
- Kuznetsov, Y. A., and C. Piccardi. 1994. “Bifurcation Analysis of Periodic SEIR and SIR Epidemic Models”. *Journal of Mathematical Biology* 32:109–121.
- Ma, W., Y. Takeguchi, T. Hara, and E. Beretta. 2002. “Permanence of an SIR Epidemic Model with Distributed Time Delays”. *Tohoku Mathematical Journal* 54:581–591.
- Mokhtari, A., C. Mineo, J. Kriseman, P. Kremer, L. Neal, and J. Larson. 2021. “A Multi-Method Approach to Modeling COVID-19 Disease Dynamics in the United States”. *Scientific Reports* 11:12426.

- Ozik, J., J. Wozniak, N. Collier, C. Macal, and M. Binois. 2021. "A Population Data-Driven Workflow for COVID-19 Modeling and Learning". *International Journal of High Performance Computing Applications* 35(5):483–499.
- Ray, D., M. Salvatore, R. Bhattacharya, L. Wang, J. Du, S. Mohammed, S. Purkayastha et al. 2020. "Predictions, Role of Interventions and Effects of a Historic National Lockdown in India's Response to the COVID-19 Pandemic: Data Science Call to Arms". *Harvard Data Science Review* (Special Issue 1):1–36.
- Riley, S., C. Fraser, C. A. Donnelly, A. C. Ghani, L. J. Abdu-Raddad, A. J. Hedley et al. 2003. "Transmission Dynamics of the Etiological Agent of SARS in Hong Kong: Impact of Public Health Interventions". *Science* 300.5627:1961–1966.
- Silva, P. C., P. V. Batista, H. S. Lima, M. A. Alves, F. G. Guimarães, and R. C. Silva. 2020. "COVID-ABS: An Agent-Based Model of COVID-19 Epidemic to Simulate Health and Economic Effects of Social Distancing Interventions". *Chaos, Solitons & Fractals* 139:110088.
- Wolfram, C. 2020. "An Agent-Based Model of COVID-19". *Complex Systems* 29(1):87–105.

AUTHOR BIOGRAPHIES

ADITYA A. PARANJAPE is a scientist in the Software Systems & Services Research Area at Tata Consultancy Services Ltd. He is also Honorary Lecturer in the Department of Aeronautics at Imperial College London. His research interests include control theory, nonlinear dynamics and multi-agent systems. His e-mail address is a.paranjape@imperial.ac.uk.

SOUVIK BARAT is a Principal Scientist at Tata Consultancy Services Research. He holds a PhD in Computer Science and his research interests include digital twin technology, modelling and simulation of complex systems, agent and actor technology, enterprise modeling, and business process modelling. His email address is souvik.barat@tcs.com.

ANWESHA BASU is a Researcher at Tata Consultancy Services Research. Her research at TCS includes digital twins and modelling and simulation of complex systems. Her email address is anweshha.basu1@tcs.com.

ROHAN SALVI is pursuing a Masters's in Information Science at the University of Illinois at Urbana-Champaign. His research interests lie in the area of Natural Language Processing, spanning text mining and text generation. His email address is rcsalvi2@illinois.edu

SUPRATIM GHOSH is a scientist with TCS Research. He earned his Ph.D and M. S. in Electrical Engineering, and M. A. in Mathematics from the Pennsylvania State University. His research interests include applications of control and optimization techniques to supply chain and robotics. His email address is supratim.ghosh2@tcs.com.

VINAY KULKARNI is a Distinguished Chief Scientist at Tata Consultancy Services Research. An alumnus of Indian Institute of Technology Madras, Vinay is a Fellow of Indian National Academy of Engineering. He is a Visiting Professor at Middlesex University London, and Indian Institute of Technology Jodhpur. His research interests include Digital Twins, Adaptive Enterprises, Model Driven Engineering, Artificial Intelligence, and Software Engineering. His e-mail address is vinay.vkulkarni@tcs.com.