medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

Which lockdowns are the best lockdowns?

Optimal interventions and the trolley problem in heterogeneous populations

Pratyush K. Kollepara[1], Rebecca H. Chisholm[1], István Z. Kiss[2], and Joel C. Miller[1]
[1] *Department of Mathematical and Physical Sciences, La Trobe University, Melbourne, Australia and*
[2] *Department of Mathematics, University of Sussex, Brighton, United Kingdom*

Interventions to mitigate the spread of infectious diseases, while succeeding in their goal, have economic and social costs associated with them. These limit the duration and intensity of the interventions. We study a class of interventions which reduce the reproduction number and find the optimal strength of the intervention which minimises the attack rate of an immunity inducing infection. The intervention works by eliminating the overshoot part of an epidemic, and avoids a second-wave of infections. We extend the framework by considering a heterogeneous population and find that the optimal intervention can pose an ethical dilemma for decision and policy makers. This ethical dilemma is shown to be analogous to the trolley problem and we discuss how the dilemma can be avoided.

I. INTRODUCTION

Infectious disease epidemics have been suppressed and mitigated using a combination of non-pharmaceutical interventions (NPIs) such as lock downs, social distancing, mask wearing and contact tracing, and by immunising the population using vaccines. In the absence of vaccines, NPIs are the primary option. However, NPIs, and in particular, lockdowns, can have significant economic, mental health and social costs associated with them. Instead of protracted or repeated lock downs (as observed during the COVID-19 pandemic), a one-shot intervention has been suggested as a possible alternative. An intense but short-duration lockdown is imposed near the peak of the epidemic to stop the transmission during the overshoot phase of the epidemic and reduce the attack rate to the herd immunity threshold of the epidemic [1]. The overshoot phase is when the number of active infections start to decline (effective reproduction number is less than one), but a significant number of new infections are created. The overshoot is the difference between the attack rate and herd immunity threshold.

In this work, we explore an alternative strategy to achieve the same outcome through a prolonged but weaker intervention instead of a short and intense intervention. Such an intervention, if implemented early, will have the added benefit of reducing and delaying the peak of the epidemic as well, in contrast to the one-shot intervention [1]. If the economic and social costs of an intervention grow super-linearly with the strength of intervention, moderately intense but prolonged interventions may prove to be cost effective.

As with the one-shot intervention, the rationale of this strategy is to calibrate the intervention in such a manner that the attack rate of the mitigated epidemic is identical to the herd immunity threshold of the original epidemic. Therefore, when the intervention ends, there is no risk of further introductions developing into future epidemics. We show that this strategy is an optimal strategy for minimising the attack rate in the long term.

In the context of COVID-19 modelling, research on optimal interventions has attempted to include economic costs along with the objective of reducing infections: using detailed agent-based models [2] and fine-tuned intervention strategies [3], a balance is sought between socio-economic and health costs to minimise the total cost [4], or the claim that interventions reduce the economic well-being

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

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

of a society has been challenged [5, 6]. An approach similar to the one presented here was suggested for COVID-19 in European countries [7]. Optimal interventions have also been studied as resource allocation problems where a limited stockpile of vaccine is available or a limited 'amount' of social distancing is acceptable and the objective is to find distribution of the intervention that minimises the reproduction number or a health-related objective function [8–10].

We do not include economic costs in an explicit manner in our model. The amount of reduction in $\mathcal{R}_0$ can be interpreted as the cost – the higher the reduction in $\mathcal{R}_0$, higher the social and economic cost of intervention. The calculations involved in finding the optimal strategy mainly rely on the knowledge of the basic reproduction number (or the next generation matrix).

We show that implementing an optimal intervention to minimise the attack rate could involve a moral/ethical dilemma for decision makers, which is analogous to the commonly known trolley problem [11, 12]. The dilemma arises as a result of transmission heterogeneity in the population.

In the following sections, we explain the modelling framework, results of our analysis, and conclude with a discussion of our modelling assumptions and the ethical dilemma that decision-makers could face.

II. METHODS

We used deterministic SIR and SIR-like models to study the optimal intervention. In sub-section II A and II B, we explain the models used for a homogeneous population and for a heterogeneous population, respectively, in addition to describing the calculations for finding the optimal intervention. In II C we explain how an optimal intervention can be found if there is a delay in the start of the intervention.

A. Homogeneous population

We use an SIR model with the variables s, i and r to represent the fractions of individuals in the total population who are susceptible, infected and recovered respectively [13, 14]. The size of the population is assumed to be constant such that, $s + i + r = 1$.

In this case, the attack rate of the epidemic is completely determined by the basic reproduction number $\mathcal{R}_0$ and can be obtained using the following equation [13, 14]:

$$\ln \frac{s(t_2)}{s(t_1)} = -\mathcal{R}_0(r(t_2) - r(t_1)), \tag{1}$$

where $s(t_1)$ and $s(t_2)$ are the fractions of susceptible and $r(t_1)$ and $r(t_2)$ are the fractions of recovered individuals in the population at time instants, t_1 and t_2. Using the conditions $i(t_1) \approx 0$, $r(t_1) = 0$ and $i(t_2) \approx 0$, which describe the population at the start and end of an epidemic, the attack rate relation can be obtained

$$r(\infty) = 1 - e^{-\mathcal{R}_0 r(\infty)}. \tag{2}$$

An intervention that reduces transmission would affect the basic reproduction number as $\mathcal{R}_0 \to \mathcal{R}_0(1 - c)$ where $0 \leq c \leq 1$. In the case of a homogeneous population, the herd immunity is achieved when the fraction of susceptible individuals in the population is less than $\frac{1}{\mathcal{R}_0}$. Therefore, we substitute $s(t_2) = 1/\mathcal{R}_0$ and $s(t_1) = 1$ and solve for c. The optimal reduction in the basic

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

reproduction number is

$$c = 1 - \frac{\ln \mathcal{R}_0}{\mathcal{R}_0 - 1}. \tag{3}$$

We verify this analytical result in section III A by simulating an epidemic where $\mathcal{R}_0$ is changed to $\mathcal{R}_0(1-c)$ in the early stage of the epidemic and the intervention is switched off once the active infections ($i(t)$) decline to a negligible number.

B. Heterogeneous population

In a heterogeneous population, individuals may be further stratified into groups. To represent fraction of individuals in the total population who belong to a group k, we use the variables s_k, i_k and r_k such that $s_k + i_k + r_k = n_k$, where n_k is the proportion of the population who belong to group k and $\sum_k n_k = 1$. Heterogeneity in transmission characteristics can affect the behaviour of epidemics in a significant manner. Epidemics in populations with different transmission structures but identical reproduction numbers can have widely different attack rates. An epidemic in a heterogeneous population can be described by the following SIR-like model, assuming identical duration of infection for all groups and measuring time in the units of the average infection duration,

$$\frac{\mathrm{d}s_k}{\mathrm{d}t} = -s_k \sum_l B_{kl} i_l, \tag{4}$$

$$\frac{\mathrm{d}i_k}{\mathrm{d}t} = \left(s_k \sum_l B_{kl} i_l \right) - i_k, \tag{5}$$

$$\frac{\mathrm{d}r_k}{\mathrm{d}t} = i_k. \tag{6}$$

The term B_{kl} is the average number of infectious contacts that an individual in group l causes in group k. The next generation matrix, $\mathbf{G}$ [15], can be constructed for this system with entries:

$$G_{kl}(t) = s_k(t) B_{kl}. \tag{7}$$

The term G_{kl} is the expected number of infections caused in group k by an infected individual in group l. The dominant eigenvalue of $\mathbf{G}$ gives the reproduction number of the system [15]. The attack rates for this model are given by [16, 17]

$$r_k(\infty) = n_k(1 - e^{-\sum_l B_{kl} r_l(\infty)/n_l}). \tag{8}$$

It should be noted that we are defining the attack rate for the heterogeneous population as the number of infections in a group as a fraction of the total population.

The recipe for finding the optimal intervention is similar to the homogeneous case. Instead of a single reproduction number, we have the next generation matrix $\mathbf{G}$. The herd immunity threshold is computed by finding the proportion of susceptibles of each group, s_k, which would ensure that the reproduction number (top eigenvalue of $\mathbf{G}$) is equal to one. In this case, there could be (infinitely) many such states that would ensure herd immunity. From all the possible herd immunity states, we select the one which maximises $\sum_k s_k$ (or equivalently, minimises the attack rate). The optimal

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

intervention transforms $\mathbf{B}$ to $\mathbf{B}'$ such that $\mathbf{B}'$ results in the selected state as the attack rate.

A crucial point to note here is that unlike with the homogeneous population, it is possible for certain elements of $\mathbf{B}$ to increase in order to minimise the attack rate. In other words, the optimal intervention corresponds to an increase in transmission among certain groups or among pairs of groups. Thus, the change in reproduction number can not be a measure of of the economic or social cost in such cases. Nonetheless, this leads to some interesting results which are presented in the next section.

Instead of minimising the total number of infections, one can also consider a weighted objective function which is a weighted sum of the number of infections in each group. Such an objective function is useful when the disease affects different groups in different ways. The optimisation problem of finding the state of the population which minimises a general objective function and fulfils the herd immunity condition can be solved semi-analytically for the case of two groups and is presented in the Appendix.

C. Delayed intervention

In the above sections, we have assumed that the basic reproduction (or the next generation matrix) is a known entity and therefore an intervention is implemented right at the start of the epidemic. Calculating the strength of the optimal intervention requires knowledge of the reproduction number, the intervention would have to start after the epidemic has been established and enough observational data has been collected to calculate the reproduction number. While the basic principle would still hold, a delay could change the strength of the optimal intervention. To find the optimal strength for a delayed intervention, we use the attack rate relation with the final state $s = 1/\mathcal{R}_0,\ i = 0$ and an initial arbitrary state $s_L,\ i_L$ at a time instant t_L when the intervention begins. We replace basic reproduction number in equation (1) with $\mathcal{R}_0(1-c)$ and solve for

$$c = 1 - \frac{\ln s_L \mathcal{R}_0}{\mathcal{R}_0 \left(s_L + i_L - \frac{1}{\mathcal{R}_0} \right)}. \tag{9}$$

Using a numerical solution of the SIR equations, s_L and i_L can be found and the above equation can be solved for c. The equation (9) reduces to equation (3) when $s_L = 1,\ i_L = 0$, and $c = 1$ when $s_L = 1/\mathcal{R}_0$ and $i_L > 0$. If $s_L < \frac{1}{\mathcal{R}_0}$, $c > 1$ which is biologically meaningless and reflects the fact that the population is already below the herd immunity threshold.

III. RESULTS

A. Homogeneous population

Simulation of the SIR model differential equations confirms our assertion in equation (3). As shown in FIG. 1, a 'weak' intervention reduces the attack rate but does not reduce the overshoot to zero. A strong intervention, on the other hand, reduces the attack rate as long as the intervention is in place but a resurgence occurs as soon as the intervention ends. The final health outcome under the strong intervention is worse than (or at least comparable to) the weak intervention, while incurring a higher social and economic cost during the intervention. The resurgence occurs because the small number of infections and sufficient number of susceptibles remaining in the population

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

lead to new infections after the intervention is lifted. An intervention which is strong enough to minimise the final number of infections, while avoiding a resurgence, is the one whose attack rate (during the intervention) matches the herd immunity threshold of the unmitigated epidemic. This is the optimal intervention.

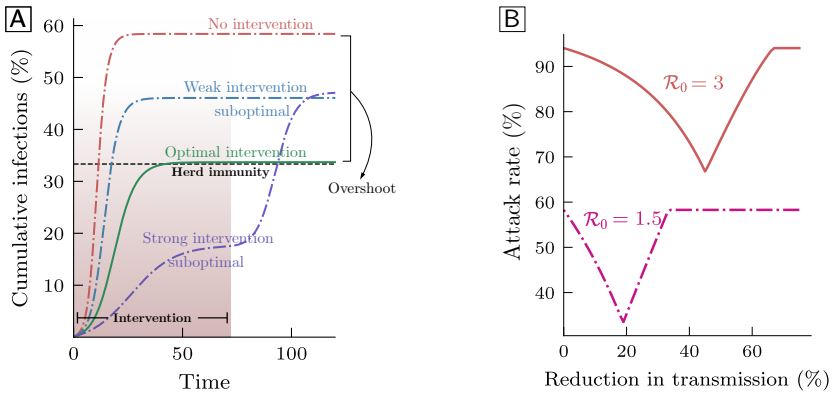

FIG. 1. **(A)** Simulation of four types of intervention for an epidemic with $\mathcal{R}_0 = 1.5$ in a homogeneous population: (i) No intervention – leads to largest attack rate, (ii) Weak intervention – reduces the attack rate, (iii) Strong intervention – reduces attack rate during the intervention, but leads to a resurgence in infections once the intervention is removed, (iv) A moderate intervention but optimal – attack rate is same as the herd immunity threshold. **(B)** The existence of a global minima for the attack rate shows that an optimal intervention strength exists. The resurgence of infections under a strong sub-optimal intervention is subject to certain assumptions which are discussed in section IV.

B. Heterogeneous population

Introducing heterogeneity in the model opens up a space of interventions that is not seen in the homogeneous case. In the homogeneous case, the herd immunity threshold is defined by a single point, but in the case of a structured population, the threshold is given by a collection of points. This can be seen by considering the following: the condition required for reaching herd immunity is that the typical infected individual must not infect more than one individual. In the homogeneous case, one can randomly chose a sufficient number of individuals and immunise them to ensure that the number of infectious contacts is less than one. If the population is structured, the typical infected individual must not infect more than one individual, on an average. As long as the average number of infectious contacts is less than one, herd immunity is achieved, irrespective of how the immunisation has been distributed among the various groups in the population. Thus there are infinitely many interventions that lead to herd immunity and prevent resurgence. Out of all these possibilities, we define the optimal intervention to be the one which minimises the attack rate.

When the population can be described using two sub-populations or groups, the optimal intervention belongs to one of the following types: the first group is fully infected, the second group is fully infected, or none of the groups are fully infected. This creates the possibility that under the optimal intervention, the number of infections in one of the groups is larger than what would have occured in the unmitigated epidemic, subject to the structure of the population. In FIG. 2 we show an example in which this occurs. This leads to an ethical dilemma wherein a certain group in the

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

population incurs a higher cost (due to an increased number of infections) in order to minimise the cost for the whole population. Thus, the non linearity of the infectious disease dynamics, combined with population structure, lead to an ethical dilemma for policy/decision-makers which is analogous to the well-known trolley problem (see FIG. 2). The trolley problem involves a setup in which a train is going to hit a group of people who are lying on the tracks. The train can not be stopped, but a lever can be pulled to switch the train onto a different track on which only one person is lying. The dilemma that is posed by this situation is whether it is ethical to save a few lives at the cost of one?

In our case of optimal intervention, we are dealing with infections rather than deaths. So the dilemma maybe avoided by considering the cost of an infection for each group. Diseases often lead to a worse health outcome (mortality rate, hospitalisation rate, chance of leading to chronic conditions etc.) in certain groups of the population (the elderly age groups for instance). Instead of minimising the attack rate of the epidemic (which is the sum of attack rates in each group), it may be more prudent to minimise a cost function which is a linear combination of the attack rates in the groups, such that a group with a worse outcome of infection is given a higher weight in the cost function. For the example shown in FIG. 2, the ethical dilemma can be avoided using such a weighted cost function, although this is not always the case.

C. Delayed intervention

We calculate the optimal strength of the intervention and simulate the model to confirm the mathematical analysis in Section II C. Using equation (9) we observe that the strength of optimal intervention increases in a super-linear manner with the duration of delay. The results are presented in FIG. 3. As the population approaches the herd-immunity threshold, the strength of intervention approaches one – corresponding to the one-shot intervention [1].

IV. DISCUSSION AND ONGOING WORK

In this work, we have examined a strategy of optimal interventions which uses the epidemic to cause just enough infections to induce herd immunity, eliminate the overshoot, and prevent future introductions from becoming epidemics. In addition to minimising the attack rate, this intervention would also slow down the growth of the epidemic and reduce the peak, which allows time to develop treatments and increase healthcare capacity.

In the case of heterogeneous transmission, our results indicate that the optimal strategy may require increasing the transmission in some of the groups and decreasing it in others, in order to minimise the attack rate for the whole population. This is analogous to the trolley problem, and it calls for a discussion around the ethics of subjecting certain groups to a higher rate of disease incidence, and the feasibility of this policy. If increasing transmission in certain groups is not viable either due to operational reasons or ethical considerations, herd immunity can still be achieved (and resurgence prevented) by reducing transmission in all groups. We have also presented some scenarios in which a weighted cost function avoids the dilemma. Our work shows that even without an explicit consideration of economic and social costs of an intervention, there are challenging questions to be answered for the first order problem of minimising the attack rate.

In addition to the ethical dilemma shown through our modelling here, interventions which require increasing transmission prompt an ethical discussion in relation to disadvantaged groups. Cultural, economic and social conditions factor into the contact structure of any human population – a high

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

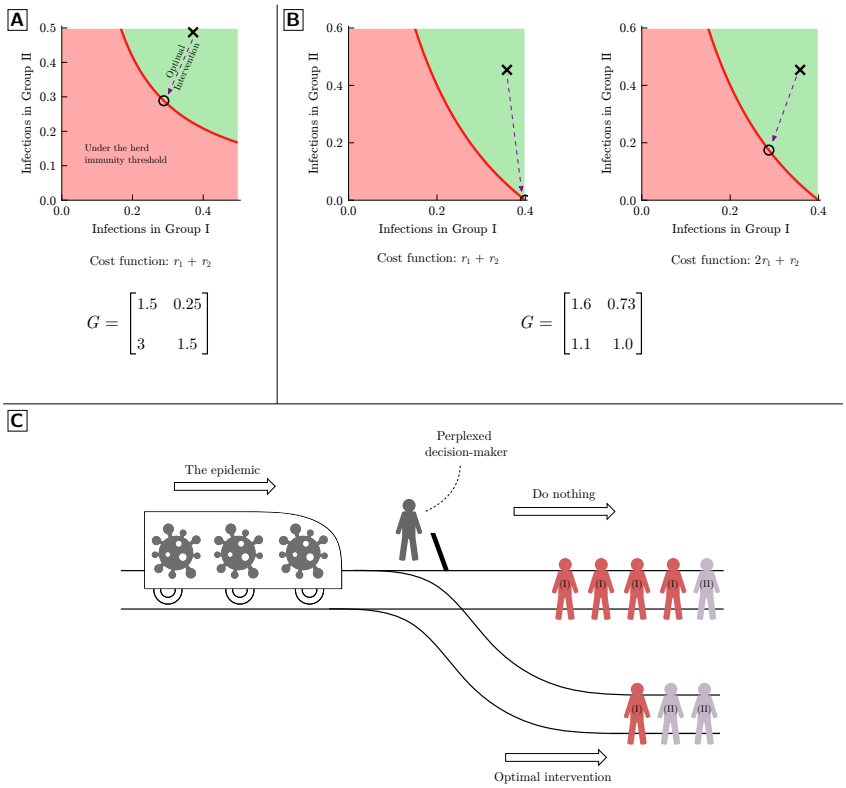

FIG. 2. Population structure and non linearity of infectious disease dynamics lead to an ethical dilemma. (**A**) A plot of the infections in the two groups that constitute the population. The corresponding cost function and next generation matrix are shown. The red curve shows the herd immunity threshold, the point marked with a cross shows the attack rate without an intervention, the point marked with a circle shows the attack rate if the optimal intervention is used. (**B**) A comparison of the interventions when the cost functions are different. The first plot shows that the optimal intervention leads to an increase in the number of infections in the first group. It is an example of the ethical dilemma of implementing the optimal intervention, which is explained further in (**C**). The second figure shows the plot for the cost function when the first group is given twice the weight as the second group, which means prevention of infections in the first group takes precedence over the second group. In this case, the intervention reduces the infections in both the groups. Thus, a biased cost-function may allow us to avoid the ethical dilemma. (**C**) This figure shows that the ethical dilemma involved in implementing optimal interventions is analogous to the well-known 'trolley problem'. If the decision-maker does not act, the incoming epidemic (represented as a trolley) is going to cause many infections. If the decision-maker implements an optimal intervention (switches the tracks), the number of infections in the total population is minimised, but someone who otherwise would have been safe from infection, becomes infected (shown by the increased number of infections of group (II)).

number of contacts due to living in close spaces, a high susceptibility to infection due to preexisting health conditions, or poor access to healthcare facilities, etc. Mathematical models of epidemics can throw light on possible choices of policy and may even help us pick the ones that lead to optimal outcomes. But the decisions made by policymakers are intertwined with political will, their popularity, and social attitudes. These eventually determine whether a particular intervention

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

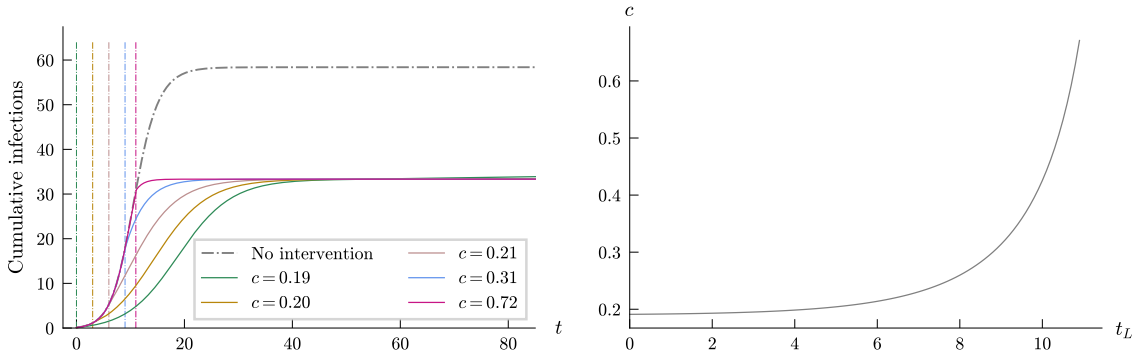

FIG. 3. **Left:** The time series of cumulative infections $(i + r)$ for optimally controlled epidemics for a range of delays in the intervention. The vertical dash-dot lines show the time at which the intervention starts, and its corresponding time series is represented by the same color. **Right:** The strength of optimal intervention c plotted against the delay in intervention t_L using equation (9). The super-linear increase in c shows the need for an early implementation. **Parameters:** Homogeneous SIR model with basic reproduction number $\mathcal{R}_0 = 1.5$ and $\gamma = 1$. Intervention is implemented for a duration of 50 time units.

is favoured by a decision-making body [18, 19]. Disadvantaged groups, across the world, do not exercise sufficient political power to represent their interests in decision making bodies. In such a case, a decision-making body may find it convenient to subject a disadvantaged group to a higher attack rate in order to decrease the net attack rate for the whole population and achieve herd immunity, even if it does not minimise the net attack rate. The intervention strategy presented here, always carries such risks with it; and representation of disadvantaged groups thus becomes essential.

In addition to the ethical concerns with the optimal interventions presented here, there are some practical limitations to this strategy. There would be a natural tendency for individuals to protect themselves from getting infected even if interventions are not in place, so asking individuals to increase their transmission may not be a feasible strategy [20]. The optimal interventions could require a group of individuals to fully isolate themselves from the rest of the population. Such interventions are difficult to implement, as there would always be a small possibility for infections to be introduced into the isolated group [20]. If the transmission in other groups is increased, it would imply a larger chance of introduction into the isolated group.

We have assumed an SIR structure for disease progression in an individual. But, as long as the disease can be reasonably described by a model in which individuals do not become susceptible after getting infected, we would expect our results to be valid. A crucial detail that we have ignored is the stochastic and discrete nature of disease spread since it can capture the elimination behaviour of outbreaks, i.e., it can incorporate the difference between existence and absence of infections. The deterministic assumption and the use of continuous variables in our model means that after an intervention is over, the small number of infections present in the population will lead to another epidemic if herd immunity is not achieved (shown in FIG. 1, strong intervention). This however, is one of the possible outcomes. It is possible that the intervention completely eliminates all infections in the population, in which case a new epidemic does not result from any residual infections. However, even in this case the population remains vulnerable to an epidemic due to lack of herd immunity. Thus, a new epidemic can occur if new infectious individuals are introduced into the population. Another possibility is that the epidemic may get established with a delay due to the stochastic

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

dynamics. Factors around contact-tracing and surveillance capacity (to eliminate the disease) and travel restrictions (to prevent introduction of new infections) are important for the selecting the optimal policy response, in addition to the results presented here.

The choice of two sub-populations was made to demonstrate the ethical dilemma arising out of the population structure, in the simplest manner. Generally, the population will be described by N sub-populations (or an $N \times N$ next generation matrix) with N increasing as the resolution of the information about the population increases. Therefore, describing an N-group structured population using an M-group model (where $M < N$) amounts to using a coarse-grained description of the system. We have shown that the optimal interventions in homogeneous population and in a structured population with two homogeneous sub-populations can be widely different. Thus, the optimal intervention can change if the amount of information about the structure changes. We intend to follow-up on this work by attempting to characterise how the optimal intervention changes as the scale of observation is changed, if the interventions could fail and if they can be robust.

APPENDIX: FINDING THE OPTIMAL INTERVENTION

The generalised cost function depends on the weight vector $\mathbf{a} = (a_1, a_2, \ldots)$ and the attack rate vector $\mathbf{r} = (r_1, r_2, \ldots)$ and is given by:

$$C(\mathbf{r}, \mathbf{a}) = \mathbf{a} \cdot \mathbf{r} = \sum_k a_k r_k. \tag{10}$$

$$\tag{11}$$

The vector $\mathbf{r}$ is constrained such that $\rho(\mathbf{G}) = 1$, where $\mathbf{G}$ is the next generation matrix and $\rho(\cdot)$ is the dominant eigenvalue.

We present a semi-analytical solution to this optimisation problem when there are two groups in the population. Minimising the cost function C is equivalent to maximising the function f :

$$f(s_1, s_2, a_1, a_2) = a_1 s_1 + a_2 s_2, \tag{12}$$

while the constraint is given by equating the eigenvalue of the next generation matrix to one

$$g(\mathbf{G}) = 1 - \operatorname{Tr} \mathbf{G} + \det \mathbf{G} = 0, \tag{13}$$

$$= 1 - (s_1 B_{11} + s_2 B_{22}) + s_1 s_2 \Delta = 0, \tag{14}$$

where $\Delta = \det \mathbf{B}$. Values of s_1 and s_2 that satisfy this equation may correspond to cases where the non-dominant eigenvalue is unity. Therefore, the dominant eigenvalue needs to be computed numerically in order to impose this restriction.

To solve this optimisation problem we define a new function $h = f + \lambda g$, where λ is a Lagrange multiplier. The partial derivatives of h with respect to $s_1,\ s_2$ and λ will be zero at points of interest

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

TABLE I. List of possible extremal points for a two-group population. The first column shows the solution type referred to in FIG. 4. The second and third column give the expression for the extremal values of s_1 and s_2. The quantities B_{kl} are the elements of the transmission matrix, $\{a_1, a_2\}$ are the cost function weights. The quantity $\Delta = B_{11}B_{22} - B_{12}B_{21}$.

Type	s_1^*	s_2^*	Description
0	0	$\frac{1}{B_{22}}$	Boundary point: Group 1 fully infected
1	n_1	$\frac{1-n_1B_{11}}{B_{22}-n_1\Delta}$	Boundary point: Group 1 fully protected
2	$\frac{1}{B_{11}}$	0	Boundary point: Group 2 fully infected
3	$\frac{1-n_2B_{22}}{B_{11}-n_2\Delta}$	n_2	Boundary point: Group 2 fully protected
4	$\frac{B_{22}+\sqrt{B_{12}B_{21}a_2/a_1}}{\Delta}$	$\frac{B_{11}+\sqrt{B_{12}B_{21}a_1/a_2}}{\Delta}$	Local extrema
5	$\frac{B_{22}-\sqrt{B_{12}B_{21}a_2/a_1}}{\Delta}$	$\frac{B_{11}-\sqrt{B_{12}B_{21}a_1/a_2}}{\Delta}$	Local extrema

and the three equations will be solved simultaneously to obtain

$$\frac{1}{\lambda} = \pm\sqrt{\frac{B_{12}B_{21}}{a_1a_2}}, \tag{15}$$

$$s_1^* = \frac{B_{22} \pm \sqrt{\frac{B_{12}B_{21}a_2}{a_1}}}{\Delta}, \tag{16}$$

$$s_2^* = \frac{B_{11} \pm \sqrt{\frac{B_{12}B_{21}a_1}{a_2}}}{\Delta}. \tag{17}$$

For a given matrix $\mathbf{B}$, group sizes n_1 and n_2, and costs a_1 and a_2, the function f will be evaluated for the two points given by (s_1^*, s_2^*) and the four boundary points. The six 'types' of possible extrema for the optimisation problem, are shown in TABLE I and in FIG. 4. Depending on the parameters, one of these points will maximise the function f subject to the constraint $g = 0$.

The boundary extremal points represent scenarios where one of the groups is either fully infected or is fully protected from getting infected. Due to the large number (six) of parameters governing the optimal solution, we present graphical representations of the outcomes under the optimal intervention, in each group for sections of the parameter space in FIG. 4. We use a transmission matrix with a symmetric structure $B_{kl} = b_k b_l((1-\alpha)\delta_{kl} + \alpha)$ where α can be used to control the relative magnitudes of the diagonal and off-diagonal elements. The changes in the type of solution occur primarily based on the validity of the extremal points ($r_1^* \in [0, n_1]$ and $r_2^* \in [0, n_2]$) and which of the valid solutions minimise the cost. As an example, the straight line separation between type 5 (local minima) and type 2 (boundary point) and between type 5 and type 0 (boundary point), in the first row of FIG. 4, can be explained by imposing $s_1^* \geq 0$ and $s_2^* \geq 0$ on the expressions listed in TABLE I

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

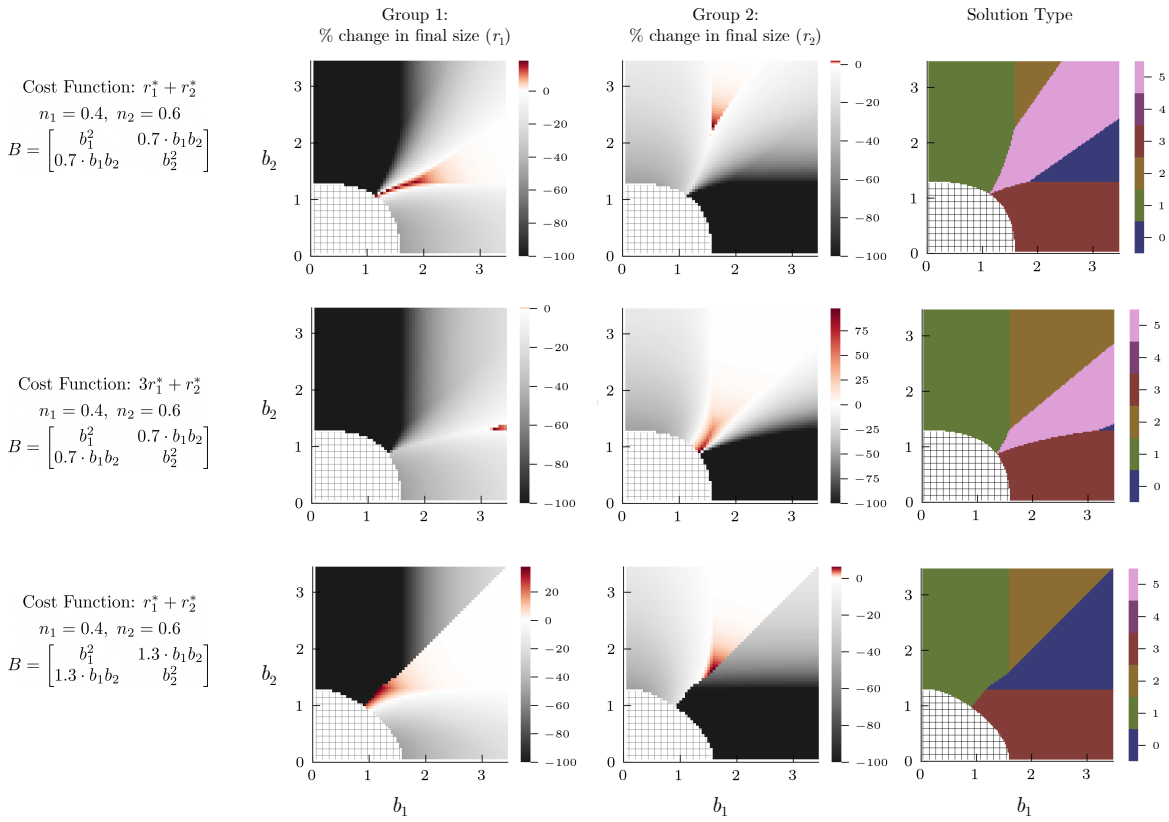

FIG. 4. A section of the parameter and optimal outcome space: The first column of each row shows the cost function, group sizes and the structure of the transmission matrix B. The next two columns are heat-maps which show the percentage change in the attack rate of Group 1 and 2 if the optimal intervention is implemented. The last column shows the type of solution obtained from the optimisation problem – the mathematical expression and description for the solutions can be looked up in TABLE I. The red coloring in the heat-map show regions of the parameter space where the optimal intervention would cause a worse outcome in one of the groups while minimising the cost function for the overall population. The model is identical for the first two rows, but the cost function is different. In the first row the cost function is the attack rate, while in the second row it is a weighted function of the two attack rates. The group sizes are 0.4 and 0.6 for all the rows.

V. ACKNOWLEDGEMENTS

We thank Michael Meehan, Sarah Becirevic and Somya Mehra for helpful discussions. This work was supported by a La Trobe University Graduate Research Scholarship and a La Trobe University Full Fee Research Scholarship for Pratyush Kollepara, and startup funding from La Trobe University for Joel Miller.

medRxiv preprint doi: https://doi.org/10.1101/2023.02.18.23286135; this version posted February 23, 2023. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted medRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC 4.0 International license .

VI. SOFTWARE REPOSITORY

The software used in this paper is available at `https://github.com/Joel-Miller-Lab/optimal-intervention`.

REFERENCES

[1] Di Lauro, F., Kiss, I. Z. & Miller, J. C. Optimal timing of one-shot interventions for epidemic control. *PLOS Computational Biology* **17**, e1008763 (2021).

[2] Blakely, T. *et al.* Determining the optimal COVID-19 policy response using agent-based modelling linked to health and cost modelling: Case study for Victoria, Australia. Preprint, Epidemiology (2021).

[3] Britton, T. & Leskelä, L. Optimal intervention strategies for minimizing total incidence during an epidemic. Preprint, Infectious Diseases (except HIV/AIDS) (2022).

[4] Samuel, J. & Sinha, S. Optimal control in pandemics. *Physical Review E* **103**, L010301 (2021).

[5] Dorn, F. *et al.* The common interests of health protection and the economy: Evidence from scenario calculations of COVID-19 containment policies. *The European Journal of Health Economics* (2022).

[6] Siegenfeld, A. & Bar-Yam, Y. Minimizing economic costs for COVID-19 (2020).

[7] Ragonnet, R. *et al.* Optimising social mixing strategies to mitigate the impact of COVID-19 in six European countries: A mathematical modelling study. Preprint, Infectious Diseases (except HIV/AIDS) (2020).

[8] Wallinga, J., van Boven, M. & Lipsitch, M. Optimizing infectious disease interventions during an emerging epidemic. *Proceedings of the National Academy of Sciences* **107**, 923–928 (2010).

[9] Goldstein, E. *et al.* Distribution of vaccine/antivirals and the 'least spread line' in a stratified population. *Journal of The Royal Society Interface* **7**, 755–764 (2010).

[10] Abell, I. R., McCaw, J. M. & Baker, C. M. Understanding the impact of disease and vaccine mechanisms on the importance of optimal vaccine allocation. Preprint, Infectious Diseases (except HIV/AIDS) (2022).

[11] Foot, P. *Virtues and Vices* (Oxford University Press, 2002).

[12] Thomson, J. J. Killing, Letting Die, and the Trolley Problem. *The Monist* **59**, 204–217 (1976).

[13] Kermack, W. O. & McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character* **115**, 700–721 (1927).

[14] Hethcote, H. W. The Mathematics of Infectious Diseases. *SIAM Review* **42**, 599–653 (2000).

[15] Diekmann, O., Heesterbeek, J. & Metz, J. On the definition and the computation of the basic reproduction ratio R 0 in models for infectious diseases in heterogeneous populations. *Journal of Mathematical Biology* **28** (1990).

[16] Miller, J. C. A Note on the Derivation of Epidemic Final Sizes. *Bulletin of Mathematical Biology* **74**, 2125–2141 (2012).

[17] Andreasen, V. The Final Size of an Epidemic and Its Relation to the Basic Reproduction Number. *Bulletin of Mathematical Biology* **73**, 2305–2321 (2011).

[18] Palus, S. America Never Did a Real Lockdown. *Slate* (2020).

[19] Yerushalmy, J. Zero-Covid policy: Why is China still having severe lockdowns? *The Guardian* (2022).

[20] Smith, C. A., Yates, C. A. & Ashby, B. Critical weaknesses in shielding strategies for COVID-19. *PLOS Global Public Health* **2**, e0000298 (2022).