

NBER WORKING PAPER SERIES

ADAPTIVE CONTROL OF COVID-19 OUTBREAKS IN INDIA:
LOCAL, GRADUAL, AND TRIGGER-BASED EXIT PATHS FROM LOCKDOWN

Anup Malani
Satej Soman
Sam Asher
Paul Novosad
Clement Imbert
Vaidehi Tandel
Anish Agarwal
Abdullah Alomar
Arnab Sarker
Devavrat Shah
Dennis Shen
Jonathan Gruber
Stuti Sachdeva
David Kaiser
Luis M.A. Bettencourt

Working Paper 27532
<http://www.nber.org/papers/w27532>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2020

We thank Jake Kramer for comments on this draft. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w27532.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Anup Malani, Satej Soman, Sam Asher, Paul Novosad, Clement Imbert, Vaidehi Tandel, Anish Agarwal, Abdullah Alomar, Arnab Sarker, Devavrat Shah, Dennis Shen, Jonathan Gruber, Stuti Sachdeva, David Kaiser, and Luis M.A. Bettencourt. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Adaptive Control of COVID-19 Outbreaks in India: Local, Gradual, and Trigger-based Exit Paths from Lockdown

Anup Malani, Satej Soman, Sam Asher, Paul Novosad, Clement Imbert, Vaidehi Tandel, Anish Agarwal, Abdullah Alomar, Arnab Sarker, Devavrat Shah, Dennis Shen, Jonathan Gruber, Stuti Sachdeva, David Kaiser, and Luis M.A. Bettencourt

NBER Working Paper No. 27532

July 2020

JEL No. I1

ABSTRACT

Managing the outbreak of COVID-19 in India constitutes an unprecedented health emergency in one of the largest and most diverse nations in the world. On May 4, 2020, India started the process of releasing its population from a national lockdown during which extreme social distancing was implemented. We describe and simulate an adaptive control approach to exit this situation, while maintaining the epidemic under control. Adaptive control is a flexible counter-cyclical policy approach, whereby different areas release from lockdown in potentially different gradual ways, dependent on the local progression of the disease. Because of these features, adaptive control requires the ability to decrease or increase social distancing in response to observed and projected dynamics of the disease outbreak. We show via simulation of a stochastic Susceptible-Infected-Recovered (SIR) model and of a synthetic intervention (SI) model that adaptive control performs at least as well as immediate and full release from lockdown starting May 4 and as full release from lockdown after a month (i.e., after May 31). The key insight is that adaptive response provides the option to increase or decrease socioeconomic activity depending on how it affects disease progression and this freedom allows it to do at least as well as most other policy alternatives. We also discuss the central challenge to any nuanced release policy, including adaptive control, specifically learning how specific policies translate into changes in contact rates and thus COVID-19's reproductive rate in real time.

Anup Malani
University of Chicago
Law School
1111 E. 60th Street
Chicago, IL 60637
and NBER
amalani@uchicago.edu

Satej Soman
University of Chicago
1155 E 60th St
Chicago, IL 60637
satej@uchicago.edu

Sam Asher
Johns Hopkins SAIS
1717 Massachusetts Ave NW
Office 743a
Washington, DC 20433
samuel.asher@gmail.com

Paul Novosad
Economics Department
Dartmouth College
6106 Rockefeller Center Room 301
Hanover, NH 03755
paul.novosad@dartmouth.edu

Clement Imbert
Department of Economics
The Social Sciences Building
The University of Warwick
Coventry CV4 7AL
United Kingdom
C.Imbert@warwick.ac.uk

Vaidehi Tandel
IDFC Institute
Mumbai
India
vaidehi.tandel@gmail.com

Anish Agarwal
Massachusetts Institute of Technology
32-D666
32 Vassar Street
Ray and Maria Stata Center,
Cambridge, MA 02139
anish90@mit.edu

Abdullah Alomar
Center for Complex Engineering
Systems (CCES)
King Abdulaziz City for Science and
Technology (KAST)
a.alomar@cces-kacst-mit.org

Arnab Sarker
Massachusetts Institute of Technology
Institute for Data, Systems, and Society
arnabs@mit.edu

Devavrat Shah
Massachusetts Institute of Technology
devavrat@mit.edu

Dennis Shen
Massachusetts Institute of Technology
deshen@mit.edu

Jonathan Gruber
Department of Economics, E52-434
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142
and NBER
gruberj@mit.edu

Stuti Sachdeva
Massachusetts Institute of Technology
ss5368@columbia.edu

David Kaiser
Massachusetts Institute of Technology
182 Memorial Dr
Cambridge, MA 02142
dikaiser@mit.edu

Luis M.A. Bettencourt
University of Chicago
1155 E 60th St
Chicago, IL 60637
bettencourt@uchicago.edu

Adaptive control of COVID-19 outbreaks in India: Local, gradual, and trigger-based exit paths from lockdown

Anup Malani, Satej Soman, Luis M. A. Bettencourt, Sam Asher, Paul Novosad, Clément Imbert, Vaidehi Tandel, Anish Agarwal, Abdullah Alomar, Arnab Sarker, Devavrat Shah, Dennis Shen, Jonathan Gruber, Stuti Sachdeva, David Kaiser*

May 12, 2020

Abstract

Managing the outbreak of COVID-19 in India constitutes an unprecedented health emergency in one of the largest and most diverse nations in the world. On May 4, 2020, India started the process of releasing its population from a national lockdown during which extreme social distancing was implemented. We describe and simulate an adaptive control approach to exit this situation, while maintaining the epidemic under control. Adaptive control is a flexible counter-cyclical policy approach, whereby different areas release from lockdown in potentially different gradual ways, dependent on the local progression of the disease. Because of these features, adaptive control requires the ability to decrease or increase social distancing in response to observed and projected dynamics of the disease outbreak. We show via simulation of a stochastic Susceptible-Infected-Recovered (SIR) model and of a synthetic intervention (SI) model that adaptive control performs at least as well as immediate and full release from lockdown starting May 4 and as full release from lockdown after a month (i.e., after May 31). The key insight is that adaptive response provides the option to increase or decrease socioeconomic activity depending on how it affects disease progression and this freedom allows it to do at least as well as most other policy alternatives. We also discuss the central challenge to any nuanced release policy, including adaptive control, specifically learning how specific policies translate into changes in contact rates and thus COVID-19's reproductive rate in real time.

Contents

1	Introduction	2
2	Probabilistic SIR model for adaptive control	5
2.1	Starting with a basic SIR model	6
2.2	Probabilistic SIR model	6
2.3	Updating parameters	7
2.4	SIR model with district- or ward-specific compartments	8

*University of Chicago (SS, LB, AM), MIT (AA, AA, AS, DS, DS, JG, SS, DK), IDFC (AM, VT), John Hopkins (SA), Dartmouth (PN), Warwick (CI), COVIN Working Group (all). We thank Jake Kramer and Gyanendra Badgaiyan for comments on this draft.

3	Parameters and simulations	11
3.1	Contact matrix	11
3.2	Simulation	13
3.3	Approximating lockdown and release	14
3.4	Parameter sensitivity	15
3.5	Caveats on SIR models and role of simulations	15
4	Simulated policies	16
4.1	Immediate release after May 3, 2020	16
4.2	Delayed release after May 31, 2020	16
4.3	Adaptive control starting May 3, 2020	16
4.3.1	Gradual change	16
4.3.2	Starting points	16
4.3.3	Response to disease triggers	17
4.3.4	Travel restrictions	17
5	Results of simulations	18
5.1	Tamil Nadu (TN), Andhra Pradesh (AP), Madhya Pradesh (MP)	18
5.2	Maharashtra and Mumbai	19
5.3	Robustness checks	22
6	Implementation of adaptive control	22
6.1	Possible triggers	24
6.2	Policies to regulate contact rates short of a lockdown	24

1 Introduction

After initial introductions of COVID-19 cases by foreign travelers in late January in Kerala and then again in early February in a range of other states, India responded aggressively to the outbreak. On March 24, the Prime Minister declared a lockdown (“Lockdown 1.0”) involving closure of all non-essential private activity and much governmental activity.¹ This response was more severe than in most countries, with Google Mobility reports indicating a decline in travel for retail by 70% and to workplaces by nearly 40%.² Although this response had a salutary effect on COVID-19 deaths, it also imposed a serious burden on the economy, especially among the poor. Wary of these negative consequences, the national government relaxed the lockdown in phases. First, it allowed States to ease back on April 22 (“Lockdown 2.0”). Then, on May 4, the government exited from full lockdown to a zoned lockdown (“Lockdown 3.0”).³

To prepare for this zoned stage, the Ministry of Home Affairs (MHA), with assistance from the Ministry of Health, Family and Welfare (MOHFW), issued an order that classifies

¹Ministry of Home Affairs Order No 40-3/2020-D and Annexure (March 23, 2020) (closing all government, commercial and social activity except essential services) (annexure available at https://prsindia.org/files/covid19/notifications/145.IND_Citizens_Guidelines_LockdownAnnexure_Mar_24.pdf). See also <https://www.nytimes.com/2020/03/24/world/asia/india-coronavirus-lockdown.html>.

²See https://www.gstatic.com/covid19/mobility/2020-05-29_IN_Mobility_Report_en-GB.pdf.

³Ministry of Home Affairs Order No. 40-3/2020-DM-I (A) and Consolidated Revised Guidelines (April 15, 2020) (guideline available at https://prsindia.org/files/covid19/notifications/2732.IND_Revised_Consolidated_Guidelines_April_15.pdf).

districts across the country based on their recent case counts and described smaller areas called containment zones within districts.⁴ Districts that had no confirmed cases in the last 21 days or more were classified as green zones and could begin all activity except air and inter-state bus and rail transport and opening of schools, hotels, malls, etc. Districts whose total active cases, doubling rate of cases, and testing rate met certain thresholds were classified as red zones and, on top of the restrictions imposed on green zones, were required to limit the use of local taxis and buses, opening of salons, allow only limited manufacturing and only in special zones or rural areas, limit e-commerce to essentials, and limit office work.⁵ All other districts were classified as orange zones and, on top of the green zone restrictions, had to limit local taxis and local bus capacity.⁶ Within each of these color-coded zones states may establish “containment zones” as small as a building, floor or even a home, which would be subject to even more severe restrictions than red zones.

This classification and associated restrictions continue until the end of May with 2 changes. First, after May 17, states would determine which districts fall in each color category (“Lockdown 4.0”).⁷ Second, starting June 8 the country will begin opening up – in a phased manner – all areas aside from containment zones in all color-coded zones (“Lockdown 5.0”).⁸

In this document we present a framework for adaptive control of the COVID-19 outbreak in India, as the nation exits from lockdown. The framework is compatible with the MHA approach during Lockdowns 3.0 and 4.0, but adds a rigorous and transparent methodology for counter-cyclical measures that can keep the outbreak under control. The hallmark of adaptive control, as we define it, is the re-authorization of socioeconomic activity that is gradual, is varied across locations, is and conditional on observed local progress in containing the epidemic. Importantly, adaptive control requires the ability to both expand and contract contacts that may transmit the disease as circumstances warrant. It is this ability to adjust up and down the degree of social distancing that allows a balance between the benefits of economic activity and the costs of COVID-19 infection.

Simulations based on a stochastic Susceptible-Infected-Recovered (SIR) model (and a Synthetic Intervention (SI) approach reported elsewhere) suggest that adaptive control is at least as good as alternative, more extreme or lockstep approaches. For example, simulations of an SIR model suggests that, given a wide array of parameters, adaptive control does better than either immediate release from lockdown or delaying lockdown by another 4 weeks. Doing better is measured either by (a) the area under the infection curve, i.e., the total number of person infected through the epidemic, or the spread with which prevalence of infection declines after reaching its peak. Examples for the period immediately after May 4 for different states are provided in Figure 1. There are two critical factors that produce this result. First, adaptive control has option value: states can mimic any other policy if it turns out to be better by adjusting the temporal scheduling of what activities they authorize. More extreme or lockstep

⁴Ministry of Home Affairs Order No. 40-3/2020-DM-I(A) (May 1, 2020).

⁵The Ministry of Home Affairs, in consultation with the Ministry of Health and Family Welfare, categorized district by color but did not clarify the exact standards used to distinguish Red and Orange Zone districts.

⁶States can code districts into more restrictive classes, but cannot lower their classification. The Union government can also recode a district as circumstances warrant.

⁷Ministry of Home Affairs Order No. 40-3/2020-DM-I(A) (May 17, 2020) (item 5i) (available at https://prsindia.org/files/covid19/notifications/IND_MHA_Lockdown_Extension_upto_May31_17052020.pdf).

⁸Ministry of Home Affairs Order No. 40-3/2020-DM-I(A) (May 30, 2020) (item 1 Phase I) (available at https://prsindia.org/files/covid19/notifications/IND_Lockdown_Extension_June30_31052020.pdf).

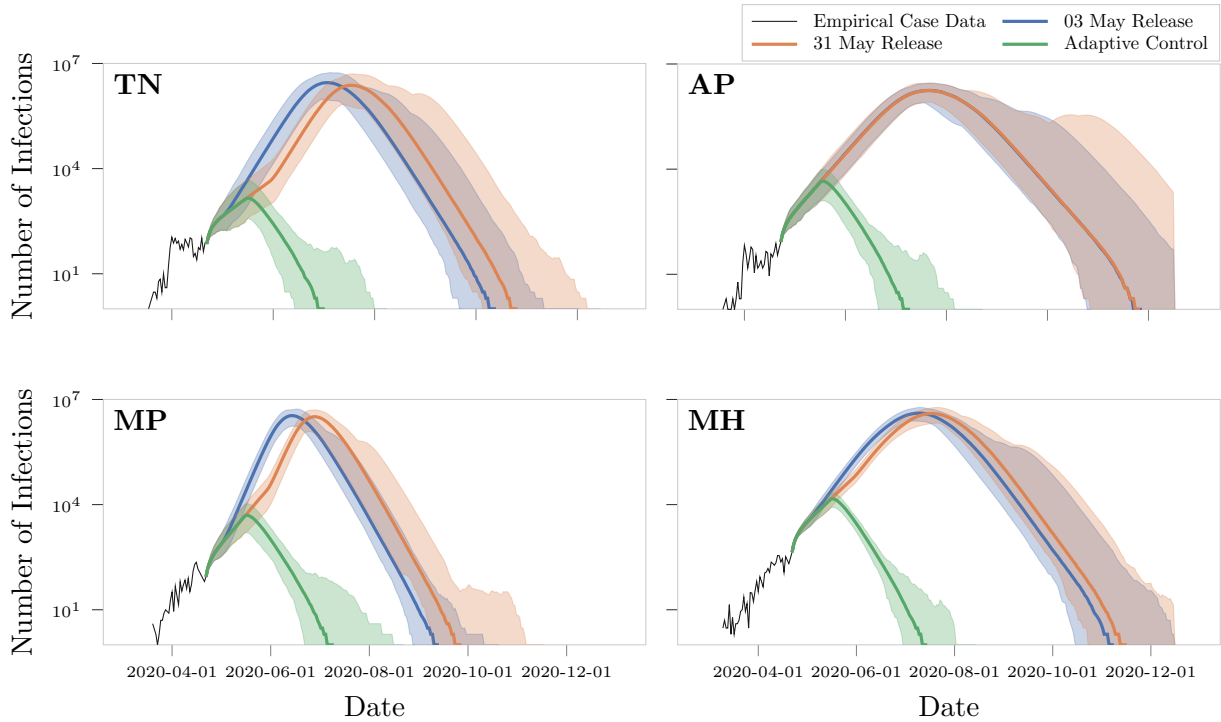


Figure 1: Impact of adaptive control, immediate and full release after May 3, and immediate and full release after May 31 on number of infections, by state, for select states: Tamil Nadu (TN), Andhra Pradesh (AP), Madhya Pradesh (MP), and Maharashtra (MH). Results are based upon 1000 simulations. Boundaries of shaded areas show outcome under minimum and maximum infections across simulations for each policy.

policies do not do that. Second, all policies are judged against the same set of parameters, i.e., the contact and recovery rate in the population. As long as adaptive control plays by the same “rules” it performs better as measured by total infections of speed with which the epidemic retreats.

Adaptive control policies are common in engineering and applied economics. They are characterized by the application of counter-cyclical or “feedback control” signals (policies) that drive a dynamical system to a desired operating point. They do so not by having a perfect model of the system, but more simply by varying control measures while observing the deviations to the desired outcome, and adapting up and down the signal to reach the desired result automatically. There are very simple and familiar examples of this approach in engineering, such as thermostats or vehicle cruise control, and there are more complex and policy driven approaches such as monetary policy controlling the supply of money with the intent of attaining high employment and/or lowering inflation and spurring economic growth.

A useful analogy for understanding how adaptive control differs from other policies are the steps involved in driving a car. When you are driving a car, you have three basic goals: reach a destination, while maintain a safe speed and avoiding obstacles, e.g., pedestrians. Likewise, when a government manages an epidemic outbreak, it has three goals. It wants to eliminate the epidemic, while avoiding extreme interventions, or letting outbreaks get out of hand so that they overwhelm the health system. Discovering the most effective policies and their degree of

severity may require adjusting their “speed” in response to observed cases in real time as well as towards achieving final objectives. The key way to achieve both is to gradually accelerate and be prepared to brake as required by both present observations and final objectives. If you accelerate too quickly (i.e., allow too much activity too fast), you will likely just have to brake most dramatically later (i.e., lockdown more harshly).

The adaptive control approach is a slight modification and a formalization of the approach the Ministry of Home Affairs (MHA) took in the May 1 order announcing Lockdown 3.0. First, like that order, adaptive control initially classifies districts based on risk. Low risk is the starting point for release. Adaptive control can use the MHA classifications – or others that are based on metrics such as reproductive rate (transmissibility) of the disease, trajectory of deaths, or the ratio of hospitalizations to available beds – to assign districts to initial risk levels. Second, just as MHA allowed states to use updated local information to re-assign districts to color-coded zones after May 16, the adaptive control approach recommends continuously revisiting risk levels, as new information becomes available, and adjusting appropriate controls. Third, in both the MHA order and adaptive control, the level of socioeconomic activity allowed or prescribed is a function of the current risk level. A natural progression of adaptive control is that one should continuously estimate how controls applied by different states have affected mobility and disease progression. Through this learning process, adaptive control can become more effective than other approaches.

Section 2 presents a probabilistic SIR model that is the “engine” underneath the simulations of adaptive control results previewed in Table 1. (In a separate technical paper, Agarwal et al. (2020) present simulations from an synthetic interventions model, which is based on the synthetic control method used in causal inference to demonstrate that the logic of adaptive control does not depend on specific model choices.⁹) Section 2 also discusses caveats about modeling COVID outbreaks. Section 4 defines the different policy options that we simulate. Importantly, it formally describes the way we implement adaptive control in our analysis. Section 5 presents our simulation results. While all states do as well with adaptive control as with immediate or delayed full release, the advantage of adaptive control varies across states depending on the local progression of COVID as of the date the state exited lockdown. Section 6 discusses how to implement adaptive control in practice. We also discuss the possible triggers that can be used as well as how states can encourage social distancing short of total lockdown.

2 Probabilistic SIR model for adaptive control

We simulate a variant on the Susceptible-Infected-Recovered (SIR) model to evaluate adaptive control. We first present a simple SIR model and then describe our modifications to it.

⁹The technical paper is available at <https://arxiv.org/abs/2005.00072>.




Figure 2: Schematic of basic SIR compartmental model.

2.1 Starting with a basic SIR model

A standard, deterministic SIR model for a focal population with internal homogeneous mixing has the structure set forth in Figure 2 and is characterized by the following set of equations.¹⁰

$$\frac{dS}{dt} = -\beta \frac{S}{N} I \quad (1)$$

$$\frac{dI}{dt} = \beta \frac{S}{N} I - \gamma I \quad (2)$$

$$\frac{dR}{dt} = \gamma I \quad (3)$$

Individuals start in the susceptible state S . They interact with individuals who are infected at a rate of I/N , where I is the population that is infected and N is the size of the overall population. Conditional on interaction, the probability a susceptible person becomes infected is β . Homogeneous mixing means that individuals have an equal probability of interacting with every other member of the population and the probability of infection is constant across all interactions with infected persons. Once infected, an individual recovers at a rate γ , which implies that the expected time to recovery is $1/\gamma$. Once recovered, a person has no risk of reinfection.

2.2 Probabilistic SIR model

We build on the basic SIR model in a few ways. First, we make the model probabilistic. To build in this structure, we follow Bettencourt and Ribeiro (2008) and start by integrating (2) over a period τ between t and $t + \tau$ to get

$$I(t + \tau) = I(t) \exp \left[\gamma \int_t^{t+\tau} \left(R_0 \frac{S(s)}{N(s)} - 1 \right) ds \right] \approx I(t) \exp [\tau \gamma (R_t - 1)] \quad (4)$$

Then we define the branching parameter as

$$b(t) = \exp \left[\gamma \int_t^{t+\tau} (R(s) - 1) ds \right] \approx \exp [\tau \gamma (R_t - 1)] \quad (5)$$

Now we can write $I(t + \tau) \approx b(t)I(t)$.

¹⁰Although we build on the basic SIR model by specializing it to smaller geographies, one could add other variations. For example, one could add an exposed category, E between the susceptible and infected state, or keep track of the number of deceased, D . Alternatively or additionally, one can add a quarantine compartment Q (Chatterjee et al., 2020) such that some subset of those infected end up in quarantine and therefore have a lower probability of interacting with the susceptible population. Yet another alternative is that one could fragment the susceptible and infected classes by age (Singh and Adhikari, 2020). We do not make these changes either because our prior robustness analysis (Figuroa et al., 2020) suggested one does not gain much additional predictive or planning value from these variations or they are reliant on parameters, e.g., mixing across age groups, that are very hard to measure and that may in fact degrade predictability. Our model is most similar to Ray et al. (2020), also known as the University of Michigan model of COVID in India.

While our purpose here is to define a stochastic process for simulation, we also want to keep an eye out for the ease of estimating of parameters. It will be easier to estimate $b(t)$ if we do not have to keep track of recoveries or deaths. Therefore, we focus on $T(t) = I(t) + R(t)$, which we define as the total number of cases up until time t , then (2) and (3) imply the change in the number of new cases follows $dT/dt = -\beta(S/N)I$. To take this equation to data, we will want to discretize the last equation at times $t + \tau$ and at t :

$$\frac{T(t + \tau) - T(t)}{\tau} = \beta \frac{S(t + \tau)}{N(t + \tau)} I(t + \tau) \approx \beta \frac{S(t)}{N(t)} b(R_t) I(t) \quad (6)$$

$$\frac{T(t) - T(t - \tau)}{\tau} = \beta \frac{S(t)}{N(t)} I(t) \quad (7)$$

assuming $S(t)/N(t)$ is piece-wise constant over the period τ . Dividing (6) by (7), we get

$$\Delta T_{t+\tau} = T_{\tau} - T_t = b(R_t) \Delta T_t \quad (8)$$

Either (4) or (8) can be the basis for estimating $b(t)$, a topic we will turn to momentarily.

For our present purposes, (8) can be used to make the SIR model probabilistic by assuming that

$$\Delta T(t + \tau) \sim P(b(R_t) \Delta T(t))$$

where $P(\lambda)$ is a discrete probability distribution with mean $\lambda = b(R_t) \Delta(t)$ where $b(t)$ is from (5). We assume P is Poisson because that is the least informative/maximum entropy distribution if the only available information the average number of cases per day. As we shall see when we discuss data, what we have is a moving average of cases. To implement this stochastic process, we will need not just data on the average number of total cases but also an estimate of $b(R_t)$. We turn to this next.

2.3 Updating parameters

A second way we build on the basic SIR model is to continually update our estimate of key parameters. Following Bettencourt and Ribeiro (2008), we use Bayes Rule to update beliefs about the branching parameter $b(t)$ or the effective reproduction number R_t . For example, we could write

$$\begin{aligned} P(R_t | \Delta T(t + \tau) \leftarrow \Delta T(t)) &= \frac{P(\Delta T(t+\tau) \leftarrow \Delta T(t) | R_t)}{P(\Delta T(t+\tau) \leftarrow \Delta T(t))} P(R_t) \\ &= L(\Delta T(t + \tau), \Delta T(t), R_t) P(R_t) \end{aligned} \quad (9)$$

where $P(R_t)$ is our prior on R_t , $L(\Delta T(t + \tau), \Delta T(t), R_t)$ is our likelihood function for data given R_t , and the left hand side is our posterior on R_t . We can estimate this using Bayes methods.¹¹ For that we use Uniform[0,3] as a prior on R_t . The denominator of (9) does not contain R_t , so it is just a normalizing constant. The Bayesian posterior can provide provide mean, median, modes and credible intervals for R_t .

¹¹This is not the only way to update parameters. Wallinga and Teunis (2004) and Cori et al. (2013) provide a method that does not depend on an SIR, but does require assumptions about the generation interval, i.e., the time between symptom onset in one infected person and symptom onset by a person infected by the first person. Since we are working in an SIR framework, Bettencourt and Ribeiro (2008) seem a more natural model for estimation.

Taking this method to data leaves a few choices. First, one could construct a posterior for R_t given I instead of for ΔT using (4). Which posterior to work with depends on the data that is most easily available. This may vary by application. Second, one could construct posteriors for R_t or for $b(R_t)$. Which one to work with depends on the nature of information on the recovery rate γ . If this is just an assumed parameter, either approach works as R_t and $\log b(R_t)$ are an affine transforms of each other. An advantage of directly estimating b is that it is a cleaner function of I or ΔT . An advantage of working with R_t is that it has a better known interpretation. Third, one must choose a prior. If the prior is uninformative, one can just use the Poisson maximum likelihood estimator to obtain posteriors.¹² If the conjugate prior is Gamma, the posterior will also, conveniently, be a Gamma distribution.

Fourth, data on cases has daily variation that may not reflect only cases. For example, a jurisdiction may report some cases late. To address this, it makes sense to employ a few periods of data to estimate, e.g., R_t . One way to do this is to use an informative prior. The prior will contain information from prior periods. An alternative approach is to use a few periods of data and use an uninformative prior before then. Either way, our assumption that $S(t)/N(t)$ is piece-wise constant requires us to assume a period for which this is true. One can think of this as two choices: a period over which we “smooth” our estimate of R_t and a method of implementing that smoothing.

Given the data accessible to us, here are the choices we make. We work with information on infected persons because we have access to recovered cases and deaths. For our main simulation, we start with an uninformative prior on $b(R_t)$. We use a 5 day “smoothing” window, i.e., we observe 6 days of data on infections and from that estimate a 5 day average for the current reproductive rate. Specifically, for each of the 6 past days we calculate $\log(b_t) = \log(I_t/I_{t-1})$. We then calculate the mean of the posterior as $\log(\hat{b}_t) = \frac{1}{5} \sum_{i=0}^4 \log(b_{t-i})$. We estimate R_t using $\log(\hat{b}_t)/\gamma + 1$. When we want to project out $k \in [1, 14]$ days from t , we use $[\exp(\log(\hat{b}_t))]^k I_t$.

Figure 3 shows estimates of $R_{t,j}$ from available case data. Additionally, Figure 4 shows estimates for predicted and observed net new cases, and preview an updating process using a Gamma prior, which we are in the process of implementing. Results are aggregated at the state level (j), for the 4 states we examine in this paper from March 11 to May 11, 2020.

2.4 SIR model with district- or ward-specific compartments

A third way we build on the basic SIR model is to specialize it to different geographies, e.g., districts in a state or wards in a city.¹³ The specific equations that describe our model are:

$$\frac{dS}{dt} = \dot{S}_j = -\frac{S_j}{N_j} \sum_{j'} \beta_{j,j'} I_{j'} \quad (10)$$

$$\frac{dI}{dt} = \dot{I}_j = \frac{S_j}{N_j} \sum_{j'} \beta_{j,j'} I_{j'} - \gamma I_j \quad (11)$$

$$\frac{dR}{dt} = \dot{R} = \gamma \sum_{j'} I_{j'} \quad (12)$$

¹²This also highlights the a frequentist approach that can be taken to update parameters.

¹³In some sense, people building SIR models for different countries already do this. The difference between what they do and what we do is that we allow mixing between different geographies. This is important because our geographic units are smaller than countries and there is, in fact, more mixing across geographic units as they become smaller, such as between city wards.

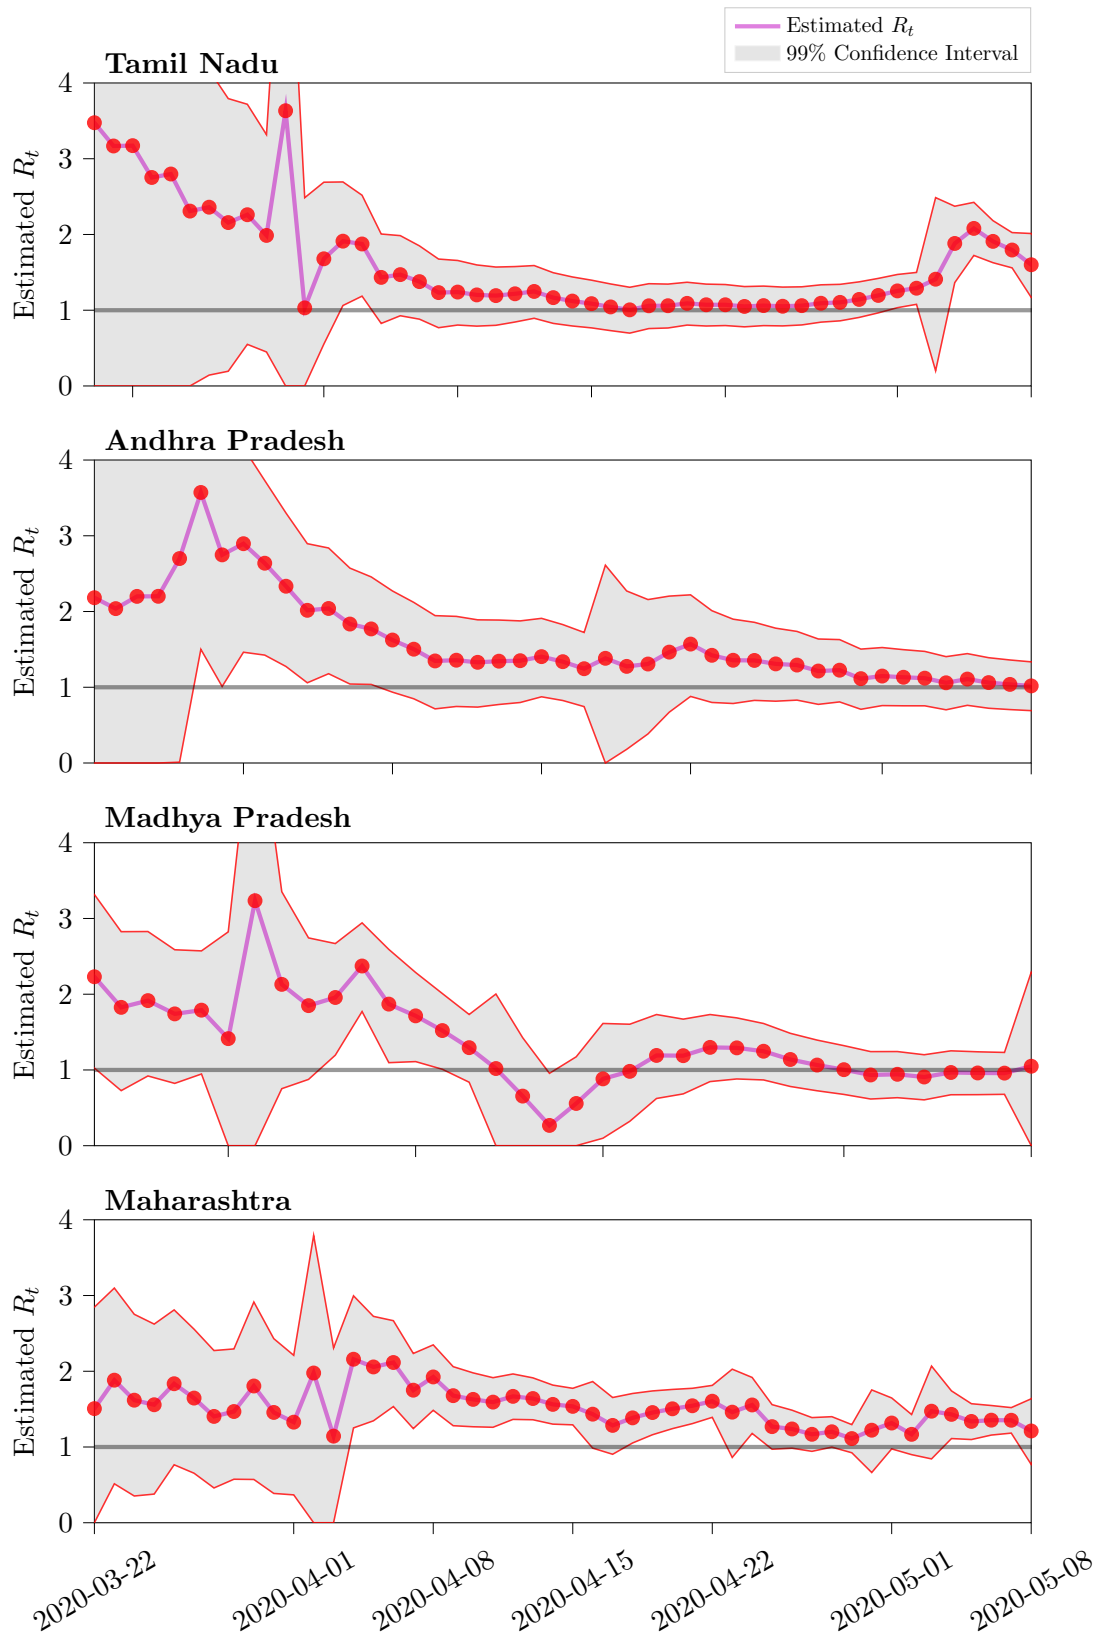


Figure 3: Estimated reproductive rate, by state. (Data range from March 11, 2020, to May 11, 2020. Code and files available at <https://github.com/mansueto-institute/covin-c2-adaptive-control-wp>.)




Figure 4: Observed and predicted new net daily cases, by state. (Data range from March 11, 2020, to May 11, 2020. Code and files available at <https://github.com/mansueto-institute/covin-c2-adaptive-control-wp>.)

where j indexes geographic sub-units, such as districts, S_j is the number of susceptible persons in district j , N_j is the population of district j , $\beta_{j,j'}$ is the contact rate between susceptibles in district j and infected persons in district j' , and $I_{j'}$ is the number of infected persons in location j' . The model does not need to partition the population of recovered persons by location because they do not affect the progression of the disease short of herd immunity and, in any case, can be captured entirely by $R_j = N_j - S_j - I_j$. While one can add a compartment to reflect the death of infected persons, this is not necessary if the death rate is low, as it is for COVID. Likewise, one need not add births and non-COVID deaths to the model as that rate of those two events in the time scales that COVID is operating (i.e., weeks or months versus years) is very low.

We model COVID at the state, district or ward level in India. State and district level modeling each assume that there is no movement across state lines. That is factually incorrect. However, the rate of movement across at least state lines is small relative to the movement within states. Our state level modeling assumes homogeneous mixing within states. District level modeling assumes there may be homogeneous mixing within districts and movement across districts. Ward level modeling within cities assumes there is no movement across city boundaries but there is homogeneous mixing within wards and movement across wards boundaries.

3 Parameters and simulations

Here we discuss how we estimate or set parameters for the model, other than our updated estimates for R_t . We begin with mixing among populations from different districts or wards. Second, we lay out the steps of our simulation. Third, we discuss how we approximate lockdown and release from lockdown, two extremes of our policy space. Finally, given that projections from SIR models are exponentially sensitive to specifications of parameters, namely the contact and the recovery rates (Figueroa et al., 2020), we discuss how test the sensitivity of our results to changes in these parameters.

3.1 Contact matrix

Our main modification to the standard stochastic SIR model is to allow separate compartments for different districts or wards. Later we will consider the progression of the disease under different policies on social distancing. Under one of those policies we will allow mixing between districts. When we do so, our estimate of the ratio of contacts between (a) district j and district j' and (b) j and j'' will be proportional to the ratio of (c) labor flows from j' to j and (d) labor flows from j'' to j prior to the outbreak, using data on labor migration from the Indian Census and National Statistical Survey.¹⁴ When we do city-level models, our relative rates of inter ward travel will be analogously proportional to data on passenger rail traffic from the Centre for Railway Information Systems.

In order to implement, e.g., inter-district migration in our simulations, at each time step t , we run the stochastic epidemiological model for each district in a state and obtain an estimate of the district's population, $N_j^{\text{epi}}(t) = S_j(t) + I_j(t) + R_j(t)$. We assume a proportion $\mu \cdot M_{jj'}$ of district j 's population travel to district j' , where $M_{jj'}$ is the empirical migration flow from Census data and $\mu \in [0, 1]$ is a proportionality factor that changes under different policy regimes. Our model reflects these migrations by taking a proportion of a district's population $\mu \cdot M_{jj'} \cdot N_{j,t}$,

¹⁴Specifically, we use data from the 2001 and 2011 Census and the 2007-2008 Employment Survey.

removing them from the source district's population, and redistributing them across to other districts j' . The infected populations are just the fraction of flows from each district that are infected. The adjusted numbers for total and infected population are used to evolve the model at time $t + 1$.

$$\Delta N_{j,t}^{\text{out}} = \sum_{j'} \mu M_{jj'} N_{j,t} \quad (13)$$

$$\Delta N_{j,t}^{\text{in}} = \sum_{j'} \mu M_{j'j} N_{j',t} \quad (14)$$

$$N_{j,t+1} = N_{j,t}^{\text{epi}} - \Delta N_{j,t}^{\text{out}} + \Delta N_{j,t}^{\text{in}} \quad (15)$$

$$\Delta B_{j,t}^{\text{out}} = \frac{I_{j,t}}{N_{j,t}} \sum_{j'} \mu M_{jj'} N_{j,t} \quad (16)$$

$$\Delta B_{j,t}^{\text{in}} = \sum_{j'} \frac{I_{j',t}}{N_{j',t}} \mu M_{j'j} N_{j',t} \quad (17)$$

$$\Delta B_{j,t+1} = \Delta B_{j,t}^{\text{in}} - \Delta B_{j,t}^{\text{out}} \quad (18)$$

This update changes both the total populations in each district, but also the number of infectious individuals, generating introductions from the outside as well as removals.

Presently, we use historical data for contact matrices. However, contact matrices do change during outbreaks. For example, some laborers may leave cities to evade the higher infection risk there or because work is no longer available. Upon release, they may be eager to return. Going forward we will plan to adjust the contact matrix to reflect these outbreak-specific movements by incorporating more recent mobility data from Google and perhaps ad tech data. Gruber & Sachdeva (2020) provide an evaluation of these data in the Indian context.

However our contract matrix is estimated, it requires us to alter how we make the SIR model probabilistic. Specifically, we have to account for infections from migrants, which are not cases of community transmission. We treat the migratory flows as a reservoir of infected agents interacting with the pool of residents in the geographic unit. Following Bettencourt and Ribeiro (2008), we denote the number of new introductions per unit of time as dB/dt . The equation for the time-evolution of the number of cases then becomes

$$\frac{dI}{dt} = \left[\beta \frac{S(t)}{N(t)} - \gamma \right] \cdot I(t) + \frac{dB}{dt}$$

Using Euler's method of integrating factors, we can write the solutions of this first-order inhomogeneous ordinary differential equation as:

$$\begin{aligned} I(t + \tau) &= \frac{1}{\exp\left(\int_t^{t+\tau} \left[\gamma - \beta \frac{S(t_1)}{N(t_1)}\right] dt_1\right)} \left[\int_t^{t+\tau} \exp\left(\int_t^{t_2} \left[\gamma - \beta \frac{S(t_1)}{N(t_1)}\right] dt_1\right) \cdot \frac{dB}{dt_2} dt_2 + I(t) \right] \\ &= \exp\left(\int_t^{t+\tau} \left[\beta \frac{S(t_1)}{N(t_1)} - \gamma\right] dt_1\right) \left[I(t) + \int_t^{t+\tau} \exp\left(\int_t^{t_2} \left[\gamma - \beta \frac{S(t_1)}{N(t_1)}\right] dt_1\right) \cdot \frac{dB}{dt_2} dt_2 \right] \\ &= b(R_t) \left[I(t) + \int_t^{t+\tau} \exp\left(\int_t^{t_2} -\gamma [R_t - 1] dt_1\right) \cdot \frac{dB}{dt_2} dt_2 \right] \\ &= b(R_t) [I(t) + \Psi(t, \tau, B)] \end{aligned} \quad (19)$$

where Ψ denotes the integral that evolves the migratory introductions from t to $t + \tau$.

The total number of new cases evolves according to:

$$\frac{dT}{dt} = \beta \frac{S(t)}{N(t)} I(t) + \frac{dB}{dt}$$

By discretizing 19, we can write the change in the number of cases in a time-step τ as

$$\begin{aligned} \Delta T(t + \tau) &= T(t + \tau) - T(t) \\ &= \Delta B(t + \tau) + b(R_t) [\Delta T(t) - \Delta B(t) + \tau \gamma R_t \Psi(t, \tau, B)] \end{aligned} \quad (20)$$

Assuming the vector of introductions is approximately constant across the time-step τ , we can approximate the integral Ψ as:

$$\begin{aligned} \Psi(t, \tau, B) &\approx \tau \frac{dB(t)}{dt} \equiv \Delta B(t) \\ \implies \Delta T(t + \tau) &= b(R_t(t)) \Delta T(t) + [1 - b(R_t(t)) + \tau \gamma b(R_t(t)) R_t(t)] \Delta B(t) \end{aligned} \quad (21)$$

To introduce stochasticity to the model, we treat the right-hand side of 21 as the rate parameter for a Poisson draw to generate the number of net new cases in a time-step.

Additionally, now that we allow changes in population due to migration, we can also easily introduce removals from the population due to deaths. We will assume that infected individuals die at a rate of $m = 0.02$. This will affect both the total population and the population of recovered persons. We incorporate this by assuming that new deaths $\Delta D(t + \tau)$ is also Poisson with mean

$$m \cdot \gamma \cdot (I(t + \tau) + \Delta T(t + \tau)) \quad (22)$$

and that newly recovered $\Delta R(t + \tau)$ is Poisson with mean

$$(1 - m) \cdot \gamma \cdot (I(t + \tau) + \Delta T(t + \tau)) \quad (23)$$

Together (21), (22), and (23) define the whole stochastic process with both migration and death in our model.

3.2 Simulation

We are now ready to fully describe our simulations. The following is a step-by-step description of our algorithm. For each jurisdiction j , do the following steps. (We suppress j to simplify the exposition.)

1. Pick a start date in the real world for the analysis. This is typically after an infection has taken hold, but before it has spread widely. Call this start date $t = 0$.
2. Measure $(S_0, I_0, N_0, \Delta T_0)$ from actual data, where ΔT_0 is new cases minus new recoveries and new deaths from time $t = -1$ to 0.
3. Estimate $\log(b_0)$ from $(\Delta T_0, S_0, N_0)$ using (9) for $\log(b_t)$ rather than R_t . We explain details at the end of 2.3. Deduce $R_{t,0}$ using $b_0 = \exp(\gamma(R_{t,0} - 1))$ and assuming $\gamma = 0.2$.
4. Implement re-distributions of population using contact matrix as described in Section 3.1. From this get $(\Delta N_0, \Delta B_0)$.

5. Draw ΔT_1 from Poisson with mean $b_0 \Delta T_0 + [1 - b_0 + \gamma R_{t,0}] \Delta B_0$.
6. Draw deaths ΔD_1 from Poisson with mean $m\gamma(I_0 + \Delta T_1)$ where $m = 0.02$. Define deaths at $t = 1$ as $D_1 = D_0 + \Delta D_1$.
7. Draw recovered ΔR_1 from Poisson with mean $(1 - m)\gamma(I_0 + \Delta T_1)$. Define the number recovered at $t = 1$ as $R_1 = R_0 + \Delta R_1$.
8. Define $N_1 = N_0 + \Delta N_0 - \Delta D_1$; $S_1 = N_0 + \Delta N_1 - \Delta T_1$; $I_1 = I_0 + \Delta T_1 - \Delta D_1 - \Delta R_1$.
9. Estimate $\log(b_1)$ from simulated $(\Delta T_1, S_1, N_1)$ using (9) for $\log(b_t)$ rather than R_t . We follow the steps described at the end of 2.3 with one change: we do not average b_t over the last 5 days. There is no need to average over simulated data, which is smoother than actual data. Deduce $R_{t,1}$ using $b_1 = \exp(\gamma(R_{t,1} - 1))$ and assuming $\gamma = 0.2$.
10. Repeat from step 5 except use simulated results rather than real data for periods 1 on.

We ran 1000 simulations of the SIR model for each state and 15,000 simulations of the SIR model for the city of Mumbai. In addition, we ran 1,000 simulations of each parameter variation we considered to check robustness of our conclusions about adaptive control.

In each simulation, we held parameter values constant across policies up until April 23, the latest day on which we had actual data for our analysis. We used projections from April 23 until May 4, when the different policies went into effect. After May 4, we allowed each policy to generate new updated branching parameters to reflect the simulated progression of the disease under each policy regime.

Our simulations can be found at the following repository:

<https://github.com/mansueto-institute/covin-c2-adaptive-control-wp>. Further work, such as updated and improved estimators, is centralized at <https://github.com/mansueto-institute/adaptive-control>.

3.3 Approximating lockdown and release

In order to simulate policy counterfactuals, we need to define different policies in terms of model parameters. At the extreme of the policy space are lockdown and full release. We define these here and then discuss intermediate steps when we define adaptive control later.

We define the branching parameter $b_{j,t}^L$ in location j during lockdown L as the branching parameter on April 15. India entered lockdown on March 24. However, it took some time for the lockdown to be implemented. Lockdown was partially released on April 21. We want a date that is closer to the partial release so that it reflect fuller implementation of the lockdown and incorporates addition post-lockdown observations to increase the precision of our estimate. Because we estimate branching parameters using actual case data and because it takes roughly 7 days for contact to generate infections that show up in COVID case counts, we use case data on day t to measure the branching parameter on day $t - 7$. Because the simulations in this paper use data up until April 23, that means the latest we can estimate lockdown associated branching parameters is April 15.

We define the branching parameter $b_{j,t}^R$ after full release as the branching parameter on March 23. When India releases from lockdown, we do not expect that contact rates will rise to the level observed before COVID, e.g., on May 4, 2019. The reason is that people are aware of

COVID and will engage in voluntary social distancing. Although how much they will voluntarily distance is unknown, a reasonable approximation is to look at a date on which the population was aware of the seriousness of COVID but was not required by law to, e.g., shelter-at-home. A good candidate for that is March 23. On March 22, the Prime Minister declared a Janata Curfew, which translates as voluntary curfew. Individuals were required to stay at home for 14 hours. That exhortation did not apply on March 23. And, on March 24, the government declared a legally mandatory curfew. Therefore, March 23 seems a reasonable date for modeling voluntary social distancing. Given the lag between contact and infection, we use case data from March 31 to approximate the branching parameter on March 23.

3.4 Parameter sensitivity

Although we attempt to “force” the SIR model to fit the data updating parameter estimates daily, we still rely on a pre-fixed contact matrix and best-guess proxies for full lockdown and full release. Moreover, our estimate of branching parameters rely on assumptions about the contact rate. Finally, even our data on cases could be incorrect because of incomplete and time-varying testing policy and the high rate of asymptomatic cases.

In order to ensure that our conclusions about adaptive control are robust to these assumptions we engage in two types of robustness checks. First, we address sensitivity to basic parameters such as $\beta_{j,j'}$ and γ by scaling those up and down proportionally by 10% and 25%. Likewise, we vary the proxy dates we choose to proxy for lockdown and release contact rates. Instead of March 23 for release, we try March 15. Instead of April 15 for lockdown, we try April 1.

Second, to address incomplete testing and asymptomatic infections, we scale our count of cases by 2, 4 and 8. Scaling confirmed infections by 2 (8) implies either that we only test one half (12.5%) that we only observe symptomatic cases and they are one half (12.5%) of all cases. We go as high as 8 because data from Maharashtra suggest that 85% of cases may be asymptomatic.

3.5 Caveats on SIR models and role of simulations

Our analysis is not intended to generate quantitative predictions about the number of infections under adaptive control or any other policy. Instead, we view our simulation exercise as *qualitative* analysis that suggests that adaptive control does at least as well as the alternative policies we consider. The amount of the difference between outcomes across policies is unknown.

Our caveat is borne of experience with modeling of the COVID epidemic both by members of the team and others. Our concerns about modeling are the reason we try to simulate the epidemic not just with an SIR model, but also with the synthetic interventions (SI) approach laid out in Agarwal et al. (2020). Both SIR and SI models have limitations. Simulations of policies using our SIR model depend on the adequacy of the parameters. Simulations of policies using the SI model depend on the similarity between the donor pool of countries and the experience in India or Indian states. The assumptions behind each of these simulations are difficult to directly test, especially in the short time frame we have had to study COVID. Our hope is that by looking at two different models, we can reduce our dependence on any one set of assumptions.

4 Simulated policies

Here we provide details on how we define the three policies that we simulate using our SIR model.

4.1 Immediate release after May 3, 2020

We define immediate release as a policy where the effective contact reverts on May 4, 2020, to the rate estimated for March 23. All travel within and between districts is allowed. However, no traffic between states is allowed. (In simulations for Mumbai, the analogous travel restriction is all travel within and between wards are allowed. No travel is allowed in and out of the city.)

4.2 Delayed release after May 31, 2020

We define delayed release as maintaining branching parameters equal to those that prevailed on April 15 from May 4 to May 31. Moreover, there is no travel between states or districts in state simulations, or between ward in the Mumbai simulation. After May 31, states release in the exact manner they do under the immediate release policy. The travel restrictions under immediate release apply with this policy, except they start June 1, 2020.

4.3 Adaptive control starting May 3, 2020

There are four steps to defining the adaptive control policy we simulate.

4.3.1 Gradual change

Adaptive control envisions gradual changes in social distancing and associated socioeconomic activity. We capture this in two ways. The first is that we introduce small steps between full lockdown (branching parameter on April 15) and full release (on March 23). Specifically we partition this range of branching parameter into 4 steps. At the risk of causing confusion with the MHA May 1 order, we call no release the red stop, taking a 1/3 step towards full release the orange stop, a 2/3 step the yellow stop, and a full release the green stop.

We can and will explore different partitions going forward. For example, if one supposes that one can enforce only minor deviations from the extremes, one might try a partition that is 1/4 release and then 3/4 release. Our initial efforts focus on equally spaced steps for simplicity.

The other way we capture gradual changes is that we revisit these steps in 2 week increments. We choose this duration for each step because it roughly maps onto the interval between policy shifts under lockdown in India: March 24 lockdown, April 21 partial release, May 3 end of lockdown, and May 17 next policy pronouncement. We did explore shorter intervals of one week. This introduces more variation in progression through the steps above across districts. However, it does not qualitatively change our conclusion that adaptive control is a constructive policy.

4.3.2 Starting points

Adaptive control requires specifying starting points for states on the day the policy is started (May 4 in our simulations). We initially chose starting positions for districts and wards that are different than the MHA May 1 order. Specifically, if a district's estimated $R_{t,j}$ on May 3

was greater than or equal to one, we started the district at the red stop. If the estimate for $R_{t,j}$ was less than 1, we started the district at the green stop.¹⁵

4.3.3 Response to disease triggers

Adaptive control adjusts social distancing in response to the progress of the disease. Our primary simulations map latest reproductive rate to specific stops defined above. The specific mapping is listed in the table below.

$R_{I,j}$	Stop	Release	Travel to/from other districts
≥ 2	Red	None	none
$[1.5, 2)$	Orange	1/3	none
$[1, 1.5)$	Yellow	2/3	none
< 1.0	Green	Full	Only other Green districts in state

Table 1: Mapping of 2 week ahead forecasted reproductive rate to contact rates and mobility regimes in adaptive release strategy.

This is not the only way to map disease progression to social distancing policy. One alternative is to use forecasted reproductive rate rather than current reproductive rate. (Going forward we plan to use 1-week ahead forecasted reproductive rate as a trigger, though our simulations use current rate.) Another alternative is to map reproductive rates to changes in stops. For instance, one could map from the reproductive rates (≥ 2 , $[1.5, 2)$, $[1, 1.5)$, < 1) to changes in stops (move 2 stops back, move back 1 stop, stay still, move forward). Yet another alternative is to change the trigger from reproductive rate to trajectory of confirmed cases or COVID-associated deaths or to the ratio of COVID-associated cases of serious aggressive respiratory illness (SARI) and available beds in the district. We will explore these alternatives in future simulations.

4.3.4 Travel restrictions

A final feature of adaptive control is that it is a local policy. This means that it regulates not just mobility within a district but also between districts. In our simulations, we adopt a restriction that travel between districts is only allowed if the two districts are both at a green stop. The mobility when two districts are both green is proportional to mobility prior to lockdown, as explained in Section 3.1.

In future simulations we will relax this constraint. Even under the lockdown, movement of essential goods across district boundaries was permitted. Moreover, there were some mass movements, such as of migrants from Delhi to UP. While it is hard to anticipate such super spreader events, we will be doing simulations in which, 3 months after exit from lockdown, there is exogenous introduction of cases.

The importance of inter-district movement or migration as a driver of dynamics varies over time under an SIR model. The key issue is the extent to which migration changes a person in

¹⁵As a robustness check, one could perform simulations where we started districts at the color classifications in the MHA May 1 order. Importantly, we applied red/orange/green stops laid out in Section 4.3.1 rather than the MHA controls because our controls map to specific branching parameters, while MHA policies do not. We have not implemented this check because it is past the date where it is policy-useful to revisit the MHA’s color coding.

district j 's probability of encountering an infected person. This is a function of (a) the number of migrants from another district and (b) the probability that a migrant from another district is infected. For example, if migration is symmetric, this comes down to the relative prevalence of infection inside and outside the district.

At the start of an epidemic, even a few introductions from migrants can trigger an epidemic. This is how COVID was likely introduced to India: travellers from China, Italy and other places brought the infection to different states in late January and early February. Likewise, migrants can reintroduce an epidemic once it is brought substantially under control. This is not just when the reproductive rate is below 1, but when the number of infections is very low. Movement can also matter when a state is close to the trigger for a change in social distancing policy, e.g., when reproductive rate is close to 1, 1.5 or 2 in our implementation of adaptive control.

In policy implementation, we expect that travel will only be allowed across *adjacent* geographic units falling within a maximal-mobility (Green) regime. Our model implementation at time of writing allows any Green-Green travel flows and over-estimates the amount of movement under adaptive control.

5 Results of simulations

We have simulated our SIR model and different policies for a wide array of states. Here we present the results of simulations for four states (Andhra Pradesh, Tamil Nadu, Maharashtra, and Madhya Pradesh) and the city of Mumbai. We have selected these states mainly to ensure that we present states representative of all regions in the country.

In general we find that adaptive control does at least as well, when measured by the number of infections, as immediate or one-month delayed release. However, the degree of benefit from adaptive control varies greatly across states. A main driver of the variation seems to be how bad the situation in a state is as of May 4.

5.1 Tamil Nadu (TN), Andhra Pradesh (AP), Madhya Pradesh (MP)

Figure 5 shows the results of 1000 simulations for Tamil Nadu. Specifically, it presents the total number of infected persons by date for the three policies we examine: full release after May 3 (blue), full release after May 31 (orange) and adaptive control after May 3 (green). This figure shows actual data on infections for the state in black (until April 23) and mean outcome of simulations with colored lines after that date. The three lines overlap until May 4, when the policies diverge in their effects on contact rates. To illustrate “uncertainty”, shaded regions for each policy show the range of outcomes observed under that policy. The edges of the shaded regions give the minimum and maximum level of infection observed across simulations for a specific policy. We will return to our interpretation of this figure in a moment.

Figure 5 shows that TN does better with adaptive control. Whereas delayed release mainly delays the peak of infections by approximately 1 month, adaptive control appears to reduce the peak by several orders of magnitude and leads to a quicker decline in the number of new infections. In the average simulation, the epidemic seems under control by mid-July. This does not mean that all simulations fully release the lockdown by mid-July.¹⁶ Some simulations

¹⁶Although this version of the paper was released in May, we confined our analysis to the simulation and data available as of May 4. The unlikelihood of control in Tamil Nadu by July illustrates the importance of updating




Figure 5: Simulated infection levels for Tamil Nadu under three policy regimes.

generate forecasts of reproductive rate that keep nearly half the districts with a lockdown past that date. The key finding, however, is that, even in that simulation, adaptive control yields better results, in terms of infection control, than either immediate or one-month delayed release.

Figure 6 presents disease progression across simulations in 2 other states: Andhra Pradesh (AP) and Madhya Pradesh (MP). Like Tamil Nadu, AP represents south India; MP is indicative of north India. Andhra Pradesh seems to benefit the most from adaptive control given what we know about their case counts, though the range of outcomes for alternative policies is wide. The peak there is much lower and the infection is resolved much more quickly. Madhya Pradesh benefits just a little less from adaptive control, but that is because the model suggests that the state would have a quicker recovery than Andhra Pradesh.

5.2 Maharashtra and Mumbai

Finally, we present results showing that the adaptive control approach can be applied at either the state or city level. Figure 7 reports on simulations for Maharashtra and Mumbai, the largest city in Maharashtra. In the former, adaptive control is applied at the district level, while in the latter it is applied at a ward level. According to the figures, in both there is improvement due to adaptive control, though that improvement is starker in the case of the state than the city.

We caution again that one ought not take the dates or peaks on these simulations too literally. Projections multiple months out have wide confidence bounds not fully captured by the shading in the figure, which captures within model variation only. We view these models as providing qualitative evidence that adaptive control is a better strategy than lockdown or full release. Moreover, the results are consistent with simulations using other, differently parametric models such as synthetic interventions (Agarwal et al., 2020).

the model parameters with the latest case data.




Figure 6: Impact of adaptive control, immediate and full release after May 3, and immediate and full release after May 31 on number of infections, for Andhra Pradesh and Madhya Pradesh. (Results are based upon 1,000 simulations. Boundaries of shaded areas show outcome under minimum and maximum infections across simulations for each policy.)




Figure 7: Impact of adaptive control, immediate and full release after May 3, and immediate and full release after May 31 on number of infections, for Maharashtra and Mumbai. (Results for Maharashtra and Mumbai are based upon 1,000 simulations. Boundaries of shaded areas show outcome under minimum and maximum infections across simulations for each policy.)

5.3 Robustness checks

Because SIR models are sensitive to parameters, we check to see if adaptive control does at least as well as more extreme policies under a range of parameters. For these checks, our focus is Mumbai. To summarize the relative performance of adaptive control, we will report two outcomes. The first is the average number of infections under alternative policies minus the average infections under adaptive control. Graphically, this is the area under the infection curve for alternative policies versus adaptive control. Second, to account for the fact that performance of adaptive control may vary across simulations, we report the fraction of simulations under which adaptive control does better than the alternatives we consider. The results are presented in Table 2.

The first row presents the benefit of adaptive control in our usual or base simulations. Adaptive control is always better than alternatives: it saves 38.92 and 33.35 million infection days over immediate full release and full release after one month, respectively. Panel A scales up (rows 1-2) and down (rows 3-4) the contact rate β by 10% and 25%. Adaptive control remains strictly better and the gains are similar to the base case. The reason the incremental benefit does not change is that the rate of infection scales up across all three policy regimes.

Panel B conducts a similar exercise with the recovery rate γ . Row 1 scales γ so that the time to recovery increases from 5 to 10 days. Rows 2-4 reduce the recovery time gradually to just 2 days. Again adaptive control is always better and the gains with that policy are similar to the base case, and for the same reason - all policies are affected by the change in rate of infection.

Finally, panel C attempts to account for incomplete testing or the presence of asymptomatic cases by scaling up the number of infected persons. Rows 1-2 increase the number of infections by factors of 2, 4, and 8 respectively. This is akin to allowing that testing only captures 1/2 to 1/8 of cases or that 50% to 87.5% of infected persons are asymptomatic and not tested. Again, adaptive control dominates. However, the benefit of these policies decline. The reason is that if more people are infected, then fewer people remain to benefit from adaptive control going forward.

There are many other robustness checks that can be run. For example, we want to introduce the possibility that governments make mistakes or are unable to enforce policies at a particular stop. This can be modeled as one or two sided errors in the mapping of risk level to policy.¹⁷ Another example is that we can re-introduce the infection, perhaps from foreign travel, after a few months. We will attempt these tests, as well as examine variations to our implementation of adaptive control discussed in Section 4.3, in future drafts.

6 Implementation of adaptive control

Having illustrated the benefits of adaptive control we now discuss the implementation of this policy. We first discuss the different types of triggers that can be used to prompt a change in the amount of social distancing. Then we discuss specific rules that might regulate distancing.

¹⁷We thank Sid Rupani for this recommendation.

Table 2: Robustness of simulations to variation in parameters, as applied to Mumbai.

Contact rate	Days to recovery	Infections	Details	Simulations	How often AC is better (%)	Average improvement (mil. infection days) with AC vs.	
						Immediate release	Delayed release
Estimate from data	5 days (from lit.)	Reported cases	empirical β , $\gamma = 0.2$	15,000	100%	38.92	33.35
A. Vary β							
10% higher than estimate	5 days	Reported cases	$1.1 \times$ emp. β , $\gamma = 0.2$	1,000	100%	38.92	33.33
25% higher than estimate	5 days	Reported cases	$1.25 \times$ emp. β , $\gamma = 0.2$	1,000	100%	38.9	33.31
10% lower than estimate	5 days	Reported cases	$0.9 \times$ emp. β , $\gamma = 0.2$	1,000	100%	38.94	33.35
25% lower than estimate	5 days	Reported cases	$0.75 \times$ emp. β , $\gamma = 0.2$	1,000	100%	38.95	33.37
B. Vary γ							
Estimate from data	10 days	Reported cases	empirical β , $\gamma = 0.1$	1,000	100%	38.98	33.39
Estimate from data	~ 3 days	Reported cases	empirical β , $\gamma = 0.3$	1,000	100%	38.87	33.28
Estimate from data	2.5 days	Reported cases	empirical β , $\gamma = 0.4$	1,000	100%	38.81	33.22
Estimate from data	2 days	Reported cases	empirical β , $\gamma = 0.5$	1,000	100%	38.74	33.16
C. Vary I							
Estimate from data	5 days	Reported cases $\times 2$	true $I = 2 \times$ observed	1,000	100%	36.45	29.25
Estimate from data	5 days	Reported cases $\times 4$	true $I = 4 \times$ observed	1,000	100%	32.84	24.01
Estimate from data	5 days	Reported cases $\times 8$	true $I = 8 \times$ observed	1,000	100%	28.22	18.13

6.1 Possible triggers

Adaptive control takes its name from the fact that it calibrates social distancing to progression of the disease. There are many ways in which this progression can be measured. Here we list three.

The trigger that we primarily recommend is the reproductive rate of the infection $R_{t,j}$. When the reproductive rate is below 1, the epidemic will subside exponentially in time. Whatever the existing level of prevalence and hospitalization, the future daily burden will be lower than before. Our SIR model can use updated parameter estimates to make short, 2 week ahead forecasts of the reproductive rate with reasonable error bands.

A second possible trigger is the trajectory of COVID-related deaths. The use of a trajectory captures the need to control the momentum of the infection. The focus on deaths overcomes the problem that the number of cases is a function of testing policy, which may in practice be inadequate. COVID-related deaths are harder to miss than confirmed cases. That said, deaths tend to be may be under-counted because they can be easily be attributed to common causes such as pneumonia or hearth failure, in the absence of testing. In addition, some individuals may refuse to visit a hospital for fear of getting infected and may die at home. Because of the lockdown, these deaths may not be recorded. Moreover, reporting of both COVID-related hospital and home deaths may be delayed. Just as the SIR model can project reproductive rate, it can also project deaths.

A third possible trigger focuses on the ratio of serious infections (or hospitalizations) to the number of beds. On the theory that the mortality rate is a function of adequate treatment, states may worry not about infections per se but about whether there are more infections requiring hospital beds than there are hospital beds in a district. Thus, a reasonable trigger is the ratio of possible hospitalizations to beds. This too can be projected from the SIR model using an estimate of the ratio of hospitalizations to confirmed cases.

6.2 Policies to regulate contact rates short of a lockdown

There are a wide array of policies that governments have at their disposal to regulate the level of economic activity and social distancing locally. One of these is the lockdown. Another bundle of policies was set forth for red, orange and green districts in the MHA May 1 order. We list these and other policies potential policies in Table 3.

A central challenge for model-aided policy response to COVID is how to map government orders or policies to specific contact rates or disease outcomes. We can estimate the mapping from full lockdown to contact rates because we have a few weeks of information on this policy. However, for other policies which have not yet been tried, it is a challenge. Relatedly, if there is not a lot of variation in the policies tried in different locations or at different times, one can only estimate the impact of a portfolio of policies on outcomes. One cannot estimate the impact of specific narrow policies.

Moreover, it is difficult to map policies to COVID cases, hospitalizations and deaths. While there may be a quick response of mobility to policies, cases respond more slowly because of the latency between contacts and confirmation of infections. This delay would not be a problem if we knew exactly the gap between mobility and cases: one could just look at appropriately lagged outcomes. The problem is that we do not know that latency and it may even vary across places.

Finally, even if one could overcome the problem of variation in policy and delayed impacts, one has the problem of separating correlation and causation. Policy choice is endogenous. Indeed, the MHA May 1 order and adaptive control, the approach we advocate, imposes greater restrictions in places where epidemic risk is higher. If we know the policy rule that governs restrictions, then we can, to some extent, unravel causation from disease risk to policy from policy to disease risk. However, outside these regimes, endogeneity is a problem. Moreover, even in places that formally adopt these policies, there may be irregular variation in enforcement depending on disease risk.

Going forward we will attempt to address these problems in a number of ways. First, we have begun to gather variation in policies across districts and states. While it is tempting to treat the lockdown as a monolithic policy, in reality there is variation over time and across locations in India. Moreover, while the MHA May 1 order seems to set forth a small number of policy portfolios, it masks heterogeneity in policy. For one thing, the order establishes a floor on distancing and allows states to restrict activity more than the order requires. For another thing, there is variation in enforcement across locations. If true and that enforcement level could be measured, we can obtain greater measured variation in measured distancing.

Second, because adaptive control recommends a level of contact rates and that maps nicely to mobility, we can – for our purposes – focus on the impact of policy on mobility rather than cases. The reason we are able to do this is that we use the continually re-estimated SIR model to impose a parametric structure on the mapping from mobility or contacts to cases.

Finally, when data is sparse, we recommend the use of synthetic control methods (SCM) to estimate the causal effect of policies on outcomes. With SCM, we match countries based on disease profile and then use changes in policy in treatment locations to measure the impact of those policies. This method can even be applied to experience across countries, as illustrated in Agarwal et al. (2020). As time goes on and we have more observations on policies, especially local observations, other statistical learning methods can be applied.

Table 3: Policies to regulate social distancing

Policies	Regulation	
MHA May 1 order	<p>All locations:</p> <ul style="list-style-type: none"> - No air travel - No passenger travel in trains - No interstate buses - No metro rail service - Schools closed - No hospitality services - Theatres, malls, parks, etc. closed - No social/political gatherings - Religious places closed - Longer store hours - Shelter at home for children <10 - Shelter at home for elderly >65 - Curfew from 7pm to 7am 	
	<p>Containment zones:</p> <ul style="list-style-type: none"> - Contact tracing and quarantine - Strict perimeter 	
	<p>Red zones:</p> <ul style="list-style-type: none"> - No rickshaws, taxis, intra-state buses - No salons - Only 2 passengers in cars - Urban mfg only in special zones - Permit essential good everywhere - Permit rural mfg - Only in situ construction - Only e-commerce for essential goods - Limit private office occupancy 	
	<p>Orange zone:</p> <ul style="list-style-type: none"> - No intra-state buses - Taxis with only 2 passengers - Only 2 passengers in cars across district lines 	
	<p>Green zones</p> <ul style="list-style-type: none"> - Buses limited to 50% capacity 	
	Other policies	<hr/> <ul style="list-style-type: none"> - Require masks in public - Maximum occupancy in stores - Limits on labor migration - Expand store hours
		<ul style="list-style-type: none"> - Set more, appropriate store hours <hr/>

References

- Agarwal, A., Alomar, A., Sarker, A., Shah, D., Shen, D., and Yang, C. (2020). Two burning questions on covid-19: Did shutting down the economy help? can we (partially) reopen the economy without risking the second wave?
- Bettencourt, L. M. and Ribeiro, R. M. (2008). Real time bayesian estimation of the epidemic potential of emerging infectious diseases. *PLoS One*, 3(5).
- Chatterjee, K., Chatterjee, K., Kumar, A., and Shankar, S. (2020). Healthcare impact of covid-19 epidemic in india: A stochastic mathematical model. *Medical Journal Armed Forces India*.
- Cori, A., Ferguson, N. M., Fraser, C., and Cauchemez, S. (2013). A new framework and software to estimate time-varying reproduction numbers during epidemics. *American journal of epidemiology*, 178(9):1505–1512.
- Figuroa, N. B., Kaiser, D. I., Shah, A. J., and Shah, J. A. (2020). Sensitivity of predictions from sir and seir epidemic models to parameter uncertainty. Report, MIT.
- Ray, D., Salvatore, M., Bhattacharyya, R., Wang, L., Mohammed, S., Purkayastha, S., Halder, A., Rix, A., Barker, D., Kleinsasser, M., Zhou, Y., Song, P., Bose, D., Banerjee, M., Baladandayuthapani, V., Ghosh, P., and Mukherjee, B. (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. *medRxiv*, page 2020.04.15.20067256.
- Singh, R. and Adhikari, R. (2020). Age-structured impact of social distancing on the covid-19 epidemic in india. *arXiv preprint arXiv:2003.12055*.
- Wallinga, J. and Teunis, P. (2004). Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *American Journal of epidemiology*, 160(6):509–516.