



Disease-dependent interaction policies to support health and economic outcomes during the COVID-19 epidemic

Guanlin Li, Shashwat Shivam, Michael Hochberg, Yorai Wardi, Joshua Weitz

► To cite this version:

Guanlin Li, Shashwat Shivam, Michael Hochberg, Yorai Wardi, Joshua Weitz. Disease-dependent interaction policies to support health and economic outcomes during the COVID-19 epidemic. 2020. hal-03018339

HAL Id: hal-03018339

<https://hal.archives-ouvertes.fr/hal-03018339>

Preprint submitted on 22 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Disease-dependent interaction policies to support health and economic outcomes during the COVID-19 epidemic

Guanlin Li^{1,2,*}, Shashwat Shivam^{3,*}, Michael E. Hochberg^{4,5,†}, Yorai Wardi^{3,†} & Joshua S. Weitz^{6,2,†}

¹ Interdisciplinary Graduate Program in Quantitative Biosciences, Georgia Institute of Technology, Atlanta, GA, USA.

² School of Physics, Georgia Institute of Technology, Atlanta, GA, USA.

³ School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA.

⁴ ISEM, Université de Montpellier, CNRS, IRD, EPHE, Montpellier, France.

⁵ Santa Fe Institute, Santa Fe, NM 87501, USA.

⁶ School of Biological Sciences, Georgia Institute of Technology, Atlanta, GA, USA.

† Correspondence to: Michael E Hochberg (michael.hochberg@umontpellier.fr), Yorai Wardi (ywardi@ece.gatech.edu) and Joshua S. Weitz (jsweitz@gatech.edu);

*Authors contributed equally to the work.

Abstract

Lockdowns and stay-at-home orders have partially mitigated the spread of Covid-19. However, the indiscriminate nature of mitigation — applying to all individuals irrespective of disease status — has come with substantial socioeconomic costs. Here, we explore how to leverage the increasing reliability and scale of both molecular and serological tests to balance transmission risks with economic costs involved in responding to Covid-19 epidemics. First, we introduce an optimal control approach that identifies personalized interaction rates according to an individual's test status; such that infected individuals isolate, recovered individuals can elevate their interactions, and activity of susceptible individuals varies over time. Critically, the extent to which susceptible individuals can return to work depends strongly on isolation efficiency. As we show, optimal control policies can yield mitigation policies with similar infection rates to total shutdown but lower socioeconomic costs. However, optimal control policies can be fragile given mis-specification of parameters or mis-estimation of the current disease state. Hence, we leverage insights from the optimal control solutions and propose a feedback control approach based on monitoring of the epidemic state. We utilize genetic algorithms to identify a 'switching' policy such that susceptible individuals (both PCR and serological test negative) return to work after lockdowns insofar as recovered fraction is much higher than the circulating infected prevalence. This feedback control policy exhibits similar performance results to optimal control, but with greater robustness to uncertainty. Overall, our analysis shows that test-driven

improvements in isolation efficiency of infectious individuals can inform disease-dependent interaction policies that mitigate transmission while enhancing the return of individuals to pre-pandemic economic activity.

Introduction

As of 13 August 2020, more than 20,439,000 cases of coronavirus disease 2019 (COVID-19) have been reported worldwide with more than 744,000 deaths globally ¹. Starting at the reported origin of the pandemic in Wuhan, China, control measures have been implemented in most countries where outbreaks have occurred ^{2,3,4,5,6,7}. Multiple public health strategies are being deployed to slow outbreaks, and although recommendations always include social distancing and isolation of confirmed cases, the full spectrum of measures and levels of adherence differ from country to country, making assessments of strategy efficacy difficult (see ², controversy surrounding ⁸)

The non-pharmaceutical control strategies for COVID-19 largely follow those employed in previous viral epidemics, including SARS, Ebola and MERS. Initial strategies can be broadly grouped into mitigation and suppression, where the former attempts to preserve essential health care services and contain morbidity and mortality, whereas the latter imposes more severe, emergency restrictions to prevent health care system collapse and provide conditions for easing-off towards less intense mitigation strategies ⁹. Both mitigation and suppression approaches carry considerable social and economic costs, meaning that policymakers and the public at large only adopt them for short time periods ¹⁰. A problem is that control measures have often been applied irrespective of an individual's disease status (and/or likely infection risk severity) and are driven, in part, by the absence of information-driven alternatives.

Hence, distinct from lockdowns, there is an increasing interest in implementing population-wide prevention methods that decrease transmission risk while enabling economic re-engagement. Examples of such measures include mask-wearing ^{11,12}, contact-free interactions ¹³, and restructuring of physical spaces ¹⁴. The use of masks, in particular, has been shown to be effective at reducing respiratory transmission of SARS-CoV-2, particularly when individuals in a potentially infectious interaction are wearing a mask ^{15,16}. These population-wide measures still carry uncertainty since individuals are expected to behave uniformly irrespective of their disease status. As the scale of COVID-19 testing has increased, jurisdictions may also have an opportunity to consider implementing tactical mitigation strategies informed by testing.

Testing for infected status can, in theory, be used to initiate targeted isolation, identification and tracing of contacts, quarantining of contacts, and then selected testing of contacts ¹⁷. If done rapidly and at scale, this kind of targeted PCR-based testing can provide early detection of cases

and help break new chains of transmission^{18,19}. The effectiveness of contact-tracing based control strategies hinges on the accurate identification and isolation of exposed and infectious cases. Slow return of test results (primarily) and false negatives (as a secondary factor) limit the effectiveness of test-based control policy²⁰. A complementary tactic is the strategic deployment of recovered individuals identified by serological tests in infection hot-spots to effectively dilute transmission events^{21,22}. Similar to the identification of infected individuals through PCR tests, immunity to SARS-CoV-2 is assessed through serological testing for protective antibodies^{23,24}. Population-wide interventions, testing efforts, the selective confinement or deployment of people contingent on their infection status, and inaccuracies and limited adherence to policies, combine to create a challenging landscape for the persistent control of outbreaks²⁵.

Non-pharmaceutical COVID-19 control until effective vaccines become available will necessarily involve periods of reduced social and economic activity; i.e., ‘business, but not as usual’. Control efforts are already generating hardship and could in the longer-term result in social unrest and increased mortality^{26,27,28}. Here we confront a joint problem: how to identify policies that aim to reduce fatalities arising from COVID-19 while also enabling economic engagement. Such a joint goal necessarily has a tension, can lead to dichotomous ‘extremal’ policies, and requires confrontation through careful modeling. We do so in two stages. First, we use optimal control to assess both health and economic outcomes in an SEIR disease model framework. There is a substantial and growing literature on optimal control for COVID-19, the bulk of which focuses on non-personalized release policies or policies that target age- or risk-stratified groups^{2,29,30,6,31,9}. Here, we identify optimal control policies to modulate interaction rates based on disease status identified by PCR and serological testing - unifying prior efforts centered on isolation and shield immunity. We find that intermediate policy outcomes can do nearly as well as strict public health scenarios, without incurring the severe costs as suppression-centered policies. However, optimal controls can be fragile, when applied in practice given that they rely on time- rather than state-based interventions; the consequence of mistiming interventions can be severe³⁰. Hence, guided by the optimal control analysis, we identify state-dependent policies similar to feedback control that provide actionable guidance for individual behavior. As we show, using population-wide PCR and serological testing as the basis for mitigation (rather than surveillance) has the potential to yield dual benefits if carried out at population scales: reducing COVID-19 transmission while enabling more individuals to return to work sooner and with fewer restrictions than would otherwise be possible.

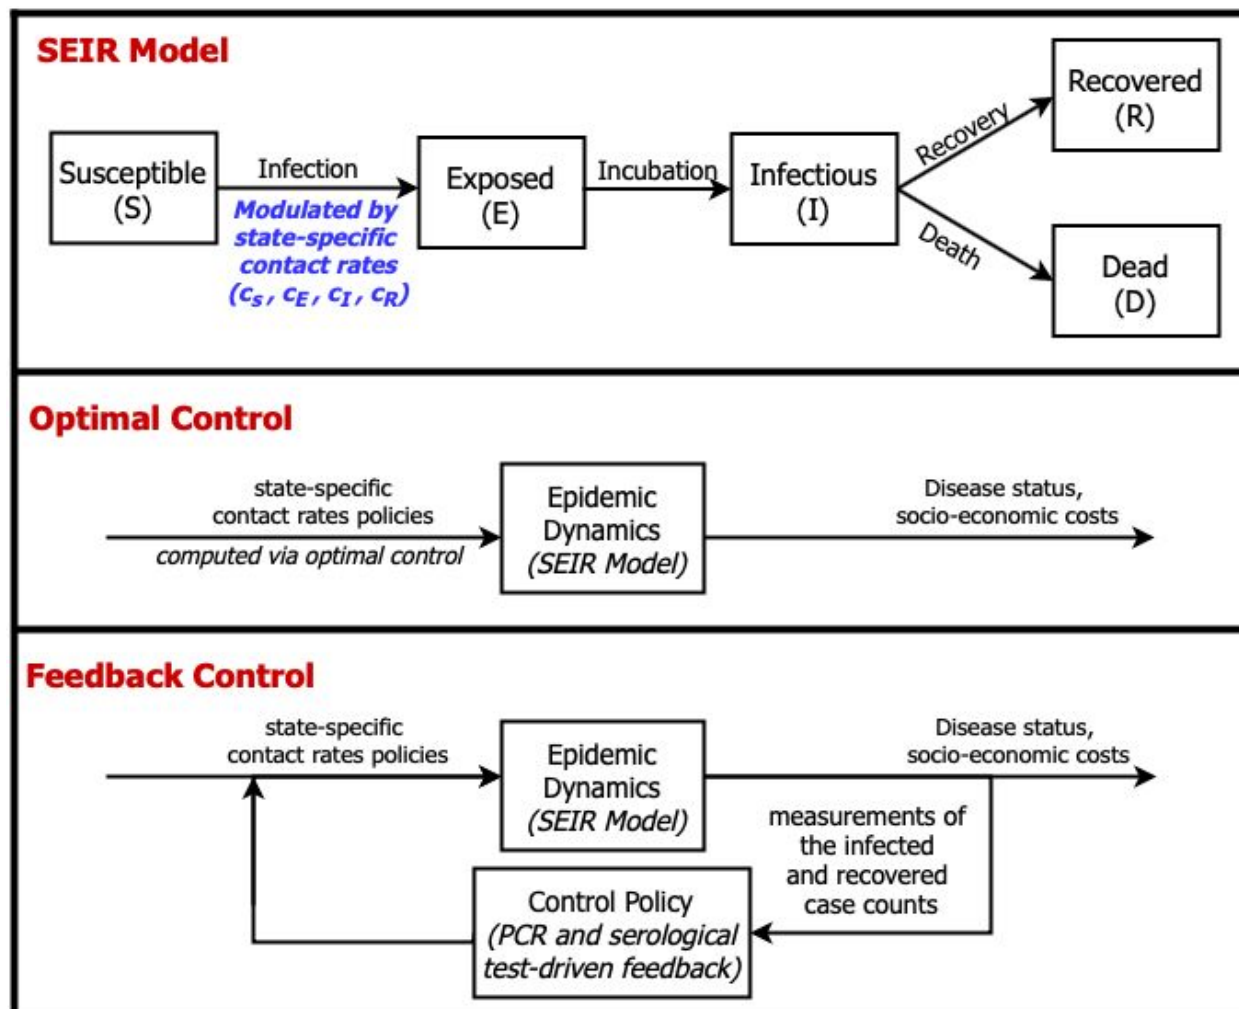


Figure 1: Epidemic dynamics with optimal and feedback control of disease-status driven contact rates. (Top) SEIR model schematic in which the force of infection is modulated by state-specific contact rates, see text and SI for details. (Middle) Diagram of optimal control approach: contact rates are pre-specified given model structure and estimate of parameters and current conditions. (Bottom) Diagram of feedback control approach: contact rates are updated in real-time based on measurements of the infected and recovered case counts via testing surveillance.

Results

Optimal control framework for state-dependent contact rates policies that balance public health and socioeconomic costs

We develop an optimal control framework to identify policies that address the tension between decreasing contacts (that reduce new infections) with increasing contacts (that are linked to socio-economic benefits). We represent the epidemic using a

Susceptible-Exposed-Infectious-Recovered (SEIR) nonlinear dynamic model (see Supplementary Information/Methods for complete details; see Figure 1). In doing so, the force of infection is influenced by state-specific contact rates c_S , c_E , c_I and c_R for susceptible, exposed, infectious and recovered individuals, respectively – these different levels form the basis for a control policy that directs individuals to interact at different levels depending on their test status.

In the optimal control framework, a set of state-specific contact rates are identified that minimize the appropriately weighted sum of what we term ‘public health’ and ‘socioeconomic’ costs. Public health costs are quantified both by average infected levels and cumulative deaths. Socioeconomic costs are quantified in terms of reductions in the total rate of interactions and by shifts in state-specific contact rates. The optimal control ‘solution’ is then a time-dependent set of disease-specific rates which are both shaped by and shape the epidemic itself (see SI/Methods for details on the gradient projection algorithm used to identify the solution). Note that we constrain the contact rate of exposed individuals to be equal to that of susceptible individuals given the challenges of timely identification of exposed individuals who are not yet infectious (and presumably have insufficient viral titer to be identified using screening tests; an issue we return to in the Discussion).

Figure 2 shows the results of comparing a baseline outbreak (i.e., neglecting public health costs, given weighting parameter $\xi=0$) to a full lockdown scenario (i.e., neglecting socioeconomic costs with 75% isolation for all, $\xi \gg 1$) and a balanced scenario with optimized contact-rates (i.e., corresponding to $\xi = 1$). As shown in Figure 2, in the baseline scenario, the disease spreads through the population leading to 94% cumulative infection (as expected given strength-size relationships for $\mathcal{R}_0 = 3$). In contrast, a full lockdown scenario with 75% reduction in contact rates of all individuals after 60 days leads to a total outbreak size of 4% of the population. The optimal control solution in the balanced case $\xi = 1$ reveals a potential route to jointly address public health and socioeconomic cost. From the perspective of public health, the optimal control solution leads to 25% cumulative infections. In addition, the socioeconomic costs in the optimal control case are higher in the short-term but approach that of the baseline scenario in the long-term. Indeed, the effective reproduction number identified via an optimal control framework in the balanced scenario gradually reduces to sub-critical levels (close to an effective reproduction number, $\mathcal{R}_{eff} = 0.75$) while gradually relaxing controls over time. The optimal control solutions are shown in Figure 2A. The optimal control solutions differ based on disease status, recovered individuals elevate their interactions, infectious individuals isolate, and susceptible individuals lock-down before gradually returning to pre-lockdown levels.

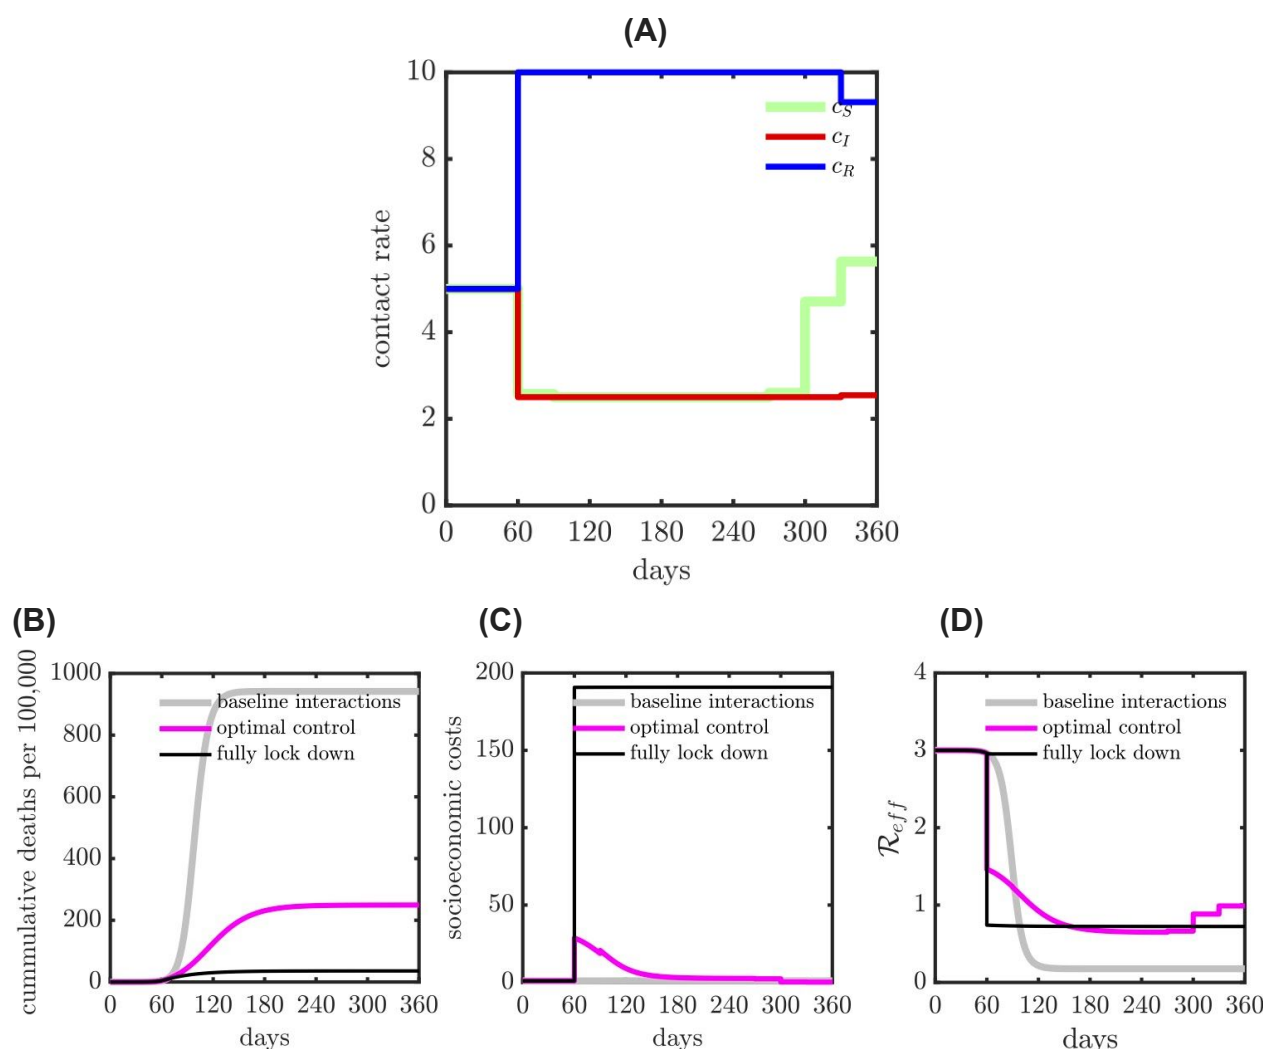


Figure 2: Comparison of health and economic outcomes of COVID-19 given various interventions: baseline interactions (i.e., no intervention); optimal contact rate intervention (balance both health and economic outcomes) and fully lock down intervention (applied to all the subpopulations) with 75% isolation efficiency. (A) The optimal contact rate with 50% isolation effectiveness and shield immunity level 2. (B) Cumulative deaths (health outcome) during the epidemic. (C) Socio-economic costs (economic outcome) during the epidemic. (D) Measure of effective reproduction number (R_{eff}) for different interventions during the epidemic.

Personalized, test-based optimal control policies and their impact on public health and socioeconomic outcomes.

In order to explore the mechanisms identified by the optimal control framework, we systematically modulate the effectiveness of isolation and evaluate its effect on the state-dependent optimal contact rates and disease dynamics. In practice, isolation effectiveness is influenced by availability, accuracy, and speed of testing as well as fundamental limitations on

an individual's ability to isolate (which can vary with socioeconomic and other factors). Figures 3A-C evaluate low, medium, and high efficiency of isolation spanning 25%, 50%, and 75% reduction in the contact rates of infected individuals, respectively. As is evident, the optimal control solutions for state-dependent contact rates vary significantly with isolation effectiveness; suggesting that COVID-19 response policies that can vary with disease status may open up new possibilities to balance public health and socioeconomic outcomes.

First, in the low (25%) or medium (50%) effectiveness cases, susceptible, exposed, and infectious individuals adopt the maximal level of isolation. Inefficient isolation of infectious individuals elevates risks of new transmission that are not outweighed by socioeconomic benefits. Notably, the optimal control solution includes an *elevated* level of interaction by recovered individuals. This finding recapitulates 'shield immunity' ^{21,22}, insofar as recovered individuals are protected from re-infection over the course of the intervention. The elevated contacts of recovered individuals have multiple effects: both diluting interactions by susceptibles (and reducing transmission risk) and by increasing socioeconomic activity. In contrast, for sufficiently high levels of isolation efficiency (75%), the optimal control solutions suggest there is no need for a general lockdown. Instead, the combination of infected case isolation and shielding by the subpopulation of recovered individuals is sufficient to rapidly reduce and contain \mathcal{R}_{eff} below 1, leading to a decreasing number of new infections. We note that irrespective of isolation effectiveness, balancing public health and economic outcomes drives \mathcal{R}_{eff} below 1, but not necessarily to 0 (albeit, given the constraints imposed by lockdown efficiency, such an extreme reduction may not even be possible), and eventually increased immunity permits an easing-off in restrictions yielding an increase in \mathcal{R}_{eff} ³².




Figure 3: SEIR dynamics with contact rate interventions for various isolation efficiencies, (A) 25% isolation efficiency; (B) 50% isolation efficiency and (C) 75% isolation efficiency. The relative importance (ξ) is 1 for all the cases (A), (B) and (C). The contact rate interventions start at 60 days, people follow baseline (or normal) interactions before that. For all the isolation efficiency scenarios (three rows), the left panel shows the population dynamics given the optimal contact rate shown in the middle panel. The gray curve in the middle panel represents the measure of corresponding effective reproduction number (\mathcal{R}_{eff}). The right panel shows the corresponding socio-economic costs. See SI/Methods for additional scenarios.

Sensitivity of optimal control approach to mis-timed implementation of policies

Despite its potential to balance public health and socioeconomic costs, a central drawback of optimal control solutions is the potential exponential growth of errors. Given the fact the COVID-19 dynamics are only partially observed (with significant uncertainty in the actual state), application of a policy requires an estimate of the time since epidemic initiation, what we term ‘epidemic age’. In order to evaluate the sensitivity of optimal control policies due to mis-timing, we first computed the optimal control policy for a system one month after an outbreak. However, instead of implementing the policy matched to the actual epidemic age, we enforce the optimal control policy 30 days later, i.e. at the end of 60 days after the start of the outbreak. Figure S7 shows the difference in the mis-timed control policy vs. the optimal control policy; as is evident the mistimed policy relaxes stringent lockdown when the optimal policy continues to lock-down. As a consequence the total deaths are far higher for 25% and 50% isolation efficiency (see Table 1). The mistimed policy, in effect, biases the system towards minimizing socioeconomic rather than public health costs. This significant difference in performance metrics demonstrates the potential shortcomings of implementing a policy based on optimal control. However we note that with a stringent isolation efficiency, delays are less problematic. The reason is that with efficient infected case isolation, both the mis-timed and optimal control policy could enable nearly all individuals to work, given that \mathcal{R}_{eff} is held below 1 by infection isolation on its own.

Despite its fragility, we identify common features of the optimal control policy given variation in the effectiveness of infectious case isolation. First, the optimal control policy minimizes infected contact rates. The optimal control solutions also robustly identify an immune shielding strategy for recovered individuals, i.e., such that recovered individuals elevate their interactions to the maximum possible relative to baseline. Importantly, differences in the optimal control policy are primarily centered on identifying a switch point in contact rate level for the susceptible population. From Figures S4, S5 and S6, we observe the switch point as a function of time, showing that irrespective of isolated case effectiveness and shield immunity constraints, the increase in susceptible contact rates happens later in the lockdown period. Switchover points correspond to times when the infection prevalence is relatively low compared to the recovered population. This observation provides the basis for a feedback, rather than optimal, control policy.

Optimal Control Approach	Without Delay		With Delay	
Efficiency of Isolation (shielding = 2)	Total Deaths (per 100,000 individuals)	Working Fraction	Total Deaths (per 100,000 individuals)	Working Fraction
25%	600	90.95%	720	94.96%
50%	250	67.13%	510	83.65%
75%	40	99.92%	40	99.95%

Table 1: Comparison between optimal control approach for contact policy with and without delay. The comparisons are made for isolation efficiencies of 25%, 50% and 75%, with performance metrics of total deaths (per 100,000 individuals) and working fraction. The total death is significantly higher for the system with delay when the isolation efficiency is not 75%, suggesting poor robustness to delay.

Feedback-control policy for balancing public health and socioeconomic costs

We propose the use of a feedback control policy adapted from emergent features of the optimal control policy solutions: (i) infectious individuals isolate as far as is possible; (ii) recovered individuals increase their activities as much as possible, i.e., akin to shield immunity. Hence, we set out to identify a system-dependent change in the contact rate of susceptible individuals, separating lockdowns vs. return-to-work. In practice, we identify a critical curve in I - R plane (i.e., infected-recovered cases plane) via a genetic algorithm, such that the recommended behavior of susceptible individuals is dictated by surveillance-based estimates of infectious and recovered individuals (see Supplementary Information for more details).

Figure 4 summarizes the results of the feedback control policy. From a policy perspective, for both 25% and 50% isolation effectiveness the feedback control policy identifies a switch between lock-down and return-to-work when there are significantly more recovered individuals than infectious individuals (approximately 100x as many). Hence, circulating case levels should be relatively low relative to recovered individuals before shifting from lockdowns to re-openings. We also examined the same scenarios while assuming recovered individuals return to work but without the beneficial effects of shield immunity. As is apparent, the policies are more stringent, with the critical relationship shifting to requiring that there are more than 350 times as many recovered individuals as infected individuals.

Critically, the performance of the test-driven feedback policy is nearly identical for the performance metrics with or without mis-timing (see Table 2). This finding implies that state-based approaches will be less likely to have exponentially mis-timed applications, and

reinforces the need for population-scale testing for both active infections (via PCR) and prior infections (via serological and/or antigenic based testing that evaluates protective immune responses). We note that a simple policy with only two states - ‘lockdown’ and ‘open’, respectively corresponding to minimum and baseline contact rates for the susceptible cases, would be easier to implement than one with continuous ‘phases’ or state changes. In Figure S9, we document the generalizability of results given variation in infected case isolation and the level of shield immunity.

Feedback Control Approach	Without Delay		With Delay	
Efficiency of Isolation (shielding = 2)	Total Deaths (per 100,000 individuals)	Working Fraction	Total Deaths (per 100,000 individuals)	Working Fraction
25%	620	89.98%	620	89.94%
50%	250	62.40%	250	62.12%
75%	30	99.90%	30	99.90%

Table 2: Comparison between feedback control approach for contact policy with and without delay. The comparisons are made for isolation efficiencies of 25%, 50% and 75%, with performance metrics of total deaths and working fraction. Both performance metrics are nearly identical for all efficiencies, suggesting no significant effect of delay on the system.

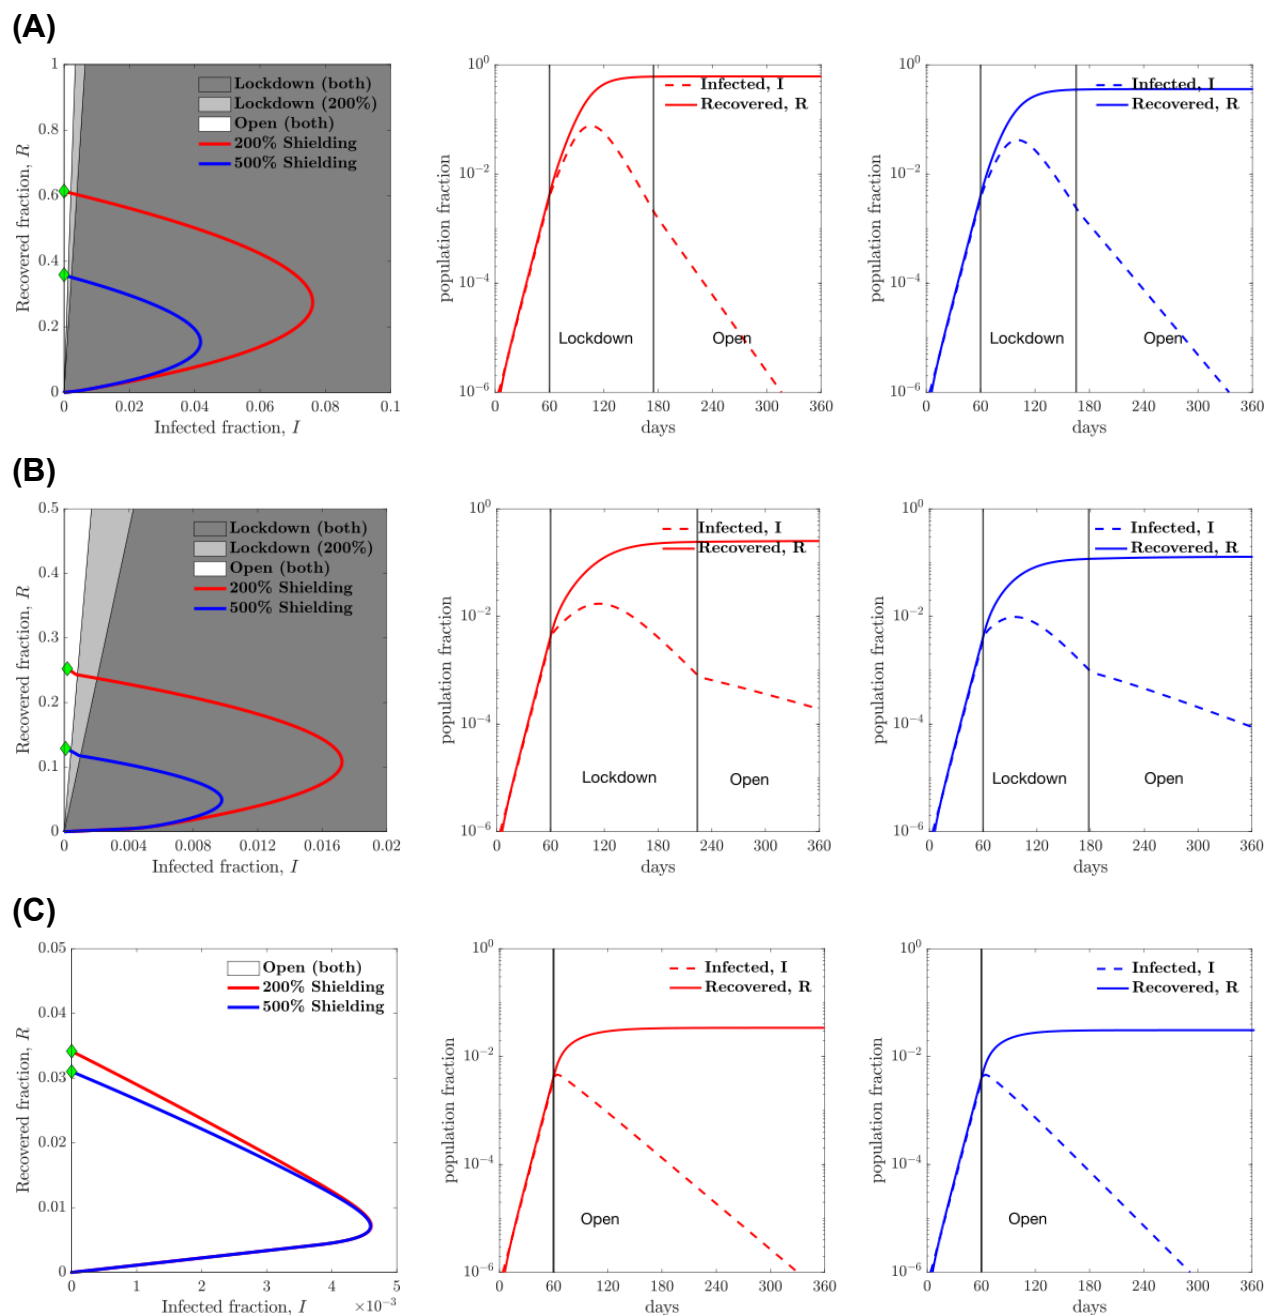


Figure 4: Heuristic state feedback intervention policies varying with isolation efficiency: (A) 25% isolation efficiency; (B) 50% isolation efficiency and (C) 75% isolation efficiency. The shielding levels considered here are 2 and 5 times the base contact rate (i.e., 200% shielding and 500% shielding respectively). An optimal line divides the plane into two regions which determines the optimal contact rate for the susceptible population for the current infected and recovered cases. The optimal policy in the dark grey region is lockdown for both shielding levels while the system with a higher level of shielding (level of 5) is open in the light grey region. The phase plots (red and blue curves) show the evolution of the infected and recovered case fractions

over the period of 360 days, while applying the control strategy described above for shielding levels of 2 and 5 respectively. See SI/Methods for a larger set of plots with increments of 5% in isolation efficiency and for shielding levels of 2, 3, 4 and 5. The second and third plots show the infected and recovered cases for shielding levels of 2 and 5 respectively as functions of time. The vertical lines mark the time instances at which the policy is implemented with lockdown imposed and also the time at which it is lifted. For the case of isolation efficiency of 75%, no lockdown is needed at all for the susceptible population.

Concluding remarks

We have developed a linked series of optimal- and feedback-control analyses to evaluate the effectiveness (and benefits) of modifying contact rates for managing the COVID19 pandemic from both health and economic perspectives. Throughout, our central goal was to optimize the interactions between individuals (in different compartments) so as to achieve a defined balance between public health and economic outcomes. By explicitly incorporating contact rates as the control variables in a SEIR model, we were able to identify optimal control policies that could, in theory, significantly reduce expected infections (and fatalities) while reducing the negative socioeconomic costs of sustained lockdowns. Optimal control policies are unlikely to be applied in practice, given the potential exponential mis-specification of policies over time. Hence, we leveraged insights from the optimal control solutions to guide a feedback control approach that performs nearly as well as the optimal control approach with significant improvements in robustness given uncertainty in estimating the epidemic state. Collectively, our control policies indicate that infected individuals should be isolated (as effectively as possible), recovered individuals should be encouraged to return to work (given benefits accrued via shield immunity), while the release of other individuals from lock-down should be guided by the epidemic state. The transition from lock-down to return-to-work occurs when circulating case loads are far lower than recovered case counts; with the scope of the epidemic sharply controlled by infection case isolation. A combination of policies, e.g., mask-wearing, physical distancing, will help to reduce transmission risk for individuals who do return to work.

We recognize that the SEIR framework is intentionally simplified. The epidemic model does not account for process or observational noise, analytic test features, heterogeneity, stratified risk, asymptomatic cases, and detailed elaboration of severe cases. By reducing the model complexity we have tried to shed light on the general problem of balancing public health with socioeconomic outcomes. In doing so, we have highlighted a middle ground between dichotomous outcomes that focus on public health or socioeconomic costs solely. Central to our work is the presumption that testing is sufficient to estimate both the circulating case counts and the number of recovered individuals. In doing so, serological testing is key. Given under-testing, the calibration of reported cases requires either model-based inference or serological test

comparisons to estimate the ascertainment bias³³. Collectively, we remain closer to the beginning than the end of the Covid-19 pandemic. As we have shown, accelerating the slow-down of transmission while restoring economic activity may be enabled by both personalized, test-driven, policies.

Data availability: All simulation and figure codes used in the development of this manuscript are available at https://github.com/WeitzGroup/COVID_Control_2020

Acknowledgements: We thank Weitz Group team members for comments and feedback. This work was supported by grants from the Army Research Office (W911NF1910384), National Institutes of Health (1R01AI46592-01), and the National Science Foundation (1806606 and 2032082).

Reference

1. World Health Organization Coronavirus Disease 2019 (COVID-19) Situation Report 206.
WHO
<https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200813-covid-19-sitrep-206.pdf> (2020).
2. Flaxman, S. *et al.* Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* (2020) doi:10.1038/s41586-020-2405-7.
3. Chinazzi, M. *et al.* The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* **368**, 395–400 (2020).
4. Wells, C. R. *et al.* Impact of international travel and border control measures on the global spread of the novel 2019 coronavirus outbreak. *Proceedings of the National Academy of Sciences* vol. 117 7504–7509 (2020).
5. Cowling, B. J. *et al.* Impact assessment of non-pharmaceutical interventions against coronavirus disease 2019 and influenza in Hong Kong: an observational study. *The Lancet Public Health* vol. 5 e279–e288 (2020).

6. Ferguson, N. *et al.* Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. (2020) doi:10.25561/77482.
7. Kraemer, M. U. G. *et al.* The effect of human mobility and control measures on the COVID-19 epidemic in China. *medRxiv* (2020) doi:10.1101/2020.03.02.20026708.
8. Zhang, R., Li, Y., Zhang, A. L., Wang, Y. & Molina, M. J. Identifying airborne transmission as the dominant route for the spread of COVID-19. *Proc. Natl. Acad. Sci. U. S. A.* **117**, 14857–14863 (2020).
9. Walker, P. G. T. *et al.* The impact of COVID-19 and strategies for mitigation and suppression in low- and middle-income countries. *Science* **369**, 413–422 (2020).
10. The territorial impact of COVID-19: Managing the crisis across levels of government. OECD Policy Responses to Coronavirus (COVID-19). *OECD 2020a*
https://read.oecd-ilibrary.org/view/?ref=128_128287-5agkkojaaa&title=The-territorial-impact-of-covid-19-managing-the-crisis-across-levels-of-government.
11. Greenhalgh, T., Schmid, M. B., Czypionka, T., Bassler, D. & Gruer, L. Face masks for the public during the covid-19 crisis. *BMJ* **369**, m1435 (2020).
12. Chu, D. K. *et al.* Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: a systematic review and meta-analysis. *Lancet* **395**, 1973–1987 (2020).
13. Brynjolfsson, E. *et al.* COVID-19 and Remote Work: An Early Look at US Data. (2020) doi:10.3386/w27344.
14. Bodkin, C. *et al.* Pandemic Planning in Homeless Shelters: A pilot study of a COVID-19 testing and support program to mitigate the risk of COVID-19 outbreaks in congregate

- settings. *Clin. Infect. Dis.* (2020) doi:10.1093/cid/ciaa743.
15. Feng, S. *et al.* Rational use of face masks in the COVID-19 pandemic. *Lancet Respir Med* **8**, 434–436 (2020).
16. Cheng, V. C.-C. *et al.* The role of community-wide wearing of face mask for control of coronavirus disease 2019 (COVID-19) epidemic due to SARS-CoV-2. *J. Infect.* **81**, 107–114 (2020).
17. Ferretti, L. *et al.* Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* **368**, (2020).
18. Yong, S. E. F. *et al.* Connecting clusters of COVID-19: an epidemiological and serological investigation. *Lancet Infect. Dis.* **20**, 809–815 (2020).
19. He, X. *et al.* Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine* vol. 26 672–675 (2020).
20. Larremore, D. B. *et al.* Test sensitivity is secondary to frequency and turnaround time for COVID-19 surveillance. *medRxiv* (2020) doi:10.1101/2020.06.22.20136309.
21. Weitz, J. S. *et al.* Modeling shield immunity to reduce COVID-19 epidemic spread. *Nat. Med.* **26**, 849–854 (2020).
22. Kraay, A. N. M., Nelson, K., Zhao, C., Weitz, J. S. & Lopman, B. A. Modeling serological testing to inform relaxation of social distancing for COVID-19 control. *medRxiv* (2020) doi:10.1101/2020.04.24.20078576.
23. Krammer, F. & Simon, V. Serology assays to manage COVID-19. *Science* **368**, 1060–1061 (2020).
24. Amanat, F. *et al.* A serological assay to detect SARS-CoV-2 seroconversion in humans.

- medRxiv* (2020) doi:10.1101/2020.03.17.20037713.
25. Fauci, A. S., Lane, H. C. & Redfield, R. R. Covid-19 - Navigating the Uncharted. *The New England journal of medicine* vol. 382 1268–1269 (2020).
 26. Flattening the unemployment curve policies to support workers' income and promote a speedy labour market recovery. OECD Policy Responses to Coronavirus (COVID-19). *OECD 2020b*
https://read.oecd-ilibrary.org/view/?ref=134_134515-cww8xuoucy&title=Flattening-the-unemployment-curve-Policies-to-support-workers-income-and-promote-a-speedy-labour-market-recovery.
 27. Douglas, M., Katikireddi, S. V., Taulbut, M., McKee, M. & McCartney, G. Mitigating the wider health effects of covid-19 pandemic response. *BMJ* **369**, m1557 (2020).
 28. Hsiang, S. *et al.* The effect of large-scale anti-contagion policies on the COVID-19 pandemic. *Nature* **584**, 262–267 (2020).
 29. Bonnans, J. F. & Gianatti, J. Optimal control techniques based on infection age for the study of the COVID-19 epidemic. (2020).
 30. Morris, D. H., Rossine, F. W., Plotkin, J. B. & Levin, S. Optimal, near-optimal, and robust epidemic control. doi:10.31219/osf.io/9gr7q.
 31. Zhao, H. & Feng, Z. Staggered release policies for COVID-19 control: Costs and benefits of relaxing restrictions by age and risk. *Math. Biosci.* **326**, 108405 (2020).
 32. Hochberg, M. E. Importance of suppression and mitigation measures in managing COVID-19 outbreaks. doi:10.1101/2020.03.31.20048835.
 33. Havers, F. P. *et al.* Seroprevalence of Antibodies to SARS-CoV-2 in 10 Sites in the United

States, March 23-May 12, 2020. *JAMA Intern. Med.* (2020)

doi:10.1001/jamainternmed.2020.4130.

Disease-dependent interaction policies to support health and economic outcomes during the COVID-19 epidemic

Supplementary Information

Guanlin Li,^{1,2,*} Shashwat Shivam,^{3,*} Michael E. Hochberg,^{4,5,†} Yorai Wardi,^{3,‡} and Joshua S. Weitz^{6,2,7,§}

¹ *Interdisciplinary Graduate Program in Quantitative Biosciences,
Georgia Institute of Technology, Atlanta, GA, USA*

² *School of Physics, Georgia Institute of Technology, Atlanta, GA, USA*

³ *School of Electrical and Computer Engineering,
Georgia Institute of Technology, Atlanta, GA, USA*

⁴ *ISEM, University of Montpellier, CNRS, IRD, EPHE, Montpellier, France*

⁵ *Santa Fe Institute, Santa Fe, New Mexico, USA*

⁶ *School of Biological Sciences, Georgia Institute of Technology, Atlanta, GA, USA*

⁷ *Center for Microbial Dynamics and Infection, Georgia Institute of Technology, Atlanta, GA, USA*

(Dated: August 24, 2020)

Contents

I. Methods	2
A. Model and parameters	2
Initial conditions and contact rate	3
B. Optimal control formulation	3
C. Optimization algorithm	4
Simulation details	6
D. Optimised time-dependent contact rate	6
Varying the relative importance of the illness-related cost vs. socioeconomic costs	7
E. Inferring state-dependent heuristic policies from optimal control solutions	8
F. Sensitivity of optimal control and feedback control to mis-timed implementation of policies	8
II. Supplementary Tables	10
References	10

*These authors contributed equally

[†]Electronic address: michael.hochberg@umontpellier.fr

[‡]Electronic address: ywardi@ece.gatech.edu

[§]Electronic address: jswertz@gatech.edu; URL: <http://ecothery.biology.gatech.edu>

I. METHODS

A. Model and parameters

The epidemic model analyzed here is an 'SEIR' model, including SEIR dynamics [1], susceptibles S , exposed E , infectious I , recovered R , with the total number of fatalities denoted by D . The *force of infection* is the contact rate of infectious individuals multiplied by the probability that the interaction is with a susceptible person multiplied by the probability that the event leads to an infection. Let c_S , c_E , c_I and c_R equal the contact rate of S , E , I , and R individuals respectively. The force of infection is

$$F(S, E, I, R; c) = \eta_I c_S \left(\frac{c_I S}{c_S S + c_E E + c_I I + c_R R} \right), \quad (1)$$

where η_I is the measure of disease transmission effectiveness from I class to S class [2]. Suppose all the contact rates are equal to a baseline contact rate c_B . Assuming relatively few fatalities per-capita, then the total population is approximately constant at N , such that the force of infection can be approximated as $F \approx \eta_I c_B (I/N)$, which recovers the standard SEIR model by defining $\beta = \eta_I c_B$ as *infection rate*. In the generalized case, the model is given by the following system of nonlinear differential equations

$$\begin{aligned} \frac{dS}{dt} &= -F(S, E, I, R; c)I \\ \frac{dE}{dt} &= F(S, E, I, R; c)I - \frac{1}{T_{inc}}E \\ \frac{dI}{dt} &= \frac{1}{T_{inc}}E - \frac{1}{T_{inf}}I \\ \frac{dR}{dt} &= (1 - \mu)\frac{1}{T_{inf}}I \\ \frac{dD}{dt} &= \mu\frac{1}{T_{inf}}I, \end{aligned} \quad (2)$$

where $F(S, E, I, R; c)$ is the force of infection given in Eq. 1, and T_{inf} is the infectious period, T_{inc} is the incubation period, μ is the case fatality ratio for infected individuals, see SEIR model schematic in Fig. S1.

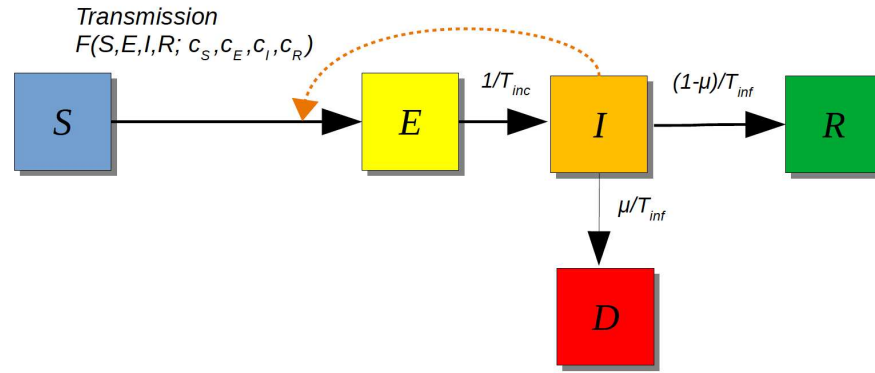


FIG. S1: SEIR Model Schematic. The interactions between susceptible and infected individuals lead to newly exposed cases. These interactions are modeled by the force of infection $F(S, E, I, R; c)$. The exposed individuals undergo an incubation period (T_{inc}) before the onset of infectiousness. Infectious individuals will either recover (and develop protective immunity) or die after an infectious period (T_{inf}), see model equations in Eq. 2.

At the start of an outbreak ($t = 0$), the contact rate of each individual is c_B . The basic reproduction number \mathcal{R}_0 of this compartmental model is

$$\mathcal{R}_0 = \overbrace{T_{inf}}^{\text{average infectious time}} \times \overbrace{\eta_I c_B}^{\text{infection rate}}. \quad (3)$$

The assumed model parameters used in the model are shown in table S3.

Initial conditions and contact rate

The model assumes an initial total population of 1,000,000 (one million) denoted as N_0 . An initial outbreak is seeded in this population given one infected individual. The simulation is run for two months (60 days) with fixed baseline contact rate (all c 's are equal to c_B) - we use this time point (which we denote time t_0 in our simulations) as the time at which intervention policies might be applied. The individual behaviors (or activities) are quantified by contact rates, $c_S(t)$, $c_E(t)$, $c_I(t)$ and $c_R(t)$, thus the control of contact rates can be mapped to intervention strategies. For examples, when c_S goes down that means shelter-in-place, when c_E or c_I goes down that can mean quarantine and isolation. When there is no intervention during the epidemic, the population dynamics of SEIR model is shown in Fig. S2.

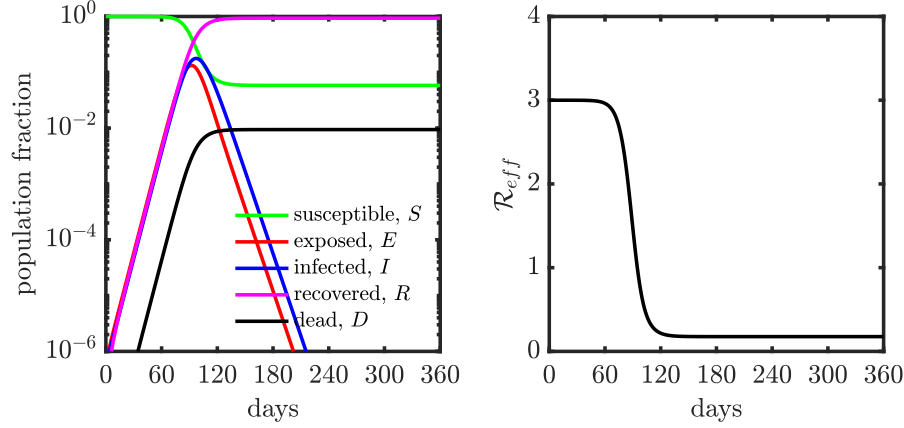


FIG. S2: Population dynamics of SEIR model without control, *i.e.*, all c 's are equal c_B from t_0 to t_f . (Left) Population dynamics of SEIR model without control. (Right) Epidemic growth as measured by the effective reproduction number (R_{eff}), the basic reproduction number R_0 is 3.

We denote $c(t) = (c_S(t), c_E(t), c_I(t), c_R(t))^T$ as the contact rate vector at time t , where $(\cdot)^T$ denotes a matrix transpose. The effective reproduction number $R_{eff}(t)$ is

$$R_{eff}(t) = \underbrace{\text{average infectious time}}_{T_{inf}} \times \underbrace{\text{infection rate in susceptible population}}_{\eta_I(t)c_I(t)\frac{c_S(t)S(t)}{Q(t)}}, \quad (4)$$

where $Q(t) = c_S(t)S(t) + c_E(t)E(t) + c_I(t)I(t) + c_R(t)R(t)$.

B. Optimal control formulation

Here, we aim to optimally deploy public health control strategies (via control of the contact rates) that minimizes deaths during the outbreak while keeping the socioeconomic costs low. The baseline average of contacts at time t , $Q(t)$, equals $c_B N(t)$, where c_B is the baseline contact rate and $N(t) = S(t) + E(t) + I(t) + R(t)$ is the total number of alive individuals. The socioeconomic costs may result from two consequences of contact rate interventions: (1) loss of essential connections; (2) shifting of roles relative to baseline, as quantified by variation in contact rates by testing status. Given the intervention policy $c(t)$ at time t , we quantify the socioeconomic costs as $\mathcal{E}_1 + \mathcal{E}_2$, where \mathcal{E}_1 and \mathcal{E}_2 may take form of

$$\begin{aligned} \mathcal{E}_1(Q(t); W_1) &= W_1 \exp \left(K \left(\frac{c_B N(t) - Q(t)}{c_B N(t)} \right) \right) \\ \mathcal{E}_2(c(t); W_2) &= \frac{1}{2} \sum_{i \in \{S, E, I, R\}} w_i \left(\frac{c_B - c_i(t)}{c_B} \right)^2 \\ &= \frac{1}{2} \left(\frac{c_B - c(t)}{c_B} \right)^T W_2 \left(\frac{c_B - c(t)}{c_B} \right), \end{aligned} \quad (5)$$

where $W_1 \geq 0$, $W_2 = \text{diag}(w_S, w_E, w_I, w_R) \succ 0$ are weight parameters, and K is a constant shape parameter for the exponential function. In the absence of any strategy, there are few deaths. This would mean that $N(t) \approx N_0$. For the sake of simplicity, \mathcal{E}_1 is approximated as $W_1 \exp(K(c_B N_0 - Q(t))/c_B N_0)$. For every t , we restrict the contact rate vector $c(t) \in [c_{\min}, c_{\max}]^4$, where c_{\min} and c_{\max} are the minimum contact rate and maximum contact rate. The set $C = [c_{\min}, c_{\max}]^4$ is a convex and compact set in \mathbb{R}^4 . We define the space of *admissible controls*, denoted by \mathcal{A} , as the set of Lebesgue-measurable functions $c : [t_0, t_f] \rightarrow C$. The (dimensionless) cost functional is written in the following Bolza form,

$$\mathcal{J}(c) = \overbrace{W_D \frac{D(t_f)}{N_0}}^{\text{costs of deaths}} + \int_{t_0}^{t_f} \overbrace{\mathcal{E}_1(Q(t); W_1) + \mathcal{E}_2(c(t); W_2)}^{\text{socioeconomic costs}} + \overbrace{W_I \frac{I(t)}{N_0}}^{\text{infectious case costs}} dt, \quad (6)$$

where $W_I \geq 0$ and $W_D \geq 0$ are weight regulators. The optimal control problem is

$$\min \{\mathcal{J}(c) | c \in \mathcal{A}\}. \quad (7)$$

The weight parameters used in Eq. 6 are shown in table S4, these weight parameters are set to balance the scale of different cost components in Eq. 6.

However, frequent and continuous updates bring in practical issues in terms of (1) establishing how subtle changes in contact rates actually translate into applicable changes in individual behaviors (*i.e.*, a continuous strategy can not be practically implemented), and (2) individual fatigue (and reduced adherence to proposed measures) due to frequent behavioral adjustments (*i.e.*, frequent update is also impractical). To address these concerns, we need to modify the space of controls \mathcal{A} .

To begin with, we fix the full time horizon of control as $[t_0, t_f]$. For example, t_0 is fixed to 60 days after outbreak (see Sect. 1A) and t_f to 360 days after outbreak, where we set the timing of the outbreak to 0. We only consider an intervention policy which is updated a finite number of times in the interval $[t_0, t_f]$. These epochs are denoted by $t_0 < t_1 < t_2 < \dots < t_M = t_f$. Let \mathcal{O} denote the corresponding partition of $[t_0, t_f]$, and denote by $\mathcal{F}(\mathcal{O})$ the space of functions $c : [t_0, t_f] \rightarrow \mathbb{R}^4$ which have constant values in the intervals $[t_{m-1}, t_m)$, $m = 1, 2, \dots$. The resulting optimal control problem is

$$\min \{\mathcal{J}(c) | c \in \mathcal{A} \cap \mathcal{F}(\mathcal{O})\}. \quad (8)$$

In most cases, the exposed individuals are unaware of their infection status, so it's impractical to control their behaviors separately from the susceptible individuals. Hence, we view them as susceptibles for the purpose of policy setting and set $c_E(t) = c_S(t) \forall t \in [t_0, t_f]$. In doing so, the contact rate vector $c(t)$ actually is a three dimensional vector consisting of $c_S(t)$, $c_I(t)$ and $c_R(t)$.

C. Optimization algorithm

We follow Polak [3] who developed a general framework for computational optimal control based on optimization on Hilbert spaces. We first discuss the conceptual algorithm in an abstract setting of infinite-dimensional spaces, and then provide approximation details in the sequel.

Let $L_2^n[t_0, t_f]$ denote the space of equivalence classes of square integrable functions from $[t_0, t_f]$ into \mathbb{R}^n ($n = 3$ in our case). We use the notation $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ for the norm and scalar product in Euclidean space (e.g. \mathbb{R}^n), and use $\|\cdot\|_{\mathcal{H}}$ and $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ for the norm and scalar product in a Hilbert space \mathcal{H} . For example, given $u, v \in L_2^n[t_0, t_f]$ (a Hilbert space), the norm $\|\cdot\|_2$ and scalar product $\langle \cdot, \cdot \rangle_2$ are: $\|u\|_2 = \sqrt{\int_{t_0}^{t_f} \|u(t)\|^2 dt}$; $\langle u, v \rangle_2 = \int_{t_0}^{t_f} \langle u(t), v(t) \rangle dt$.

In the optimal control problem under consideration, we define the state variable of the system, $x \in \mathbb{R}^5$, by $x = (S, E, I, R, D)^T$. We also define the control variable, $c \in \mathbb{R}^3$, by $c = (c_S, c_I, c_R)^T$. The dynamic equation of the system, Eq. 2, can be rewritten in a compact form:

$$\dot{x}(t) = f(x(t), c(t)), \quad (9)$$

where dot represents differentiation with respect to time t , and $f(x(t), c(t))$ is the right hand side (RHS) of equations 2. The cost functional defined by Eq. 6 can be written succinctly as

$$\mathcal{J}(c) = \phi(x(t_f)) + \int_{t_0}^{t_f} L(x(t), c(t)) dt, \quad (10)$$

where $L(x(t), c(t)) := \mathcal{E}_1(Q(t); W_1) + \mathcal{E}_2(c(t); W_2) + W_I(I(t)/N_0)$ and $\phi(x(t_f)) = W_D(D(t_f)/N_0)$. The costate (adjoint) variable, $\lambda(t) \in \mathbb{R}^5$, is defined by the following equation,

$$\dot{\lambda}(t) = - \left(\frac{\partial f}{\partial x}(x(t), c(t)) \right)^T \lambda(t) - \left(\frac{\partial L}{\partial x}(x(t), c(t)) \right)^T, \quad \lambda(t_f) = \nabla_x \phi(x(t_f)). \quad (11)$$

Next, recall the optimal control problem defined by Eq. 8, where the cost functional (performance integral) $\mathcal{J} : \mathcal{A} \rightarrow \mathbb{R}$ is defined by (10). We will propose a gradient-descent algorithm for solving this problem, where the gradient $\nabla \mathcal{J}(c(t))$ is given by the following equation,

$$\nabla \mathcal{J}(c(t)) = \left(\frac{\partial f}{\partial c}(x(t), c(t)) \right)^T \lambda(t) + \left(\frac{\partial L}{\partial c}(x(t), c(t)) \right)^T; \quad (12)$$

see [3–5].

The principal elements of most gradient-descent algorithms (including the one described here) are the descent direction and the stepsize. For our problem, the descent direction is based on $-\nabla \mathcal{J}(c(t))$; in fact, it is the projection of $c(t) - \nabla \mathcal{J}(c(t))$ into the constraint-set $C = [c_{min}, c_{max}]^3$ as defined in Sect. 1B. The stepsize that we use is the Armijo stepsize (see [3]) described in the sequel, which gives an approximate line minimization of \mathcal{J} in the computed descent direction from the control $c(t)$.

Consider first the conceptual algorithm. In its formal presentation by Algorithm 1, below, we use the term “compute” in a conceptual sense, since the “computations” refer to elements in the Hilbert space $L_2^n[t_0, t_f]$. Also, we use the notation $\mathcal{P}_{\mathcal{A}}$ to mean the projection from $L_2^n[t_0, t_f]$ into the space \mathcal{A} .

Algorithm 1: Conceptual.

Given a control iteration $c^{(i)} \in \mathcal{A}$, compute from it the next iteration, $c^{(i+1)} \in \mathcal{A}$, as follows.

Parameters: Constants $\alpha \in (0, 1)$ and $\beta \in (0, 1)$. (These parameters play a role in the computation of the step size.)

1. Compute the state trajectory $x(t)$ and costate trajectory $\lambda(t)$ by Eqs. 9 and 11, respectively.
2. Compute $\nabla \mathcal{J}(c(t))$ by Eq. 12.
3. (Here we define, for all $y \geq 0$, $h^{(i)}(y) = \mathcal{P}_{\mathcal{A}}(c^{(i)} - y \nabla \mathcal{J}(c^{(i)}))$). Compute $\ell(i) > 0$ defined as follows,

$$\ell(i) = \min\{\ell = 0, 1, \dots | \mathcal{J}(h^{(i)}(\beta^\ell)) - \mathcal{J}(c^{(i)}) \leq -\alpha \langle \nabla \mathcal{J}(c^{(i)}), c^{(i)} - h^{(i)}(\beta^\ell) \rangle_2\}.$$

4. Set $\gamma(i) = \beta^{\ell(i)}$, and set

$$c^{(i+1)} = \mathcal{P}_{\mathcal{A}}(c^{(i)} - \gamma(i) \nabla \mathcal{J}(c^{(i)})).$$

We remark that $\mathcal{P}_{\mathcal{A}}(c^{(i)} - \gamma(i) \nabla \mathcal{J}(c^{(i)})) - c^{(i)}$ is descent direction for \mathcal{J} from $c^{(i)}$, and $\gamma(i)$, the Armijo stepsize, provides an approximate-line minimization in that direction.

The main modification of Algorithm 1 towards implementation is in numerical solutions of the state equation and costate equation. Furthermore, the time-interval $[t_0, t_f]$ has to be discretized by a suitable grid in order to adequately represent various time-dependent functions such as $c^{(i)}(t)$ and $\nabla \mathcal{J}(c(t))$. All of this can be achieved by a common grid, $\mathcal{G}_0 := \{t_0, t_1, \dots, t_N\} \subset [t_0, t_f]$. Furthermore, the projection $\mathcal{P}_{\mathcal{A}}(c^{(i)} - \beta^\ell \nabla \mathcal{J}(c^{(i)}))$ in Step 3 of Algorithm 1 can be computed as follows: For a given grid-point $t_j \in \mathcal{G}_0$, project the point $c^{(i)}(t_j) - \beta^\ell \nabla \mathcal{J}(c^{(i)})(t_j) \in \mathbb{R}^3$ into the set C , where $C = [c_{min}, c_{max}]^3 \subset \mathbb{R}^3$. This can be done by a simple co-ordinate projection since C is a box. Denote the result of this pointwise projection by $\mathcal{P}_C(c^{(i)}(t_j) - \beta^\ell \nabla \mathcal{J}(c^{(i)})(t_j)) \in C$, and observe that the function comprised of a zero-order or first-order interpolation of these points can serve to approximate the functional projection $\mathcal{P}_{\mathcal{A}}(c^{(i)} - \beta^\ell \nabla \mathcal{J}(c^{(i)}))$.

These modifications of the conceptual algorithm give an implementable version for it. The resulting implementable algorithm falls in Polak’s framework of numerical optimal control [3] which includes results pertaining to asymptotic convergence.

As noted, the control $c(t)$ must be piecewise constant and maintain constant values for substantial periods. At the same time, the state trajectory typically is continuous and not piecewise constant. Therefore, it is natural to maintain the present framework of continuous-time optimal control and not to consider the problem in the setting of discrete-time systems. To do that, we have to consider a form of projection of the control functions, $c(t)$, into the space of piecewise-constant functions. We next describe the specific way we did that in our simulations.

Consider an increasing set of points $\{t_0 = t_0 < t_1 < t_2 < \dots < t_M = t_f\}$, and the corresponding partition of the time-interval $[t_0, t_f]$ by the subintervals $[t_{m-1}, t_m]$, $m = 1, \dots, M$. Denote the interval $[t_{m-1}, t_m]$, $m = 1, \dots, M$, by Δ_m ; right-close the last interval to be $\Delta_M = [t_{M-1}, t_M]$, and denote the partition by \mathcal{O} . Let $\psi_m(t)$, $m = 1, \dots, M$ denote the indicator function of the subinterval ψ_m , namely $\psi_m(t) = 1$ for $t \in \psi_m$ and $\psi_m(t) = 0$ otherwise. A continuous-time signal can be approximated over an interval by a constant equal to its average value over that interval. Given a continuous-time policy $c : [t_0, t_f] \rightarrow \mathbb{R}^3$, the piecewise constant reconstruction is

$$\tilde{c}(t) = \sum_{m=1}^{M-1} \theta_m \psi_m(t), \quad \theta_m = \frac{1}{|\Delta_m|} \int_{t_m}^{t_{m+1}} c(s) \, ds, \quad (13)$$

here $\theta_m \in \mathbb{R}^3$. Let $\mathcal{V}_{\mathcal{O}}$ denotes the operator comprising this piecewise- constant reconstruction procedure. Then Eq. 13 can be written as $\tilde{c} = \mathcal{V}_{\mathcal{O}}(c)$. Now we modify the algorithm by replacing the projection operator $\mathcal{P}_{\mathcal{A}}$ in Step 3 by the operator $\mathcal{V}_{\mathcal{O}} \circ \mathcal{P}_{\mathcal{A}}$, comprised of $\mathcal{P}_{\mathcal{A}}$ followed by $\mathcal{V}_{\mathcal{O}}$.

In summary, the implementable algorithm that we use consists of the following modification of Algorithm 1:

1. Compute $x(t)$ and $\lambda(t)$ by numerical integrations using the forward-Euler method.
2. Compute $\nabla \mathcal{J}(c(t))$ only at the points t_j on the grid \mathcal{G}_0 .
3. Redefine $h^{(i)}(y)$ as $h^{(i)}(y) = \mathcal{V}_{\mathcal{O}} \circ \mathcal{P}_{\mathcal{A}}(c^{(i)} - y \nabla \mathcal{J}(c^{(i)}))$.
4. Replace $\mathcal{P}_{\mathcal{A}}$ by $\mathcal{V}_{\mathcal{O}} \circ \mathcal{P}_{\mathcal{A}}$.

A flow chart of the algorithm is depicted in Fig. S3.

Simulation details

The model assumes an initial total population of 1,000,000 (one million) denoted as N_0 . An initial outbreak is seeded in this population given one infected individual. The simulation is run for two months (60 days) with fixed baseline contact rate (all c 's are equal to c_B) - we use this time point (which we denote time t_0 in our simulations) as the time at which intervention policies might be applied: we have $c(t) = c_B \, \forall t \in [0, t_0]$, and we set the initial state to $x_0 = x(t_0)$.

For the algorithm, we set $\alpha = 0.1$ and $\beta = 0.5$, the grid's time increments to 0.05, $\Delta_m = 30$ days, $t_0 = 60$ and $t_f = 360$ days. For the uncontrolled part of the simulation we take $c(t) = c_B$ for all $t \in [0, t_0]$. The term $\Theta(i) := |\mathcal{J}(c^{(i+1)}) - \mathcal{J}(c^{(i)})| / \mathcal{J}(c^{(i)})$ acts as a convergence indicator for stopping the algorithm. The algorithm is terminated whenever $\Theta(i) \leq 10^{-6}$.

D. Optimised time-dependent contact rate

One way of estimating the economic impact is measuring the *fraction of working-days recovered*, i.e., how many days are people working. We assume that if a person has a contact rate less than the base contact rate (c_B), it implies a proportional reduction of in-person contacts that have a direct economic benefit. For example, for a susceptible individual, the working fraction at day t can be computed as $c_S^*(t) = \min\{1, c_S(t)/c_B\}$, where c_B is the baseline contacting rate. This ensures that contact rates above the baseline does not correspond to an increase in hours worked. For a given contacting rate policy $c = (c_S, c_E, c_I, c_R)^T$, the working days restored in the period $[t_0, t_f]$ is

$$\Psi_{restored} = \int_{t_0}^{t_f} c_S^*(t)S(t) + c_E^*(t)E(t) + c_I^*(t)I(t) + c_R^*(t)R(t) \, dt, \quad (14)$$

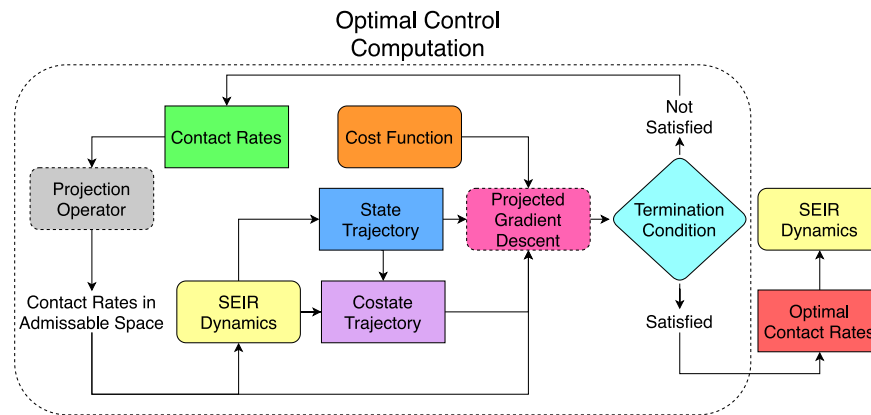


FIG. S3: Flowchart of optimal contact rate computation. The contact rates are projected into the allowed subspace, and then fed to the system dynamics to generate the state and costate trajectory. Using the cost function provided and the state and costate trajectories, a projected gradient descent algorithm is employed to update the contact rates. If the termination condition is satisfied, these new contact rates are accepted as the optimal contact rates. Otherwise, the contact rates are replaced with the computed values and the loop is rerun.

where $c_k^*(t) = \min\{1, c_k(t)/c_B\}$, $k \in \{S, E, I, R\}$, is the working fraction of an individual in subpopulation class k at day t . When there is no policy intervention, the contacting rate of each individual is c_B for all time, the working days will be maximized, which can be written as

$$\Psi_{max} = \int_{t_0}^{t_f} S(t) + E(t) + I(t) + R(t) dt. \quad (15)$$

Then, the fraction of working days restored can be written as

$$\Psi(c) = \frac{\Psi_{restored}}{\Psi_{max}}. \quad (16)$$

Economic costs are proportional to $1 - \Psi(c)$; such that economic impacts of policies are minimized as $\Psi(c)$ approaches 1

Varying the relative importance of the illness-related cost vs. socioeconomic costs

The relative importance given to the socioeconomic impact of the epidemic and the cost associated with deaths and infection spread is critical to design intervention policies. From Eq. 6, we note that the cost function \mathcal{J} can be decomposed into the two parts, socioeconomic costs (\mathcal{J}_E) and costs of infectious diseases (\mathcal{J}_I), $\mathcal{J} = \mathcal{J}_E + \mathcal{J}_I$, where

$$\begin{aligned} \mathcal{J}_E &= \int_{t_0}^{t_f} \mathcal{E}_1(Q(t); W_1) + \mathcal{E}_2(c(t); W_2) dt \\ \mathcal{J}_I &= W_D \frac{D(t_f)}{N_0} + \int_{t_0}^{t_f} W_I \frac{I(t)}{N_0} dt. \end{aligned} \quad (17)$$

In order to explore how to manage and control disease outbreak with low socioeconomic costs, we need to balance the importance of these two seemingly conflicting objectives. To this end, we define a new cost function where the relative importance of \mathcal{J}_E and \mathcal{J}_I can be parameterized. With this parameter, optimal policy decisions can be made based on the relative importance assigned to the socioeconomic effect and public health impacts. Define the relative weight ratio as ξ and the altered cost function is $\mathcal{J}(c; \xi) = \mathcal{J}_E + \xi \mathcal{J}_I$. As ξ increases, the relative importance of costs associated with death and spread of infection are increased vis-a-vis the socioeconomic impact and vice versa. The optimal intervention policies with different relative importance ratios ($\xi = 10^{-4}$, 1, and 10^4) vary with isolation efficiencies (25%, 50%, 75%) are shown in Figs. S4, S5, S6.

For example of balance case ($\xi = 1$) in the case of 50% isolation efficiency, Fig. S5 (middle) shows that the way of managing the epidemic while considering the socioeconomic impact is to isolate all the infected cases and enhance the interactions of recovered cases (*i.e.*, *shielding*, [6]), the susceptible cases stay home first and slowly go back to normal towards the end of intervention period. Using the optimized intervention policy (*i.e.*, optimal contact rate), the reduction in final deaths ($D(t_f)$) is substantial ($\sim 75\%$). The fraction of working-days recovered is about $\sim 66\%$. The results suggest advantages of isolation of infected cases and shielding enhancement of recovered cases.

E. Inferring state-dependent heuristic policies from optimal control solutions

In the previous sections, an optimal policy based on contact rates of different sub-populations was derived. However, as is often the case with optimal control solutions, these policies are open-loop, that is, once computed, they are prescribed to the system without feedback. Therefore, open-loop policies can have large sensitivities to modeling errors and signal-delays in the loop, especially if the system is unstable, as is the case with dynamic models of epidemics near the disease-free equilibrium at the onset of an outbreak. To address this, we design a feedback law for optimizing a ‘reasonable’ performance metric which, though different from the one of the optimal control problem, gives similar respective measures of two important performance metrics: total mortality ($D(t_f)$), and working fractions (Ψ). In Sect. 1F we verify this point by comparative simulations of the optimal control solution and the feedback control system, and not surprisingly, the sensitivity of the former with respect to a one-cycle (30-days) delay is much larger than that of the latter.

The optimal control solutions reveal a general structure that guides us in the choice of the feedback control law. The patterns observed in the computed optimal-policy solutions (see Figs. S4, S5, S6) suggest that the contact rates for the infected cases should be the smallest possible while the recovered population should be used as shields. Further, they display switch overs between the low and high contact rates for the susceptible populations, and these tend to occur towards the end of the lock-down period in a way that depends on the numbers of infected and recovered individuals.

To find the state feedback policy, the I - R (Infected-Recovered) phase plane is divided into 2 parts - one where the susceptibles need to isolate and another where they can get back to work. If the number of infected cases is high, the susceptible population should isolate to reduce the infection spread. For simplicity, the I - R plane is divided by a line, though better classification boundaries may exist. On one side of the line, the contact rate for the susceptible population is set to the minimum while on the other side, it is equal to the baseline (corresponding to opening the system).

The optimal slope of the dividing line is found via a *genetic algorithm* [7] using `matlab`’s built-in optimization function `ga`, with the maximum generation number (set to 30) serving as the stopping criterion [8]. The intercept of the line is fixed to a small negative number (10^{-3}) as the results are not affected by the intercept. The slope parameter is called θ , and the following cost function is minimized:

$$\mathcal{J}(\theta) = \tilde{w} \frac{D(t_f)}{N_0} + \Psi(c), \quad (18)$$

where $D(t_f)/N_0$ gives the normalized casualties, $\Psi(c)$ is the fraction of working-days restored and \tilde{w} is a weight to bring the two terms on the same scale (see table S4). The output of the genetic algorithm is θ^* , the slope of the optimal dividing line. We assume that the divided plane on the side of the line with the origin, *i.e.*, zero infected and recovered cases, is the side where the susceptibles are free to work, while the divided plane on the other side corresponds to lockdown. The simulation is run for different isolation and shielding levels and the compiled results are shown in Fig. S9. The isolation efficiencies vary from 25% to 75% while the shielding ratio varies from 2 to 5.

We note that as the efficiency of isolation increases, the susceptible population can be let out for a longer duration towards the end of the epidemic. This is because the reduced contact rate of the infected population reduces the spread of infection. If the isolation efficiency is 65% or more, there is no need for the susceptible population to be in lockdown. Similar observations can be made for the shielding ratio. While the effects of it are not as dramatic, an increase in shielding ratio decreases the force of infection, resulting in a lower time necessary for lockdown.

F. Sensitivity of optimal control and feedback control to mis-timed implementation of policies

We compute an optimal control for a system with a delayed input. We do not assume any modeling uncertainty, and attribute the perturbation only to the delayed application of the computed control. Suppose that an optimal control is computed for a scenario where its application is slated to commence 30 days after the outbreak of an epidemic, but the computed input control is uniformly delayed by 30 days. Then, by simulating the epidemic with a delayed policy,

Optimal Control Approach	Without Delay		With Delay	
Efficiency of Isolation (shielding = 2)	Deaths (per 10,000)	Working Fraction	Total Deaths (per 10,000)	Working Fraction
25%	600	90.95%	720	94.96%
50%	250	67.13%	510	83.65%
75%	40	99.92%	40	99.95%

TABLE S1: Comparison between optimal control approach for contact policy with and without delay. The comparisons are made for isolation efficiencies of 25%, 50% and 75%, with performance metrics of total deaths and working fraction. The total death is significantly higher for the system with delay when the isolation efficiency is 25% and 50%, suggesting poor robustness of the computed optimal control to input-delay.

Feedback Control Approach	Without Delay		With Delay	
Efficiency of Isolation (shielding = 2)	Deaths (per 10,000)	Working Fraction	Total Deaths (per 10,000)	Working Fraction
25%	620	89.98%	620	89.94%
50%	250	62.40%	250	62.12%
75%	30	99.90%	30	99.90%

TABLE S2: Comparison between feedback control approach for contact policy with and without delay. Isolation efficiencies of 25%, 50% and 75% are used for the purpose of comparison, with performance metrics of total deaths and working fraction. Both of the performance metrics are nearly identical for all cases, which suggests high robustness to delay on the system.

we can compare total deaths and working fraction for the delayed system with the results obtained without delay and with the same initial conditions.

The results, comprising total death and working fraction, are presented in table S1 for various isolation efficiencies. These findings indicate higher death rates as well as higher working fractions for the delayed system except for the third, and highest efficiency of isolation, where the respective performance measures are similar. For the case of 50% isolation, the death count is more than doubled due to delay. This significant difference in performance metrics demonstrates the shortcomings of implementing a policy based on optimal control if the system is not ideal.

We also tested the sensitivities of the total death and working fraction to 30-day delays. The results, shown in table S2 below, indicate low sensitivity compared to the observed results for the optimal control solution. We compared these performance metrics obtained from applications of the optimal control vs. feedback control to the system without delay. The results, summarized in tables S1 and S2 respectively, indicate similar performance with respect to the performance metrics used. Thus, feedback control is nearly as good as optimal control if there is no delay; while having more robust features in the event that there is a delay in application.

II. SUPPLEMENTARY TABLES

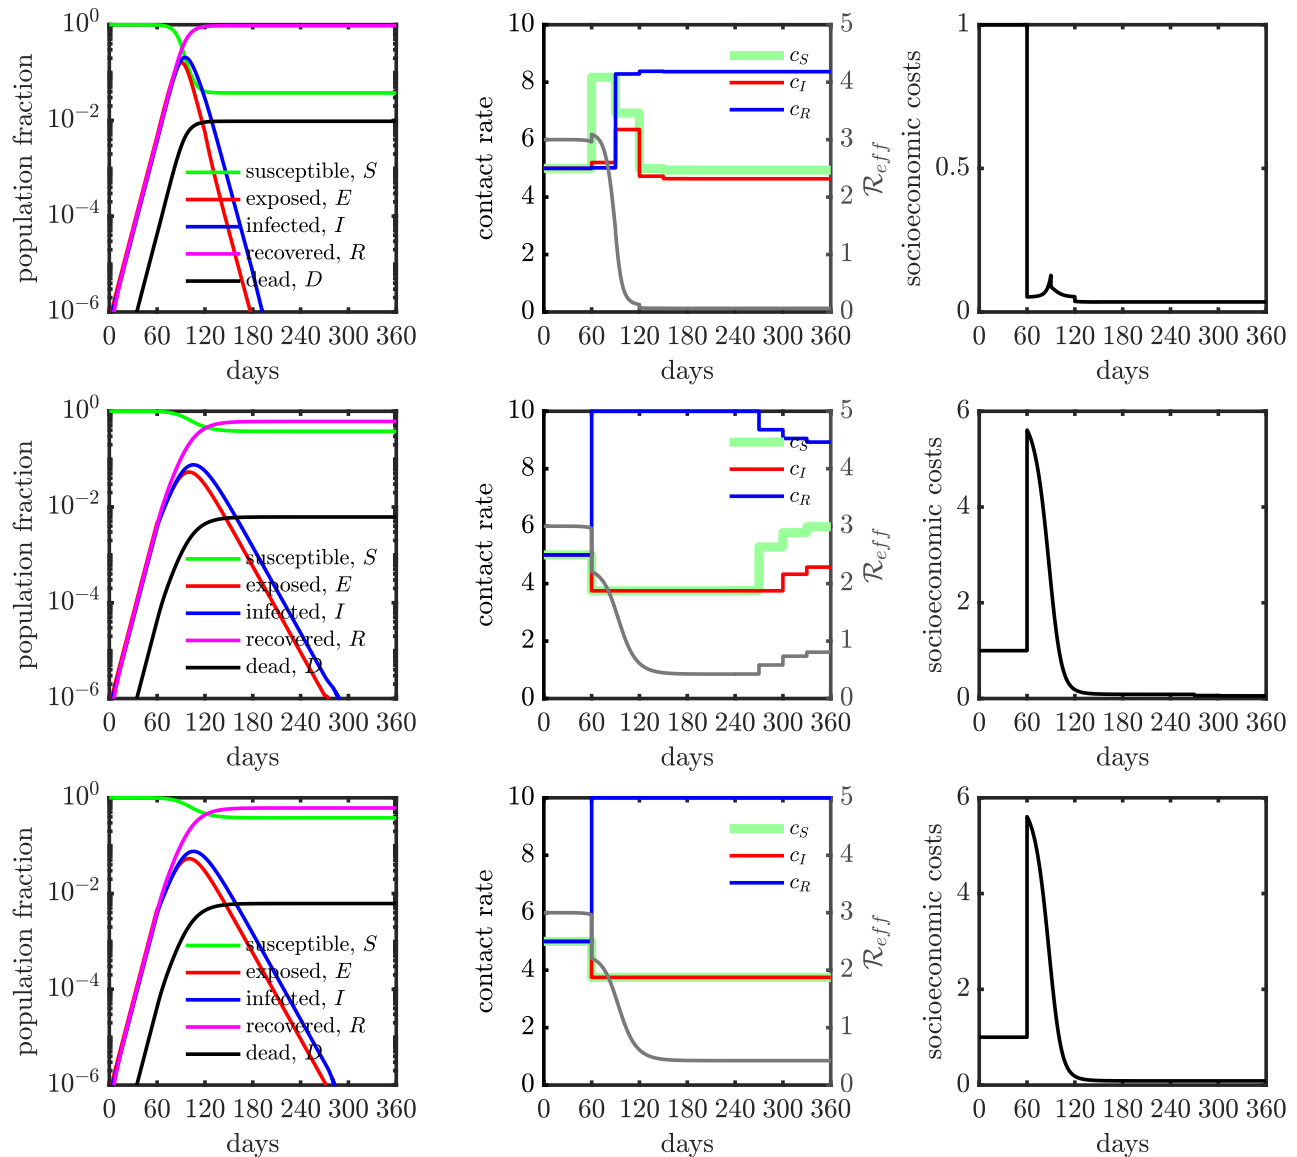
Parameter	Meaning	Value
η_I	Measure of disease transmission effectiveness	0.1
T_{inc}	Mean incubation period	4 days
T_{inf}	Mean infectious period	6 days
μ	Case fatality ratio of infected cases	0.01
c_B	Baseline potentially infectious contact rate	5/day

TABLE S3: Epidemiological characteristics.

Parameter	Meaning	Value
W_1	Weight of \mathcal{E}_1	1
W_2	Weight matrix of \mathcal{E}_2	$0.1\mathcal{I}$, \mathcal{I} is identify matrix
W_I	Weight of infection control cost	10^4
W_D	Weight of final deaths cost	10^5
K	Shape parameter in exponential function \mathcal{E}_1	7
\tilde{w}	weight of final deaths in heuristic cost	5×10^2

TABLE S4: Weight regulators and optimization parameters.

-
- [1] Anderson RM, Anderson B, May RM (1992) *Infectious diseases of humans: dynamics and control* (Oxford university press).
 - [2] Ngonghala CN, et al. (2020) Mathematical assessment of the impact of non-pharmaceutical interventions on curtailing the 2019 novel coronavirus. *arXiv preprint arXiv:2004.07391*.
 - [3] Polak E (1997) *Optimization Algorithms and Consistent Approximations* (Springer-Verlag, New York, New York).
 - [4] Kirk DE (2004) *Optimal control theory: an introduction* (Courier Corporation).
 - [5] Liberzon D (2011) *Calculus of variations and optimal control theory: a concise introduction* (Princeton University Press).
 - [6] Weitz JS, et al. (2020) Modeling shield immunity to reduce covid-19 epidemic spread. *Nature medicine* pp 1–6.
 - [7] Whitley D (1994) A genetic algorithm tutorial. *Statistics and computing* 4:65–85.
 - [8] The MathWorks I (2020) *Global Optimization Toolbox* (Natick, Massachusetts, United State).



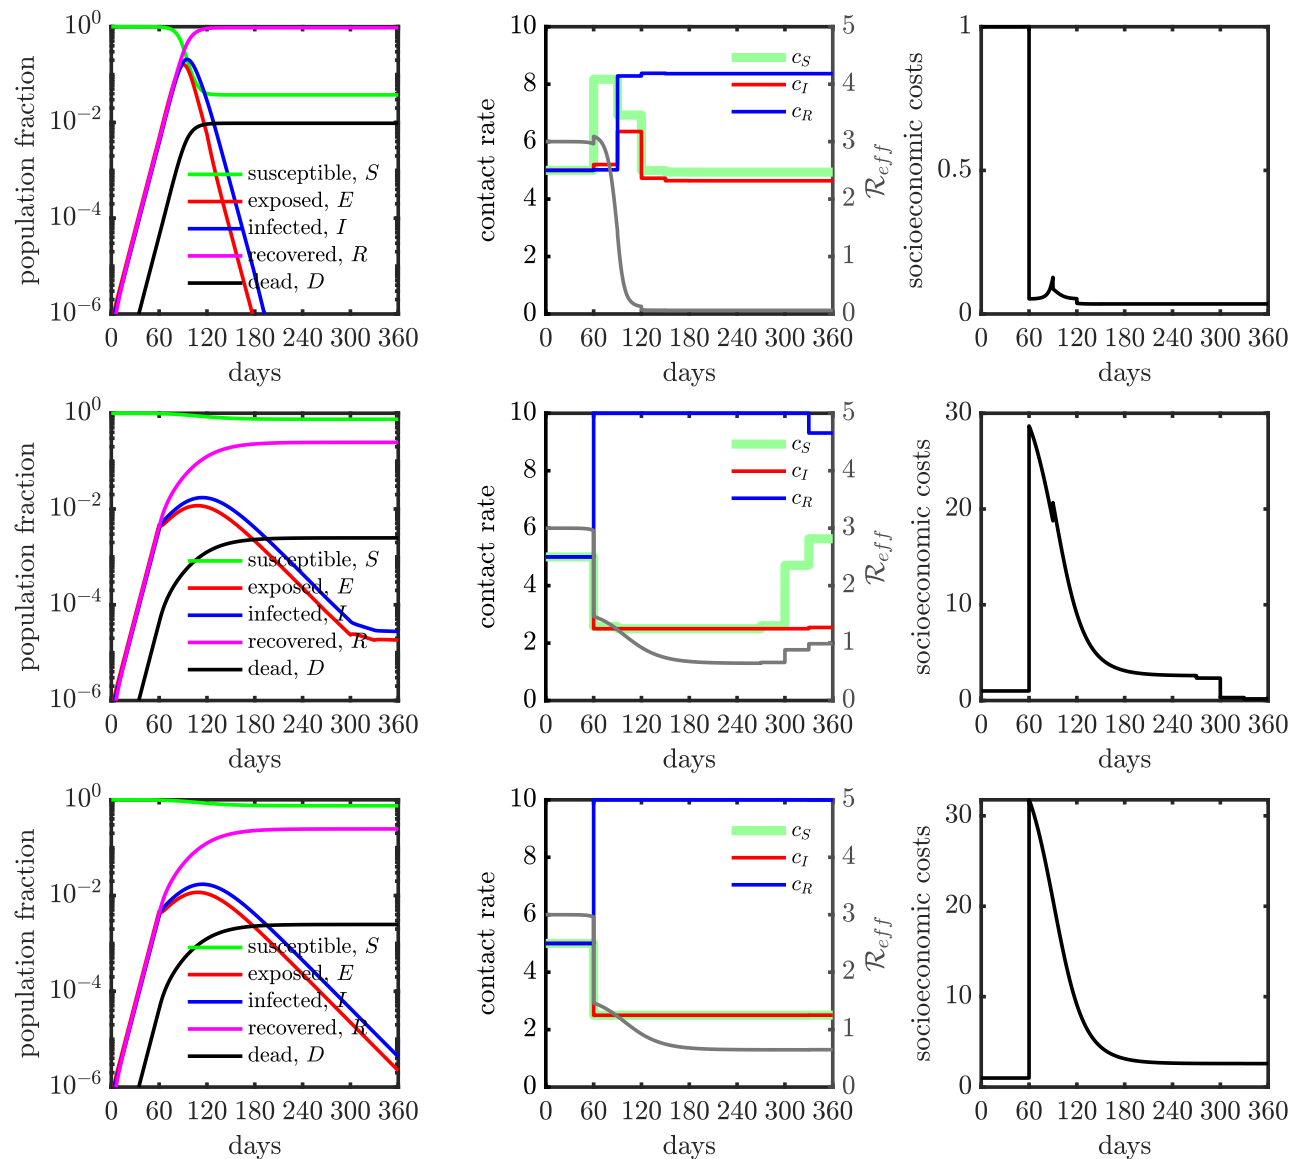


FIG. S5: Population dynamics in a SEIR model with controlled contact rate (50% isolation efficiency). (Top) The relative importance ratio (ξ) is set as 10^{-4} , the policy (*i.e.*, contact rate) is deployed to primarily minimize socioeconomic costs (\mathcal{J}_E). The socioeconomic costs are maintained at low level for all the policy intervention period (see the right panel after 60 days). Final state of the system in terms of population fraction is: $D(t_f) \approx 1\%$, $S(t_f) \approx 3\%$ and $R(t_f) \approx 96\%$, where $t_f = 360$ days. The fraction of working-days recovered is about 100%. (Middle) The relative importance ratio (ξ) is set as 1, the policy (*i.e.*, contact rate) is deployed to minimize infectious diseases costs (\mathcal{J}_I) while keeping the socioeconomic impacts low. For the first half intervention period (from 60 days to 240 days), the controlled contact rates of susceptible cases (c_S) and infected cases (c_I) are overlapped at minimum contact rate boundary (c_{min}). There are huge socioeconomic impacts at beginning of policy intervention period (due to the rapid contact rate reduction of susceptible cases and infected cases, *e.g.* isolation), the socioeconomic impacts become lower later because of *shielding* (*i.e.*, contact rate of recovered cases is enhanced). Moreover, at the last final 3 months, the socioeconomic costs drop more since the susceptible cases are back to normal (contact rate c_S is increasing). The Final state of the system in terms of population fraction is: $D(t_f) \approx 0.25\%$, $S(t_f) \approx 75\%$ and $R(t_f) \approx 25\%$, where $t_f = 360$ days. The fraction of working-days recovered is about 66%. (Bottom) The relative importance ratio (ξ) is set as 10^4 , the policy (*i.e.*, contact rate) is deployed to primarily minimize infectious diseases costs (\mathcal{J}_I). The controlled contact rates of susceptible cases (c_S) and infected cases (c_I) are overlapped at minimum contact rate boundary (c_{min}) for all time. Similar to the intermediate case ($\xi = 1$), large socioeconomic impacts are observed at beginning of policy intervention period. The Final state of the system in terms of population fraction is: $D(t_f) \approx 0.25\%$, $S(t_f) \approx 75\%$ and $R(t_f) \approx 25\%$, where $t_f = 360$ days. The fraction of working-days recovered is about 59%. Here, we set $c_{min} = c_B/2$ (*i.e.*, up to 50% isolation efficiency) and $c_{max} = 2c_B$ (*i.e.*, up to twice enhanced interactions).

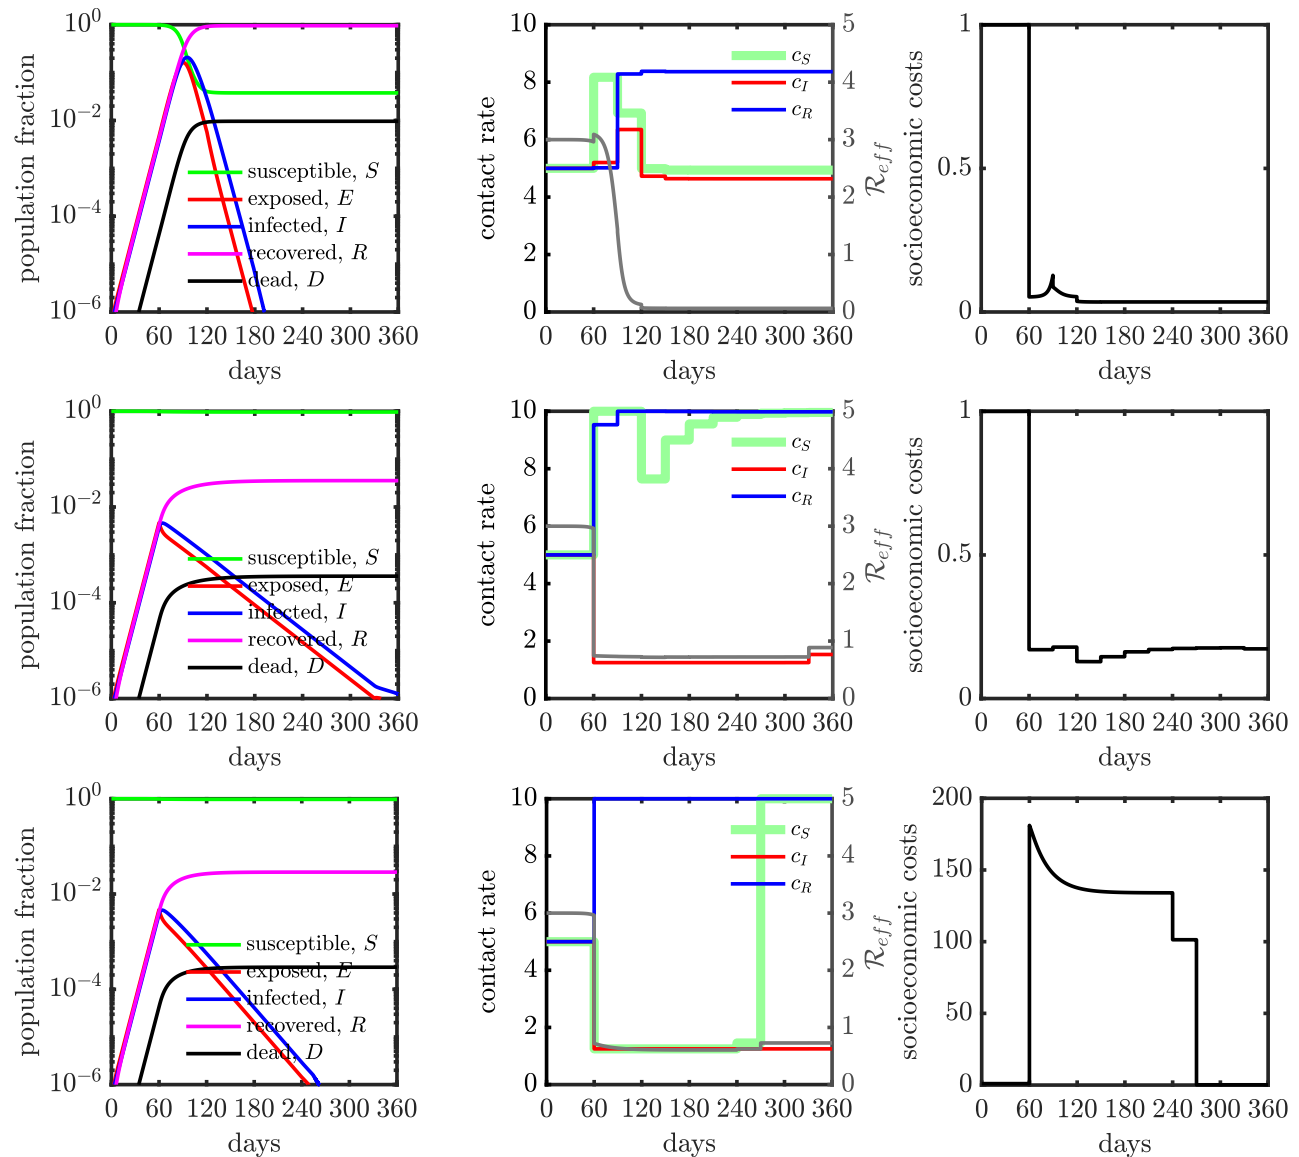


FIG. S6: Population dynamics in a SEIR model with controlled contact rate (75% isolation efficiency). (Top) The relative importance ratio (ξ) is set as 10^{-4} , the policy (*i.e.*, contact rate) is deployed to primarily minimize socioeconomic costs (J_E). The socioeconomic costs are maintained at low level for all the policy intervention period (see the right panel after 60 days). Final state of the system in terms of population fraction is: $D(t_f) \approx 1\%$, $S(t_f) \approx 4\%$ and $R(t_f) \approx 95\%$, where $t_f = 360$ days. The fraction of working-days recovered is about 100%. (Middle) The relative importance ratio (ξ) is set as 1, the policy (*i.e.*, contact rate) is deployed to minimize infectious diseases costs (J_I) while keeping the socioeconomic impacts low. Infected individuals are locked down at home for all intervention period while susceptible and recovered individuals are free to go out. The Final state of the system in terms of population fraction is: $D(t_f) \approx 0.04\%$, $S(t_f) \approx 96\%$ and $R(t_f) \approx 4\%$, where $t_f = 360$ days. The fraction of working-days recovered is about 100%. (Bottom) The relative importance ratio (ξ) is set as 10^4 , the policy (*i.e.*, contact rate) is deployed to primarily minimize infectious diseases costs (J_I). The controlled contact rates of susceptible cases (c_S) and infected cases (c_I) are overlapped at minimum contact rate boundary (c_{min}) at beginning, then the susceptibles are free to move later. The Final state of the system in terms of population fraction is: $D(t_f) \approx 0.03\%$, $S(t_f) \approx 97\%$ and $R(t_f) \approx 3\%$, where $t_f = 360$ days. The fraction of working-days recovered is about 50%. Here, we set $c_{min} = c_B/4$ (*i.e.*, up to 75% isolation efficiency) and $c_{max} = 2c_B$ (*i.e.*, up to twice enhanced interactions).

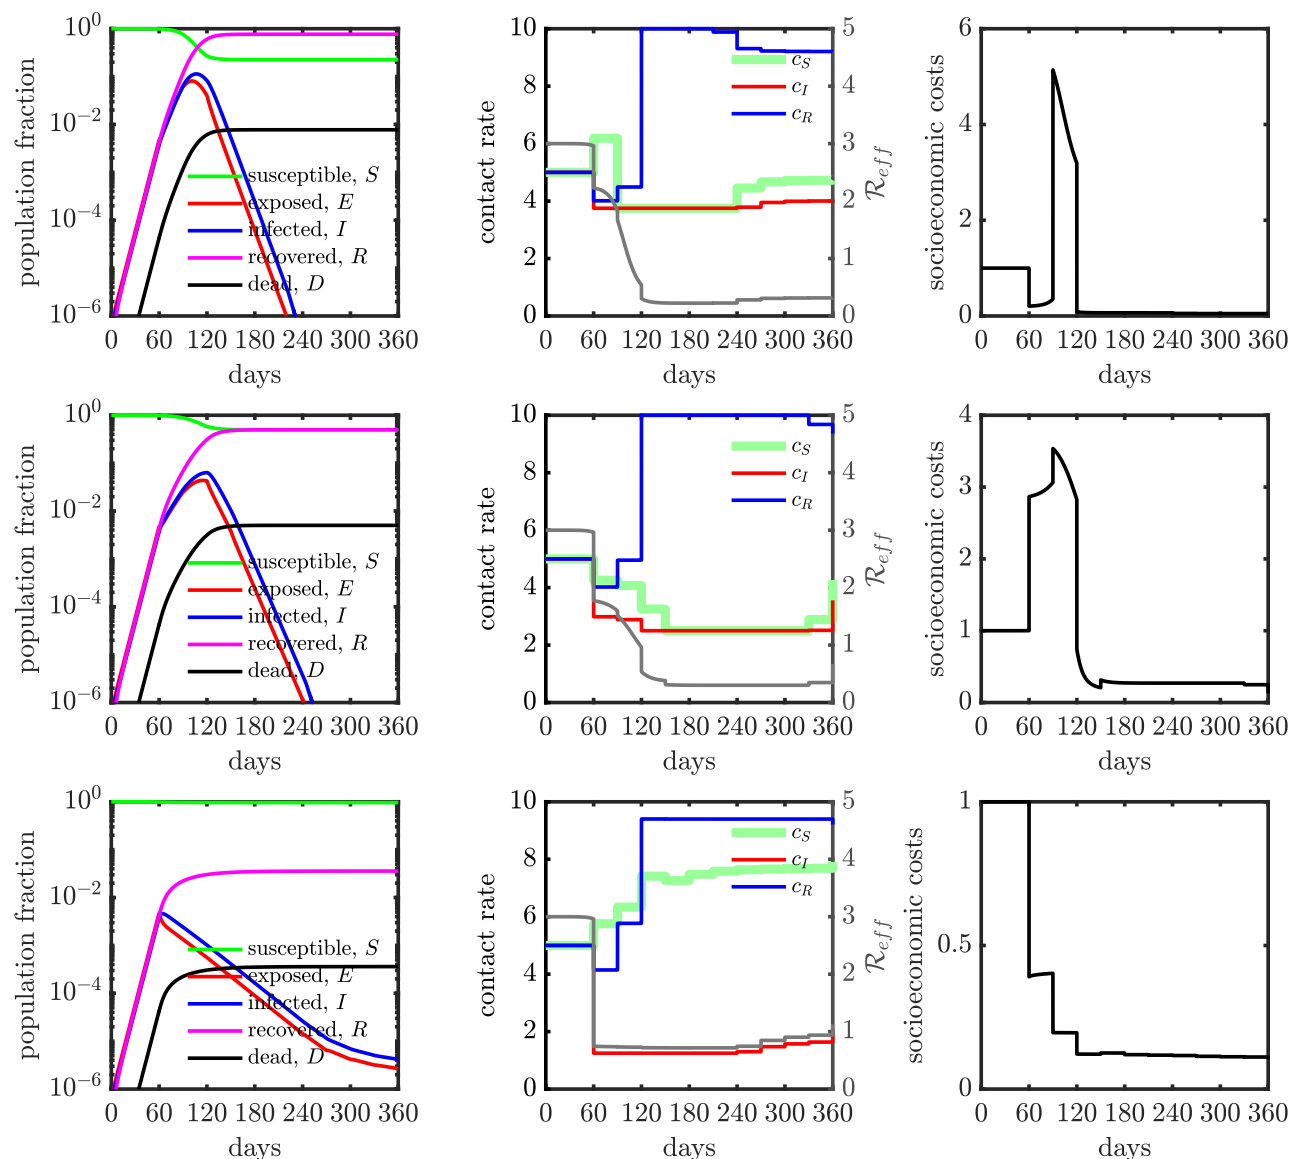


FIG. S7: Population dynamics in a SEIR model with mis-timed control policy for various isolation efficiencies: (Top) 25% isolation efficiency; (Middle) 50% isolation efficiency and (Bottom) 75% isolation efficiency. The relative importance (ξ) is 1 for all the cases. The optimal control policy is computed for a system one month after an outbreak. Then, we enforce the optimal control policy 30 days later, *i.e.*, at the end of 60 days after the start of the outbreak. The contact rate interventions start at 60 days. The total deaths and working fraction for the delayed system are presented in table S1 for various isolation efficiencies. The optimal control policies and their corresponding dynamics without mistiming are shown in the middle row of Figs. S4, S5, S6.

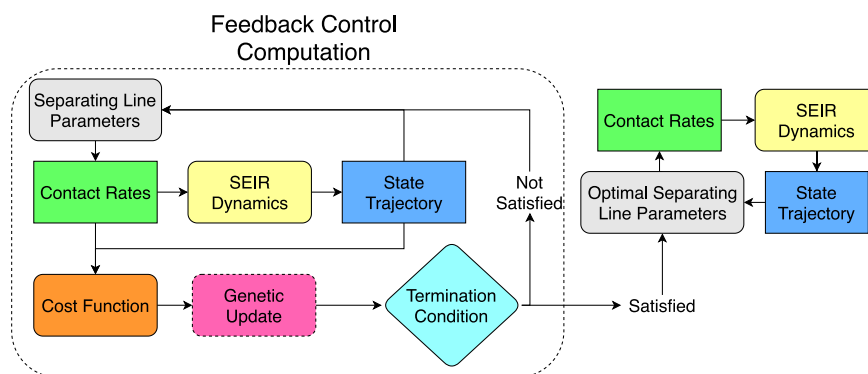
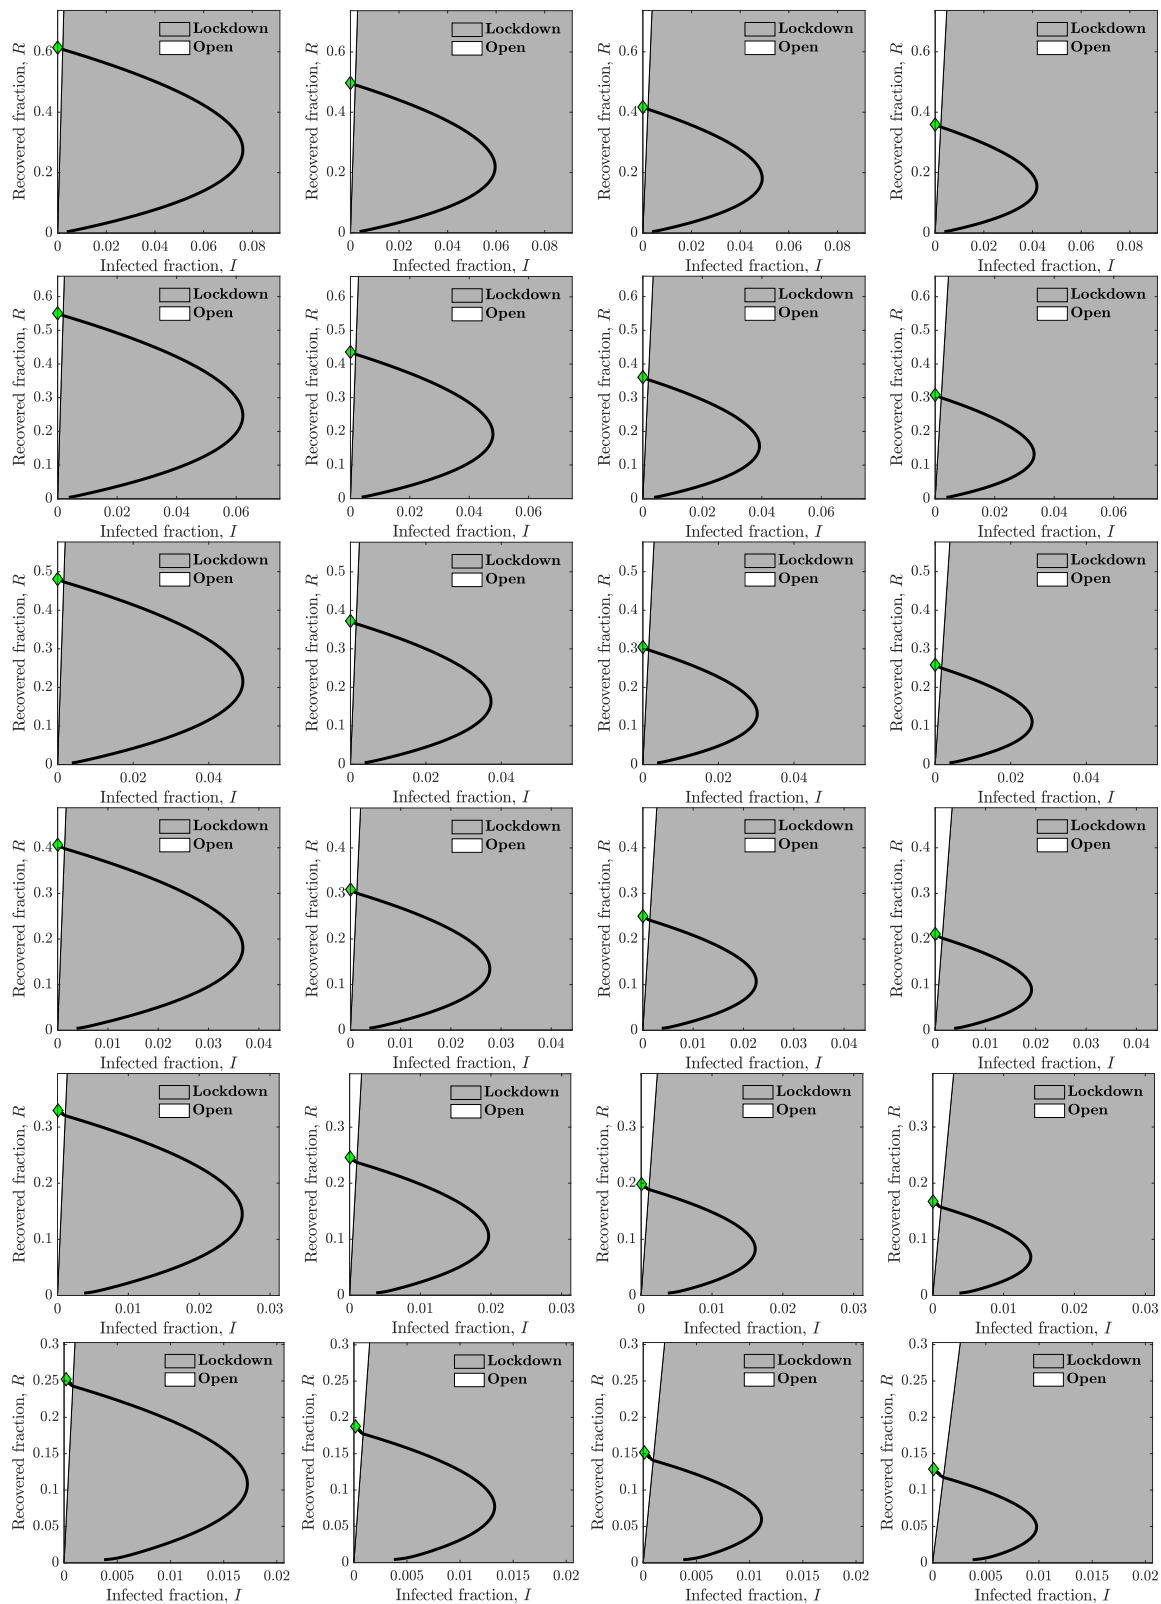


FIG. S8: Flowchart for computing state dependent heuristic policy. The line dividing the I - R plane is parameterized and for a particular parameter set, the contact rate is determined solely from the current infected and recovered cases. Starting with an initial parameter guess, the contact rate is computed from the current I and R values, giving the system's state at the next time instant. After the full state trajectory is computed, the associated cost is calculated and a genetic update step is taken which identifies the next parameters. If the termination condition is satisfied, these parameters are passed on as the optimal parameters which can be used to compute the optimal contact policy. If the termination condition is not satisfied, the line parameters are updated and the loop is run again.



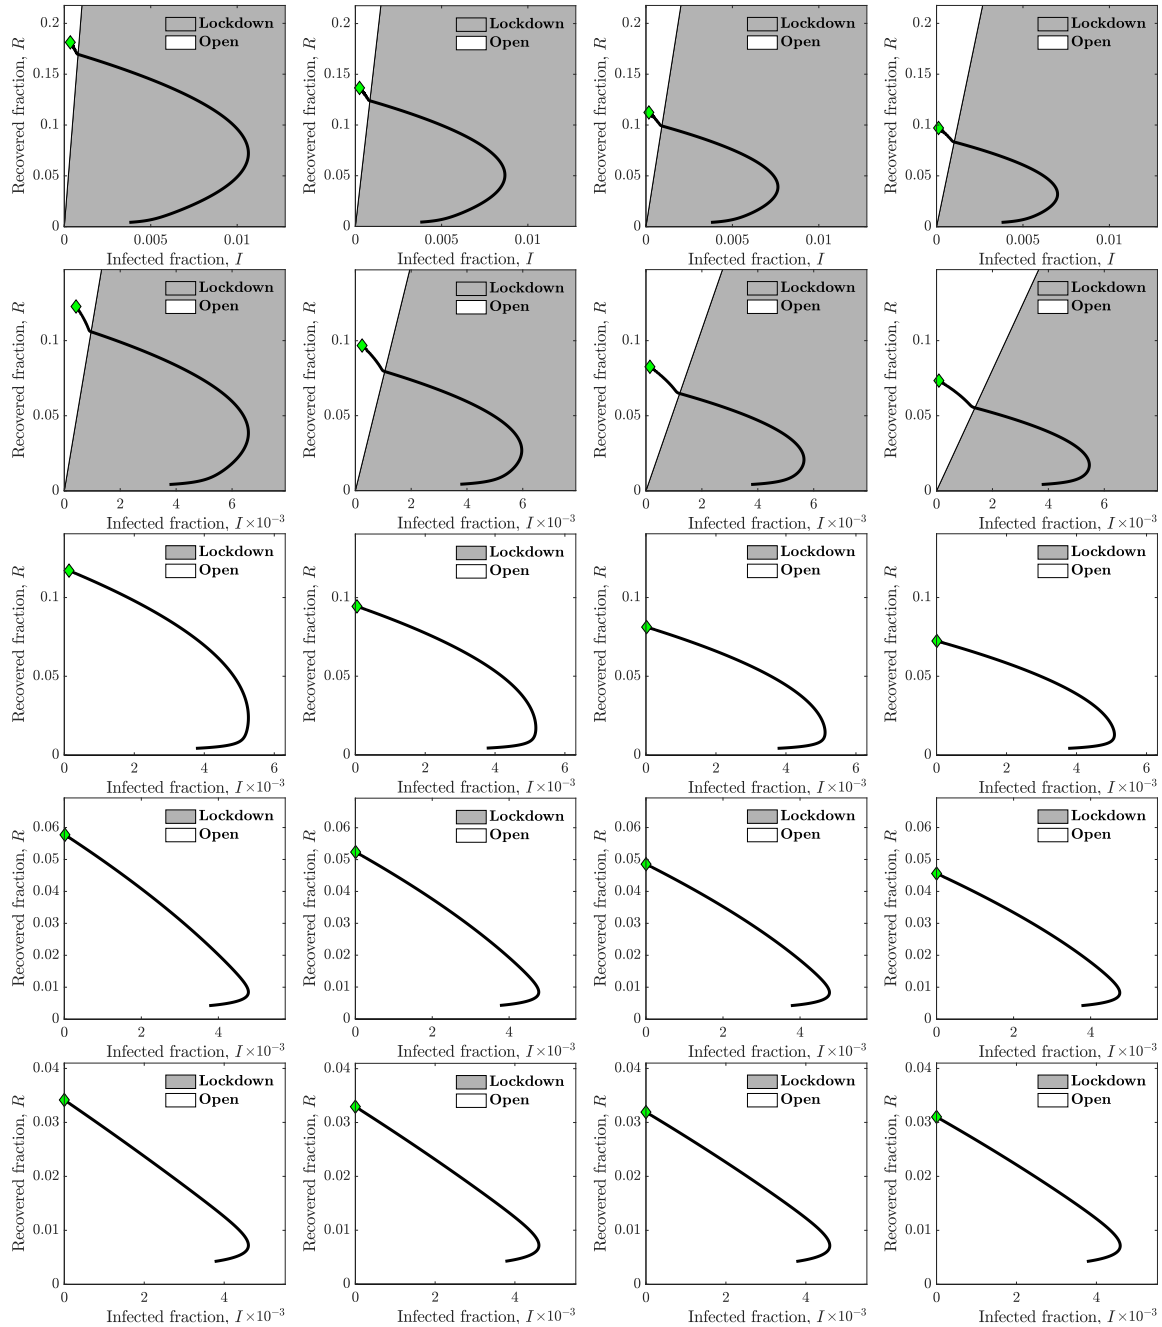


FIG. S9: Heuristic state feedback intervention policies varying with isolation efficiency and shielding level on $I - R$ phase plane. The isolation efficiency for the first row is 25% and each row increases the efficiency by 5%, which gives the efficiency of the last row as 75%. Next, the 4 columns correspond to shielding levels of 2, 3, 4 and 5 from left to right. In each plot, the grey region specifies the area of the $I - R$ plane where a lockdown policy is enforced for the susceptible population and the white region gives the area where the lockdown is lifted i.e. the system is 'open'. The phase trajectory ($I(t)$, $R(t)$), i.e., the black line, crosses from the lockdown to open region when it intersects the dividing line and this crossover happens at most once. The final state of the system is marked by a green diamond.