

RESEARCH ARTICLE

Rational social distancing policy during epidemics with limited healthcare capacity

Simon K. Schnyder^{1*}, John J. Molina², Ryoichi Yamamoto², Matthew S. Turner^{3,4*}

1 Institute of Industrial Science, The University of Tokyo, 4-6-1 Komaba, Meguro, Tokyo, Japan, **2** Department of Chemical Engineering, Kyoto University, Kyoto, Japan, **3** Department of Physics, University of Warwick, Coventry, United Kingdom, **4** Institute for Global Pandemic Planning, University of Warwick, Coventry, United Kingdom

* skschnyder@gmail.com (SKS); m.s.turner@warwick.ac.uk (MST)



Abstract

Epidemics of infectious diseases posing a serious risk to human health have occurred throughout history. During recent epidemics there has been much debate about policy, including how and when to impose restrictions on behaviour. Policymakers must balance a complex spectrum of objectives, suggesting a need for quantitative tools. Whether health services might be ‘overwhelmed’ has emerged as a key consideration. Here we show how costly interventions, such as taxes or subsidies on behaviour, can be used to exactly align individuals’ decision making with government preferences even when these are not aligned. In order to achieve this, we develop a nested optimisation algorithm of both the government intervention strategy and the resulting equilibrium behaviour of individuals. We focus on a situation in which the capacity of the healthcare system to treat patients is limited and identify conditions under which the disease dynamics respect the capacity limit. We find an extremely sharp drop in peak infections at a critical maximum infection cost in the government’s objective function. This is in marked contrast to the gradual reduction of infections if individuals make decisions without government intervention. We find optimal interventions vary less strongly in time when interventions are costly to the government and that the critical cost of the policy switch depends on how costly interventions are.

OPEN ACCESS

Citation: Schnyder SK, Molina JJ, Yamamoto R, Turner MS (2023) Rational social distancing policy during epidemics with limited healthcare capacity. PLoS Comput Biol 19(10): e1011533. <https://doi.org/10.1371/journal.pcbi.1011533>

Editor: Christian Hilbe, Max Planck Institute for Evolutionary Biology: Max-Planck-Institut für Evolutionsbiologie, GERMANY

Received: April 6, 2023

Accepted: September 20, 2023

Published: October 16, 2023

Copyright: © 2023 Schnyder et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The authors confirm that all data underlying the findings are fully available without restriction. All relevant data are within the paper and its [Supporting Information](#) files.

Funding: This work was supported by the Grants-in-Aid for Scientific Research (JSPS KAKENHI) under Grants No. 20H00129 (RY), 20K03786 (JJM), 20H05619 (RY), 22H04841 (SKS), 22K14012 (SKS), 23H04508 (JJM), the JSPS Core-to-Core Program “Advanced core-to-core

Author summary

The question of how to determine policies during epidemics is a subject of broad contemporary interest. How and when should society impose behavioural restrictions in order to reduce infections? Policymakers have to balance many objectives, suggesting a need for quantitative tools for designing optimal intervention policy. Previous work on optimal policy-making typically sidesteps the question of how the population follows any intervention chosen by a government. Furthermore, the cost of implementing an intervention is also usually ignored. Our work overcomes these shortcomings. We analyse how the population chooses behaviour in a self-organised way. This can be influenced by the government so as to optimise its own objective function. Its objectives can be different from the individuals. Our work represents a proof-of-principle that costly policy interventions can

network for the physics of self-organizing active matter" JPJSCCA20230002 (all of us), and the SPIRITS 2020 grant of Kyoto University (JJM). MST acknowledges the generous support of visiting fellowships from JSPS Fellowship, ID L19547, and the Leverhulme Trust, Ref. IAF-2019-019. Websites: JSPS: <https://www.jsps.go.jp/english/> Leverhulme Trust: <https://www.leverhulme.ac.uk/> The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

be developed in the environment of (i) a disease with well-understood epidemiological character (ii) health-care capacity limits (iii) where those interventions are costly to implement. For these reasons we believe it highlights an opportunity to develop deployable policymaking tools and further advances our understanding of epidemiology when individuals adapt their behaviour in response to perceived dangers.

Introduction

Policymakers can manage epidemics using a variety of non-clinical interventions that target behaviour and hence the rate at which the disease is passed on. At one extreme this can involve merely providing accurate information and/or conceptual tools to enable rational individuals to identify their optimum behaviour. More interventionist strategies available to policymakers include subsidising preferred behaviour and/or penalising behaviour that they wish to discourage. Recent epidemics have generated much debate about policy, including how and when to impose restrictions on behaviour. Policy is likely to fall sharply into focus as the epidemic is analysed in a historical context, informing our planning for future epidemics. The primary goal of this work is to establish a proof-of-principle that fully quantitative approaches can be used to help design optimal intervention strategies, first in a stylised model but without obvious conceptual limits to incorporating more faithful descriptions of population composition and behaviour. We show that the costs of government interventions can be incorporated into the kinds of quantitative tools that would be necessary to manage future pandemics.

Dealing with an epidemic as a policymaker requires a number of objectives to be prioritised and balanced. The goal of limiting infections may justify restrictions on the day-to-day social and economic activities of citizens or subjects. A *rational* policy design process involves policymakers who are aware of the strategies that provide the most beneficial outcomes, these being evaluated using quantitative metrics. Our motivation here is to further the development of such quantitative tools. Ultimately we would see this as an aid to policy making but here we are concerned with establishing a point of principle—that it is possible to target outcomes that are optimal in the sense that they maximise an objective function that balances costs against benefits in the specific case of (i) when these interventions can carry costs to the government, (ii) when the healthcare system has limited capacity, (iii) when the interventions (have to) take into account the endogenous behaviour of the population, including its response to said interventions, and (iv) when the interests of government and population might not be aligned.

This study is concerned with rational policymaking in and for a society of rational individuals. There already exists a literature that explores the behaviour of rational *individuals*, in the absence of policy interventions. These individuals are typically assumed to be able to adjust their behaviour in the face of an epidemic [1–9]. Broadly speaking, individuals may choose to limit their social activity when infections are high, to avoid the risk of becoming infected themselves, provided that the health risks outweigh the social and economic costs. In the opposite limit, little or no behavioural changes are made and the epidemic is assumed to run its natural course much as if the agents were unreasoning. These studies are highly stylised in several respects, including the use of population-wide mean-field compartmentalised model and little or no analysis of the role of uncertainties. While they have not yet been developed into the more sophisticated variants needed to reorientate towards real data they nonetheless lay down an important milestone in demonstrating that such analysis is possible,

at least in principle. It is generally straightforward to see how such approaches can be extended to incorporate the complexities of real data, mirroring the sophistication of epidemiological approaches that incorporate more realistic household-level descriptions. This might include multiple compartment types with different risk and behaviour profiles [2, 10–13], spatial [14] and temporal networks [15, 16], seasonal effects [17], spatial or transmission heterogeneity [11, 18–20] or agent-based models [21–24]. It is also possible to include noise, for instance in the control [25]. It is also of interest to study the inverse problem to ours where one attempts to infer the objective function underlying some observed (social distancing) behaviour [26].

Perhaps the most fundamental common assumption is that individual agents act rationally, i.e. to maximise an economic utility. Although the limitations of such approaches have been widely acknowledged, e.g. within behavioural economics generally [27], this remains one of the fundamental assumptions of modern economic theory and will be adopted in the present work, noting that conceptual tools could be provided to assist individuals in identifying rational decisions. Recent methodological advances have allowed to establish the behaviour of individuals that target a Nash equilibrium, rather than a global utility maximum that requires coordination [1, 2, 9, 28, 29].

Different from such *decentralised* decision-making, governments present an instance of *centralised* decision-making. These will typically not aim for Nash equilibria but for policy that is more socially optimal or better aligned with political or national priorities [7, 30–32]. Furthermore, subsidy and tax schemes can be used by a social planner to *decentralize optimal policy*, i.e. to bring the Nash equilibrium of individuals into alignment with the global optimum [30, 33, 34]. These approaches have so far only been applied to the special case where the subsidy and tax schemes are cost-free to the social planner, i.e. they appear only in the utility function of the individuals. Additionally, attention has been restricted to the case where the preferences of the government and the population are well aligned. We go beyond these restrictions by invoking a hierarchy of interests. This requires a nested optimisation of both the government intervention strategy and the underlying equilibrium behaviour of the population. An important aspect of our work is that we investigate the situation in which the cost of an infection relative to the cost of social distancing can be quite different for the government than for an individual. This is highly plausible, as for instance, it is likely to be more difficult for an individual to negotiate the right to work remotely than were the government to impose these arrangements.

Typical government interventions in this literature would involve taxing high social activity of infectious individuals with the aim of disincentivising them from certain behaviour, akin to a Pigouvian tax or subsidy [30, 35]. The collected taxes get redistributed equally over the whole population. The typical assumption is then that the intervention has no direct effect on the government's objective function since the *process* of redistribution is assumed to be cost-free, while of course the results of the taxation, here the reduction of social activity, do impact the government objective function indirectly. However, we argue that one must consider the *process* of redistribution itself as costly, e.g. due to the misallocation of resources and the distortion of markets caused by the collection of taxes, an effect known as the shadow or marginal cost of public funds [36]. Another factor could be that the administration of the incentivisation process is in itself costly, e.g. it requires clerical and professional resources, surveillance resources, etc.

Some recent studies have also focussed on the role of healthcare thresholds [29, 37–41], but not in combination with Nash equilibrium behaviour and costly government interventions. Ref. [29] is most similar to ours, investigating the role of government intervention on equilibrium behaviour in a situation where the case fatality rate depends on the current number of

cases. Their work differs from ours in that they study equations that are discrete in time, and that their case fatality rate is unbounded for large infection numbers. Most importantly, they are only interested in the case in which the government and the individuals have the same preferences and that the intervention is cost-free.

SIR models being compartmental models with continuous values, it is impossible to fully eradicate the disease, at best reaching an exponential decay of infections with strong social distancing or after reaching herd immunity via infections or vaccination. While eradication can in principle be incorporated, e.g. by defining a critical value of the infectious compartment below which the disease is said to have been eradicated [39], eradication is quite complicated to reach in a global pandemic in practice. This is why we choose to neglect the possibility of complete eradication in what follows.

Waves of infections are predicted to occur under certain circumstances, e.g. when fresh variants occur that (partially) escape immunity [42], waning immunity and demographics [43], or when social distancing is a more *ad hoc* response to recent changes in the infection and fatality numbers [44].

We focus on calculating the self-organised social distancing of individuals and the government incentives that enable such behaviour. We do not investigate other possible policy interventions such as vaccination and treatment strategies, [3, 4, 9, 22, 24, 31, 45–53], or isolation, testing, and active case-tracing strategies [39, 54], noting that these can be included in future variants of models like the one we analyse here. Instead, we assume that a vaccine becomes available at a time far longer than the duration of the epidemic, at which point all the remaining susceptible people become immune to the disease instantly. We do this so that we only have to study the behaviour on a finite time horizon. We ignore the situation where a vaccine becomes available during the epidemic. While the early arrival of a vaccine would have consequences for both equilibrium and globally optimal behaviour [1, 7, 29, 55], this lies outside of the scope of this work. Judging from previous work, one would roughly expect that the earlier the vaccine is expected to arrive, the more incentivised both individuals and governments would be to increase their social distancing efforts.

In what follows policymakers are also assumed to be acting rationally. They decide how to intervene so as to maximise a government-level objective function. In the spirit of a proof of principle we limit policy priorities to three of the most obvious factors: reducing direct health risks, avoiding excessive stress on the health care system and mitigating the social and financial impact associated with placing limits on individual behaviour. The primary variables are: (1) the infectiousness $k(t)$, parameterising the mean number of additional cases a single infected individual would cause in a previously unexposed population. This is assumed to have a background, or natural, level $\kappa^* > 1$ adopted by society in the absence of any behavioural changes, also known as the basic reproduction number R_0 . (2) A time-dependent government intervention $\epsilon(t)$ that can be deployed to incentivise behavioural changes in individuals. For simplicity we neglect the possibility of reinfection, although the present framework can be modified to incorporate this.

Methods

Epidemic dynamics

The epidemic dynamics represent the lowest hierarchy in our problem, see Fig 1, and inform all rational decisions made by the population and policy makers. We assume that the epidemic follows a standard SIR compartmentalised model [56] in which the fraction of the population in the susceptible, infected and recovered categories, the latter including any fatalities, obey

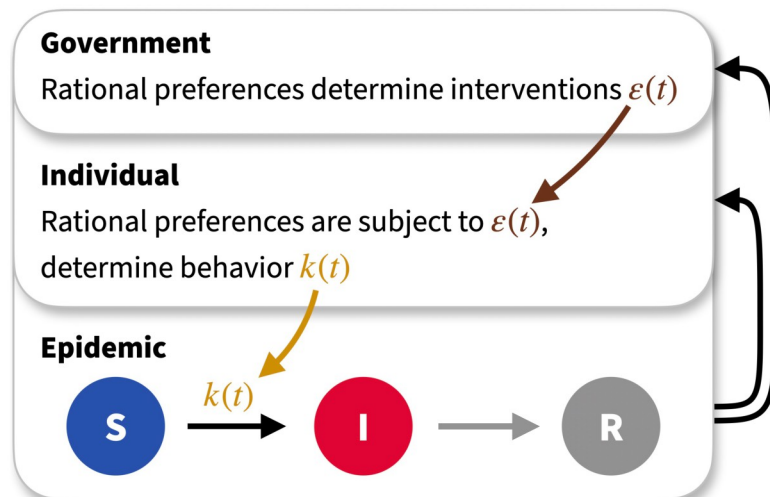


Fig 1. Causal hierarchy of the model. Epidemic dynamics are modelled using a simple Susceptible-Infected-Recovered compartmental model. This informs all decision making (black arrows). The progress of the disease depends only on the behaviour of individuals, who adopt a behaviour consistent with an infectiousness $k(t)$ at time t (gold arrow). Individuals may receive government incentives $\varepsilon(t)$ (brown arrow) to modify their behaviour. They then adopt a rational strategy $k(t)$, corresponding to a Nash equilibrium, based on some utility functional. The government maximises its own value functional and intervenes with incentives for individuals to realise this. This intervention process will, in general, itself carry costs.

<https://doi.org/10.1371/journal.pcbi.1011533.g001>

the rescaled equations

$$\begin{aligned}\dot{s} &= -k s i \\ \dot{i} &= k s i - i \\ \dot{r} &= i\end{aligned}\tag{1}$$

with initial values $s(0) = 1 - i_0$ and $i(0) = i_0$ at a time $t = 0$. We usually drop most functional dependencies, such as time here, for brevity. Here a dot denotes a time derivative and we have assumed a single timescale for recovery and the duration of infectiousness, for simplicity, measuring time t in these units. The course of the epidemic depends on the population averaged infectiousness $k(t)$, which arises from the behaviour of the whole population; as a shorthand, we directly denote $k(t)$ as behaviour. Social distancing performed by the population results in a reduction of k . At this level of the hierarchy, we take $k(t)$ as given, but we will calculate it self-consistently from individual behaviour in the next section.

Since the following results do not depend on the recovered fraction of the population, we omit it in what follows. The solution of these equations is shown for constant $k = \kappa^* = 4$ in Fig 2 as a baseline for comparison to various scenarios with behavioural modification of k . For this, we calculated the numerical solution of Eq 1 with a standard ordinary differential equation solver implemented in the `integrate.odeint` function in the `scipy` Python package [57].

Nash equilibrium behaviour

In the following we calculate the expected population behaviour if the population seeks out a Nash equilibrium. Conceptually, we are formulating a mean-field game [58, 59], which can be solved with a standard optimal control theory approach [60]. Here, we are building on the work of Reluga TC and Galvani AP [9].

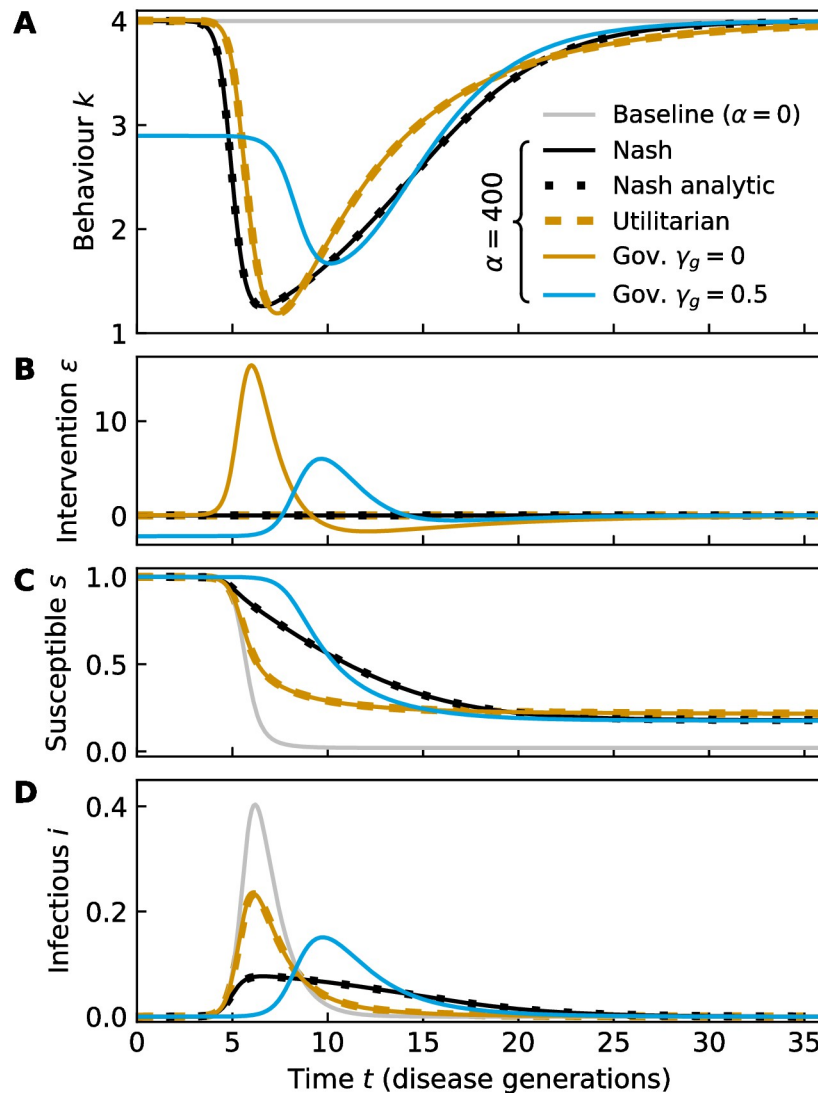


Fig 2. Comparison of social distancing behaviour. (A) Population behaviour $k(t)$, (B) government intervention $\varepsilon(t)$ and (C, D) dynamics of the disease s , i for a range of scenarios with $i_0 = 3 \cdot 10^{-8}$, $f = 1$, and $\kappa^* = 4$ throughout: a baseline where there is no behavioural modification (corresponding to equilibrium behaviour at an infection cost $\alpha = 0$, grey lines); the Nash equilibrium for $\alpha = 400$ (black lines), calculated numerically via forward-backward sweep, see section C in S1 Text (In order to demonstrate that the numerical solution is accurate, we also show the analytical solution of the same equations [62] as black dots); the utilitarian maximum for $\alpha = 400$ (gold dashes); and finally the population behaviour for two optimal government policies, one being without cost to the government, $\gamma_g = 0$ (gold lines), and one being costly, $\gamma_g = 0.5$, with $\alpha = 400$ (cyan lines). When government interventions are cost-free, they enable the population to reach the utilitarian maximum.

<https://doi.org/10.1371/journal.pcbi.1011533.g002>

Any representative individual of the population is assumed to observe the course of the epidemic in the population, and select their behaviour κ in response to it. The mean-field nature of the approach implies that the individual does not observe the behaviour of other individuals but only the averaged dynamics as described by Eq 1. The individual at any given time is either susceptible, infectious, or recovered, and their fate can be modelled as a series of discrete transition events between these states. In order to make the situation tractable, we calculate the expected probability $\psi_j(t)$ that the individual is in compartment j at time t as a continuous

time Markov process. In direct analogy to the compartmental model for the epidemic in the population, we can write

$$\begin{aligned}\dot{\psi}_s &= -\kappa\psi_s i \\ \dot{\psi}_i &= \kappa\psi_s i - \psi_i\end{aligned}\quad (2)$$

with initial values $\psi_s(0) = s(0)$ and $\psi_i(0) = i(0)$. These equations are similar to Eq 1 but involve the infected fraction of the population reservoir i , itself a solution to those equations. The equations describe how susceptible individuals become infected by coming in contact with members of the infectious compartment of the wider population. If the individual becomes infected their behaviour is assumed not to affect the course of the epidemic itself. Reducing $\kappa(t)$ has the effect of directly reducing the rate of change of ψ_s , i.e. increases the probability of remaining susceptible and lowers the probability of becoming infectious.

Alternatively, one can interpret these equations as a compartmental model for course of the epidemic in a small group of individuals, small enough compared to the whole population so as not to affect the course of the epidemic itself, being able to employ a different strategy $\kappa(t)$ as compared to the population-averaged strategy $k(t)$. One also has to assume that the individuals are dispersed in the population and cannot infect each other, only becoming infected by coming into contact with the rest of the population.

The individual knows exactly how many susceptibles, infected and recovered there are in the population, but the individual does not have any information about which group any given person belongs to. As a result, the individual cannot selectively socially distance, i.e. only distance from infected. We require everybody to socially distance.”

According to expected utility theory the individual will seek to maximise a utility functional which depends on both their own and the population behaviour, $U(\kappa(t), k(t))$. Any given individual cannot influence the behaviour of the whole population, so from the viewpoint of the individuals $k(t)$, and as a result $s(t)$ and $i(t)$ represent external or exogenous quantities, to which the individual can merely react with their own behaviour $\kappa(t)$. For the individual, the situation can be represented as a standard optimal control problem.

A Nash equilibrium for a population of identical individuals is found when one identifies a strategy $\kappa(t)$ for which, when adopted by the general population, individuals cannot find an alternative strategy $\tilde{\kappa}(t)$ that improves their utility

$$U(\tilde{\kappa}(t), \kappa(t)) \leq U(\kappa(t), \kappa(t)), \text{ for any } \tilde{\kappa}(t). \quad (3)$$

In such a situation, any given individual would be expected to react to the population strategy κ by selecting behaviour κ themselves, thus upholding the population strategy self-consistently.

The strategy to obtain explicit solutions, is to maximise $U(\kappa, k)$ over κ , treating k as exogenous. Having identified this extremum, one sets $k = \kappa$ to obtain the Nash equilibrium strategy adopted by the entire population. In more detail:

We analyse a simple stylised form for the individual utility with discounted utility per time u

$$U = \int_0^\infty u(t) dt \quad (4)$$

$$u = f^{-t} [-\alpha(i) \psi_i - \beta (\kappa - \kappa^*)^2 + (\kappa - \kappa^*) \varepsilon(t)] \quad (5)$$

Here $f \geq 1$ is the individual's discount rate (equivalent to a discount time $1/\log f$). The cost associated with infection, including the risk of death, is written $\alpha(i)$. This can reflect escalating

costs when a healthcare threshold is exceeded, e.g. as hospitals become full and as a result average treatment quality deteriorates and fatality rates increase. For simplicity we neglect the queueing process which determines whether an individual still receives state of the art healthcare such as admission to an intensive care unit with access to ventilators, etc. Instead, we assume simply that the more infectious there are on the population level, the worse on average the treatment of an individual becomes and the higher the probability of dying. Therefore the cost per single infection $\alpha(i)$ in general depends on the number of infectious i . We study two situations. One situation is characterised by the cost of an infection being always the same, i.e. $\alpha(i) = \text{const}$. The other situation represents the fact that healthcare systems have limited capacity by having the infection cost rise near a healthcare threshold i_{hc}

$$\alpha(i) = \alpha_0 + \frac{\alpha_1 - \alpha_0}{2} (\tanh[(i - i_{hc})\sigma] + 1) \quad (6)$$

with minimum cost α_0 , maximum cost α_1 and a steepness σ , see Fig 3A. If, during the course of the epidemic, the fraction of infectious i approaches the threshold, the cost to being infectious increases. This reflects the greater damage from becoming infected when healthcare resources are saturated, as well as an increased likelihood of death.

The constant β parameterises the financial and social costs associated with an individual modifying their behaviour from the baseline infectivity κ^* . Our choice of a quadratic form here ensures a natural equilibrium at $\kappa = \kappa^*$ in the absence of disease and/or intervention. In what follows we restrict ourselves to the case $\beta > 0$, i.e. where social distancing incurs a cost. The edge case $\beta = 0$ changes the control problem fundamentally and leads to so-called bang-bang style behaviour $\kappa = 0$. We can therefore choose units for all utilities and costs in which $\beta = 1$ without loss of generality.

Government incentives (if any), are written $\varepsilon(t)$. These represent state level incentives (or penalties) designed to modify behaviour. For example, if $\varepsilon < 0$, the government is incentivising cautious behaviour $\kappa < \kappa^*$ and taxing risky behaviour $\kappa > \kappa^*$. The interpretation of ε as a tax/incentive would imply that whatever balance the government earns or spends by enacting ε is ultimately equally redistributed among the population.

We have chosen a strongly idealised model and utility function, in the hope of capturing relevant behaviour without adding unnecessary model complexity. It is common to model the social activity, κ or k here, as entering linearly in the epidemic model, e.g. [1, 7], and equally common to assume a convex, and in particular quadratic control cost, e.g. [7, 29]. Others have used different, but also convex, functional forms for the social distancing cost, e.g. [1], and while there are some quantitative differences in the results, the Nash equilibrium behaviour is qualitatively quite similar. As we will see, the results greatly depend on whether there is a step in the infection cost or not, but we believe the particular functional shape of the step to not be relevant, qualitatively. One, however, finds very different outcomes when one assumes that the infection cost per infection *decreases* with the number of infections [6], which can result in infection waves. We have assumed that all individuals have to pay the cost of social distancing equally in contrast to other work, e.g. [1, 31] where the cost of social distancing is paid mostly by the s -compartment. Their choice is motivated by the fact that only susceptibles can influence their fate with their own behaviour. Our choice is motivated by the observation that no individual, regardless of compartment, can socially distance without incurring a cost and aligns more closely with the situation in which an individual doesn't necessarily know in which compartment they are, e.g. the limit where many infections occur asymptotically. This would result in individuals that are in the infectious or recovered compartments acting as if they are susceptible. Treating this precisely would require a model with significant additional complexity. Future versions of our model may include explicitly for instance asymptomatic

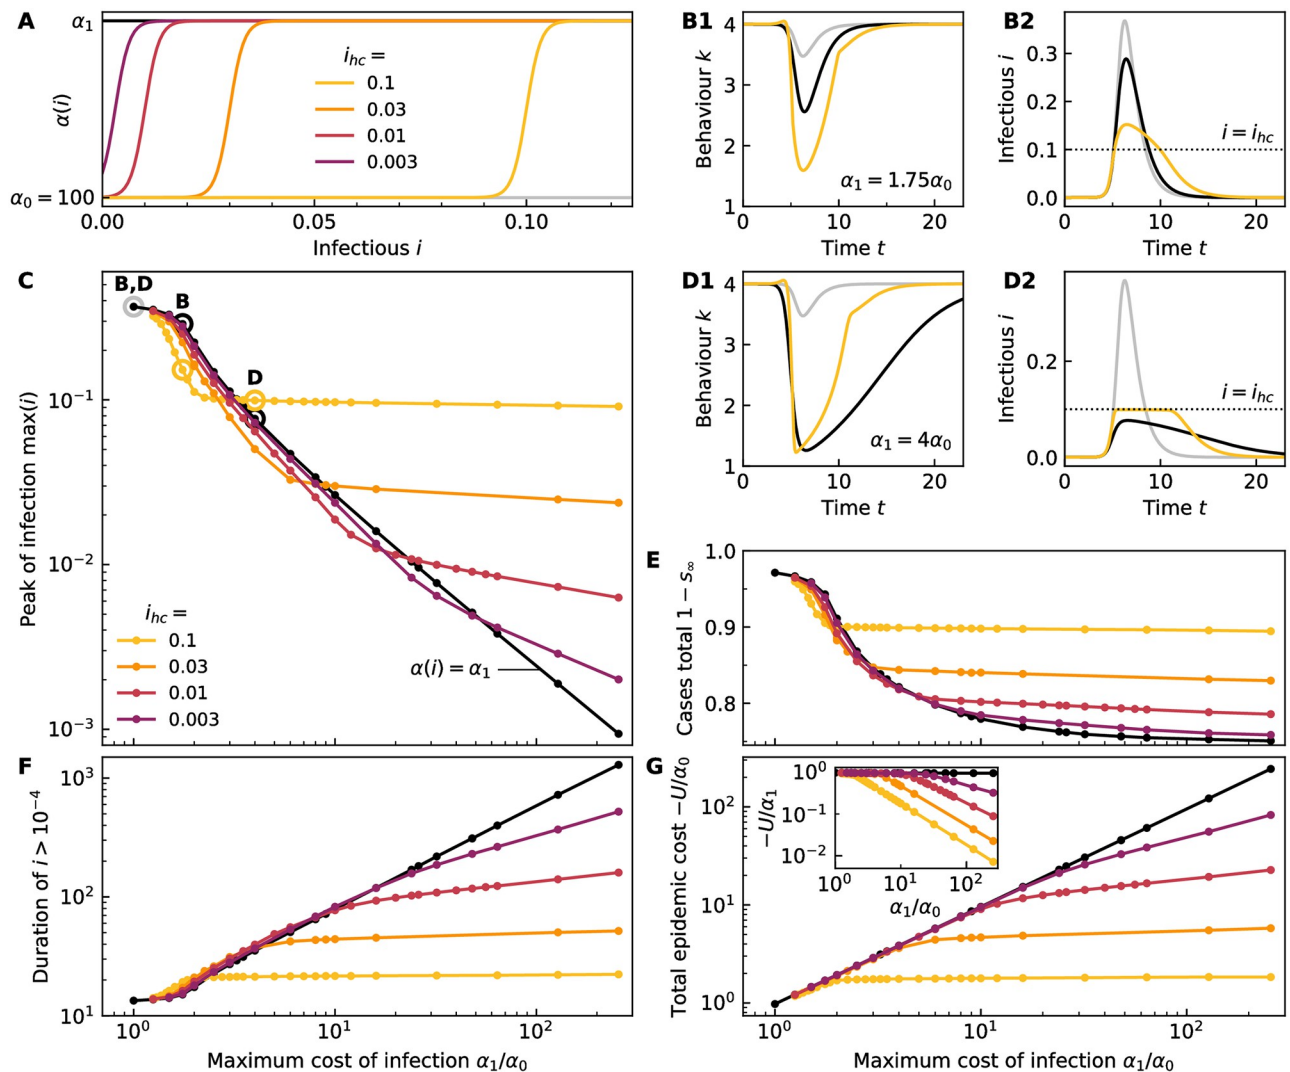


Fig 3. The rational behaviour in the presence of a healthcare threshold depends on the maximum cost of infection. (A) The infection cost $\alpha(i)$ for a range of healthcare thresholds i_{hc} , see Eq 6 with steepness $\sigma = 300$. The colours encode the position of the healthcare threshold i_{hc} for the whole figure. For comparison, two scenarios where $\alpha(i) = \alpha_0$ (grey line) and $\alpha(i) = \alpha_1$ (black) are considered as well. The base infection cost is kept constant throughout, $\alpha_0 = 100$, whereas α_1 is varied in the following panels. (B) Typical example for the equilibrium behaviour of population $k(t)$ (B1) and the corresponding infectious cases $i(t)$ (B2) over time for low maximum infection cost α_1 . We compare the behaviour for a healthcare threshold, $\alpha_0 = 100$, $\alpha_1 = 1.75\alpha_0$ with $i_{hc} = 0.1$ (yellow) to the behaviour for constant infection costs $\alpha(i) = 100$ (grey) and $\alpha(i) = 1.75\alpha_0$ (black). These cases are also marked in panel (C) by correspondingly coloured circles beneath the letters B. (C) The peak of the epidemic $\max(i)$ as a function of the maximum cost of being infectious α_1 corresponding to the infection cost scenarios shown in (A). We also mark the data points corresponding to the examples shown in panels (B1–2) and (D1–2) with circles and corresponding labels above. (D) Typical example for the equilibrium behaviour of population $k(t)$ (D1) and the corresponding infectious cases $i(t)$ (D2) over time for high maximum infection cost α_1 . We compare the behaviour for a healthcare threshold, $\alpha_0 = 100$, $\alpha_1 = 4\alpha_0$ with $i_{hc} = 0.1$ (yellow) to the behaviour for constant infection costs $\alpha(i) = 100$ (grey) and $\alpha(i) = 4\alpha_0$ (black). These cases are also marked in panel (C) by correspondingly coloured circles beneath the letters D. For the same data as in panel (C): (E) Total number of cases after the epidemic has run its course. (F) Duration of the epidemic as defined by the time interval for which $i > 10^{-4}$. (G) Total cost of the epidemic $-U$ for equilibrium behavior in units of the minimal infection cost α_0 . In the inset, the epidemic cost is shown in units of the maximum infection cost. Lines in (D–G) serve as guides to the eye.

<https://doi.org/10.1371/journal.pcbi.1011533.g003>

and exposed compartments, with separate controllable behaviours for each compartment. In addition, there can be peer-pressure effects for conformity across all compartments. As for the functional choice of the government intervention: we strongly idealised the situation and assumed that the government intervention acts as a bias on the behaviour κ linearly, with the

aim of allowing the government to both incentivise more or less activity. Alternative approaches would have been for the government to use an incentive ε to influence the cost of social distancing β by replacing the term $-\beta(\kappa - \kappa^*)^2$ with $-(\beta - \varepsilon)(\kappa - \kappa^*)^2$ (positive $\varepsilon < \beta$ allow κ to deviate more easily from κ^*) or to use ε as a tax to affect the cost of an infection $\alpha(i)$ by replacing the term $-\alpha(i)\psi_i$ by $-(\alpha(i) + \varepsilon)\psi_i$ (positive ε encourages more social distancing to avoid the increased infection cost). These ideas would be quite similar to what is explored in Ref. [30]. We believe that these choices would still allow the government to target the global optimum of its objective function by appropriately incentivising/taxing the population. Since we see our work as a proof of concept, we have only focused on one type of government intervention.

It is numerically convenient to truncate the utility integral at a final time t_f . Indeed this can be realistic if associated, e.g. with the rollout of mass vaccination. The contribution to the utility from the course of the epidemic after t_f is written U_f . Assuming the arrival of a perfect vaccine at t_f which reduces the fraction of susceptibles immediately to 0 and thus immediately reduces the incentive to social distance, $\kappa = \kappa^*$, the utility then reads

$$U = \int_0^{t_f} u(t)dt + U_f \quad (7)$$

$$U_f = \int_{t_f}^{\infty} u(t)dt = \int_{t_f}^{\infty} f^{-t}[-\alpha(i(t))\psi_i(t)]dt \quad (8)$$

which can be numerically integrated. For convenience, we approximated the salvage term

$$U_f \approx -f^{-t_f} \alpha(0) \frac{\psi_{if}}{1 + \log f} \quad (9)$$

see section D in S1 Text for a short derivation. We always choose t_f large enough so that i_f is extremely small (typically $\lesssim 10^{-8}$). As a result the approximation above is satisfied well and in addition U_f is negligible. However, the small contribution of U_f is always included in the figures and solutions we show here, for completeness. We note again that the arrival of a vaccine or treatment earlier during the course of the epidemic tends to enhance social distancing efforts [1, 7, 55]. If $\alpha(i)$ is not constant in that situation, the above approximation will not be accurate. However, in this work, t_f is assumed to always be sufficiently late for vaccination to have no behavioural or policy consequences.

The individual behaviour κ is assumed to satisfy the constrained optimisation problem

$$\begin{aligned} \kappa &= \arg \max_{\kappa} U \\ \text{subject to } \dot{\psi}_s &= -\kappa\psi_s i, \text{ with } \psi_s(0) = 1 - i_0, \text{ and} \\ \dot{\psi}_i &= \kappa\psi_s i - \psi_i, \text{ with } \psi_i(0) = i_0. \end{aligned} \quad (10)$$

The population behaviour $k(t)$ and therefore $s(t)$ and $i(t)$, as well as the government intervention $\varepsilon(t)$ are treated as external or exogenous quantities, outside of the individual's control. They merely represent an explicit time-dependence of the utility function and the individual's dynamics, to which the individual reacts by adjusting their behaviour without being able to affect them.

The solution to this optimisation problem can be calculated within a standard Hamiltonian/Lagrangian approach. Lev Pontryagin discovered that instead of solving the constrained optimisation problem directly, one can derive a simpler to solve set of differential equations that comprise a boundary value problem (BVP) [61]. As an intermediate step of deriving the

BVP one defines a Hamiltonian. What is now known as Pontryagin's Principle loosely states that an optimal control to solve the constrained optimisation problem must also solve this BVP, which in turn means that it also extremises the Hamiltonian. The BVP is equivalent to the Hamiltonian equations or Euler-Lagrange equations known in physics which can be derived when extremising an action integral. See sections A and B in [S1 Text](#), or references [9, 60], for a derivation and more details. Here we use this approach and, instead of solving the constrained optimisation problem directly, solve the BVP involving the Hamiltonian. The system's Hamiltonian for the individual behaviour is given by, see section B in [S1 Text](#),

$$\begin{aligned} H &= u + v_s(-\kappa\psi_s i) + v_i(\kappa\psi_s i - \psi_i) \\ &= -f^{-t}[\alpha(i)\psi_i + (\kappa - \kappa^*)^2 - \varepsilon(\kappa - \kappa^*)] \\ &\quad - (v_s - v_i)\kappa\psi_s i - v_i\psi_i \end{aligned} \quad (11)$$

Using this Hamiltonian, we can obtain additional differential equations and a condition on the control, which when solved together yield the optimal control. The Lagrange multipliers $v_s(t)$ and $v_i(t)$ constrain the dynamics to obey [Eq \(2\)](#). Furthermore, they can be seen as expressing the expected (economic) value of being in state s and i , respectively, at any given time. The Hamiltonian equations for the values (also called costate equations in the control theory literature) are

$$\begin{aligned} \dot{v}_s &= -\frac{\partial H}{\partial \psi_s} = (v_s - v_i)\kappa i \\ \dot{v}_i &= -\frac{\partial H}{\partial \psi_i} = f^{-t}\alpha(i) + v_i \end{aligned} \quad (12)$$

with boundary conditions

$$v_s(t_f) = \frac{\partial U_f}{\partial \psi_{s,f}} = 0, \quad v_i(t_f) = \frac{\partial U_f}{\partial \psi_{i,f}} = \frac{-f^{-t_f}\alpha(0)}{1 + \log f}. \quad (13)$$

The Nash equilibrium strategy for an individual follows from $0 = \partial H / \partial \kappa$ and reads

$$\kappa = \kappa^* - \frac{f^t}{2}(v_s - v_i)\psi_s i + \frac{1}{2}\varepsilon. \quad (14)$$

as long as this expression yields a plausible, non-negative value for κ , and $\kappa = 0$, otherwise. There are some subtleties with how this bound has to be enforced during numerical solution of the equations, which we describe in section C in [S1 Text](#).

Having obtained the optimal individual behaviour κ for any given population behaviour k which gives rise to the course of the epidemic i , we can now select the special case that constitutes a Nash equilibrium. Assuming that all individuals in the population are identical and would all independently choose the same strategy in response to a given population behaviour, we can then conclude that the average behaviour of the whole population has to be identical to each individual's behaviour, thus becoming the equilibrium behaviour, $k(t) = \kappa(t)$. Then, $s = \psi_s$ and $i = \psi_i$, as well as

$$k = \kappa = \max\left(0, \kappa^* - \frac{f^t}{2}(v_s - v_i)si + \frac{1}{2}\varepsilon\right). \quad (15)$$

Therefore, we expect social distancing to increase with how strongly the state of being susceptible is valued w.r.t. the state of being infectious, and to increase with the number of susceptibles as well as the infectious.

The equilibrium outcome of the epidemic can now easily be calculated for an exogenous government intervention field ε . This is achieved by numerically solving the boundary value problem of Eq 1 with boundary conditions $s(0) = 1 - i_0$ and $i(0) = i_0$, Eq 12 with boundary conditions Eq 13, in conjunction with Eq 15. We choose $i_0 = 3 \cdot 10^{-8}$ and $\kappa^* = 4$ and disregard discounting, $f = 1$. We use a typical numerical approach for such optimal control problems, a forward-backward sweep, see section C in S1 Text, or ref. [60] for more details and examples. Other methods for solving boundary value problems, such as a shooting method, would be applicable as well. Even though the cost of infection is a function of i , the objective function is convex in ψ_i , so we expect this optimisation problem to have a unique solution. As an example, this Nash solution is shown in Fig 2 for a constant infection cost $\alpha = 400$. The Nash behaviour leads to social distancing and therefore, compared to the non-behavioural case of $k = \kappa^*$, a longer duration for the epidemic with correspondingly lower infection rates and a smaller number of cases overall.

Utilitarian maximum

For comparison with the Nash equilibrium, we calculate the best possible population behaviour, corresponding to the limit of full cooperation on the level of individuals. This corresponds to directly optimising the corresponding population level utility of the same form

$$\begin{aligned} U_p &= \int_0^{t_f} u_p(t) dt + U_{p,f} \\ u_p &= f^{-t} [-\alpha(i) i - (k - \kappa^*)^2 + \varepsilon(k - \kappa^*)] \\ U_{p,f} &= -\frac{f^{-t_f} \alpha(0) i_f}{1 + \log f} \end{aligned} \quad (16)$$

to find the optimal k , subject to Eq 1 being satisfied. If adopted by the entire society this would yield the best possible outcome for all. For convenience, we use the same variable names for the Lagrange multipliers. Following the formalism described in section B in S1 Text again, the corresponding Hamiltonian is

$$\begin{aligned} H_p &= u_p + v_s(-ksi) + v_i(ksi - i) \\ &= -f^{-t} [\alpha(i)i + (k - \kappa^*)^2 - \varepsilon(k - \kappa^*)] - (v_s - v_i)ksi - v_i i \end{aligned} \quad (17)$$

and the Lagrange multipliers or expected values follow

$$\begin{aligned} \dot{v}_s &= -\frac{\partial H_p}{\partial s} = (v_s - v_i)ki \\ \dot{v}_i &= -\frac{\partial H_p}{\partial i} = f^{-t} [\alpha(i) + \frac{\partial \alpha(i)}{\partial i} i] + (v_s - v_i)ks + v_i \end{aligned} \quad (18)$$

with boundary conditions

$$v_s(t_f) = 0, \quad v_i(t_f) = -\frac{f^{-t_f} \alpha(0)}{1 + \log f} \quad (19)$$

The optimal strategy follows from $0 = \partial H_p / \partial k$ and yields the same decision rule as given by Eq 15 for the Nash equilibrium. The utilitarian behaviour ends up differing from the equilibrium behaviour because the equation for the Lagrange multiplier v_i gains a term $(v_s - v_i)ks$ that expresses the cost incurred from any infection causing further infections, which a self-interested individual does not consider. In general, the utilitarian optimum yields a higher utility than the Nash equilibrium, but is susceptible to defection by individuals who can gain at a

personal level at the expense of the rest of the population by adopting different strategies, up to the Nash equilibrium, see Fig 2. The Utilitarian behaviour can also be calculated with the forward-backward sweep method, see section C in S1 Text.

Government intervention strategy

The government's objectives are encoded in an objective function which has the same structure as the individual's but can have different parameter values

$$\begin{aligned} V &= \int_0^{t_f} v(t)dt + V_f \\ v &= f_g^{-t} [-\alpha_g(i)i - \beta_g(k - \kappa^*)^2 - \gamma_g \varepsilon (k - \kappa^*)] \\ V_f &= -\frac{f_g^{-t_f} \alpha_g(0) i_f}{1 + \log f_g}. \end{aligned} \quad (20)$$

with f_g a governmental discount rate and where α_g, β_g and γ_g account for the different costs assigned to outcomes, and interventions, at the government level. The sign change in the intervention term means that incentivising the population can be costly to the government. The pre-factor γ_g can account for how the cost of interventions can influence the government objective function. This is a way to model the shadow cost of public funds, i.e. the loss of utility due to the distortion of markets, etc., as caused by government intervention. The case of perfectly efficient intervention is given by $\gamma_g = 0$, while $\gamma_g > 0$ implies a loss of utility due to the process of intervention itself. We denote the lower and upper limits of $\alpha_g(i)$ as α_{g0} and α_{g1} , using the same sharpness σ as for individuals, see Eq 6. The small term V_f again models vaccination at t_f .

An important aspect of our work is that we investigate the situation in which the cost of an infection relative to the cost of social distancing can be quite different for the government than for an individual, $\alpha/\beta < \alpha_g/\beta_g$. For instance, it is likely to be more difficult for an individual to negotiate the right to work remotely than if the government imposes these arrangements.

The equilibrium behaviour expressed by Eq 15 uniquely determines the outcome of the epidemic in the presence of an imposed government policy $\varepsilon(t)$. We can therefore rewrite the SIR model as a function not of k , but of ε

$$\begin{aligned} \dot{s} &= -k(\varepsilon) s i \\ \dot{i} &= k(\varepsilon) s i - i \end{aligned} \quad (21)$$

In this spirit, it is the government determining the outcome of the epidemic with its choice of ε . In analogy to individual decision making, we now have an objective function and equations for the course of the epidemic that depend on a single control variable, but instead of optimising for κ , we optimise for ε . The complete government optimisation problem can therefore be framed as a constrained optimisation in ε, s , and i , such that

$$\begin{aligned} \varepsilon &= \arg \max_{\varepsilon} \left[\int_0^{t_f} v(t)dt + V_f \right] \\ v(t) &= f_g^{-t} [-\alpha_g(i)i - \beta_g(k(\varepsilon) - \kappa^*)^2 - \gamma_g \varepsilon (k(\varepsilon) - \kappa^*)] \\ V_f &= -\frac{f_g^{-t_f} \alpha_g(0) i_f}{1 + \log f_g} \\ \text{subject to } \dot{s} &= -k(\varepsilon) s i, \text{ with } s(0) = 1 - i_0, \text{ and} \\ \dot{i} &= k(\varepsilon) s i - i, \text{ with } i(0) = i_0 \end{aligned} \quad (22)$$

where $k(\epsilon)$ is obtained from solving its own constrained optimisation problem, Eqs 1, 12, 13 and 15, as already discussed above. We can follow the formalism described in section B in S1 Text again, noting that in the government optimisation ϵ now represents the control. The Hamiltonian for the government policy requires the introduction of two new Lagrange multipliers, λ_s and λ_i , (dropping most functional dependencies for brevity)

$$H_g = -f_g^{-t} [\alpha_g i + \beta_g (k(\epsilon) - \kappa^*)^2 + \gamma_g \epsilon (k(\epsilon) - \kappa^*)] - (\lambda_s - \lambda_i) k(\epsilon) s i - \lambda_i i \quad (23)$$

Then the differential equations for the values are, using Eq 15

$$\begin{aligned} \dot{\lambda}_s &= -\frac{\partial H_g}{\partial s} = i\Lambda \\ \dot{\lambda}_i &= -\frac{\partial H_g}{\partial i} = s\Lambda + f_g^{-t} [\alpha_g(i) + \alpha'_g(i)i] + \lambda_i \\ \Lambda &= -(\lambda_i - \lambda_s) \left(\kappa^* + \frac{\epsilon}{2} \right) + \frac{f_g^{-t} f^t}{2} (v_i - v_s) \times \\ &\quad (is[\beta_g f^t(v_i - v_s) - 2f_g^t(\lambda_i - \lambda_s)] + \epsilon(\gamma_g + \beta_g)) \end{aligned} \quad (24)$$

with boundary conditions

$$\lambda_s(t_f) = 0, \quad \lambda_i(t_f) = -\frac{f_g^{-t_f} \alpha_g(0)}{1 + \log f_g} \quad (25)$$

The optimal government strategy obeys $0 = \partial H_g / \partial \epsilon$ which yields, using Eq 15

$$\epsilon = is \frac{f^t(\beta_g + \gamma_g)(v_s - v_i) - f_g^t(\lambda_s - \lambda_i)}{\beta_g + 2\gamma_g} \quad (26)$$

We can obtain the government strategy with a nested application of a forward-backward sweep of Eqs 21, 24, 25 and 26, see section C in S1 Text. At each iteration, we use the current estimate of the optimal government strategy ϵ to calculate the Nash equilibrium behaviour $k(\epsilon)$, also with a forward-backward sweep and as described above, as part of the forward integration of the dynamics. This secondary forward-backward sweep treats government intervention as exogenous.

In the case of a constant infection cost, the government's objective function is convex and we expect numerically obtained solutions to be unique. In the case of a healthcare threshold, the government's objective function is not convex, in contrast to the individuals' objective function. For that situation it is therefore not straight-forward to establish uniqueness of our numerical solutions to the optimisation problem. In fact, by varying the initial guesses for the controls in the nested optimisation, we always found exactly two local optima for each set of parameters—never more or less—and selected the one with higher utility. We take this as indication that we successfully identify the global maximum in each case.

As an example, we calculate the government intervention strategy and the resulting incentivised equilibrium strategy for the situation where the government and individuals share the same preferences, $\alpha_g = \alpha = 400$, $\beta_g = \beta = 1$. If government intervention is free of cost for the government, $\gamma_g = 0$, the optimal government strategy $\epsilon(t)$ targets the utilitarian maximum for the population, see the gold lines in Fig 2. To achieve the utilitarian maximum, the ϵ field is used to bias the individuals' equilibrium strategy away from the unperturbed Nash equilibrium to coincide with the utilitarian maximum.

Results

Even though we strove for simplicity in our modeling choices, the model has a great number of parameters, κ^* , f , α_0 , α_1 , α_{g0} , α_{g1} , i_{hc} , σ , β , β_g , γ_g . We are therefore working in a moderately high-dimensional parameter space which would be challenging to fully explore. For simplicity, we adopted values representative of a disease like Covid-19. We selected single values for κ^* , f , α_0 , σ , β , β_g , while focusing on the effects of varying α_1 , α_{g1} , i_{hc} , γ_g to study the full range of behaviours and incentive strategies that might be expected to occur.

Results without government intervention

At first, we concentrate on the case where the cost of infection is constant, $\alpha(i) = \text{const}$ and where the government takes no role in the response to the epidemic, $\varepsilon = 0$. This situation has been already discussed for slightly different utilities, e.g. [1]. To appreciate the impact of optimal decision making, it is helpful to first establish a baseline: the course of an epidemic without any behavioural modification, $k = \kappa^*$, see Fig 2 (grey curves). This corresponds to a situation where there is no perceived risk associated with an infection, $\alpha = 0$. Since there is no behavioural modification, Fig 2A, and no government intervention, Fig 2B, the number of susceptibles s quickly drops, Fig 2C, as they become infected, Fig 2D. The peak of i is extremely large, with a fraction of roughly 0.4 of the whole population being infected at once. Consequently, the fraction of people that remains uninfected at the end of the epidemic, s_∞ , which is shorthand for $s(t \rightarrow \infty)$, reaches close to 0, Fig 2C. In contrast, for a moderate risk, $\alpha = 400$ (black lines), the population chooses a Nash equilibrium with considerable reduction in their activity k , Fig 2A, which reduces peak infection levels, Fig 2D, and the total number of cases $1 - s_\infty$, Fig 2C, at the expense of prolonging the epidemic. The utilitarian optimum (gold lines) can target a scenario where the duration of the epidemic is almost the same as for the baseline scenario, with a smaller total of cases than for the Nash equilibrium.

The higher the cost of infection α , the stronger is the behavioural modification, see black lines in Fig 3; see also examples for the equilibrium behaviour for $\alpha = 100$ (grey lines) and $\alpha = 175$ (black lines) in Fig 3B1 and 3B2, as well as $\alpha = 400$ (black lines) in Fig 3D1 and 3D2. As a consequence, the peak height of infections, Fig 3C, and the total number of cases, Fig 3E, are reduced with increasing α , whereas the epidemic duration, Fig 3F, increases. The duration on which the behaviour deviates from the pre-epidemic value κ^* is comparable to the duration of the epidemic. Even though the total number of cases decreases with rising α , the total epidemic cost $-U$ grows approximately proportionally to α , Fig 3G, see in particular the inset where the black line is almost exactly a constant. This implies that the gains in utility by avoiding cases in excess of herd immunity are almost completely offset by the cost of social distancing.

Next, we express the fact that healthcare systems have limited capacity by having the infection cost rise near a healthcare threshold i_{hc} , see Eq 6. We investigate the outcomes for a number of thresholds, see Fig 3A. We vary the value of i_{hc} while keeping the absolute steepness of the transition $\sigma = 300$ constant. This has the effect that the relative steepness σi_{hc} varies with i_{hc} , with the transitions being the steeper, the larger the threshold. This enables us to investigate the effects of threshold location and transition steepness at the same time. We set $\alpha_0 = 100$ and vary α_1 in relation to that. In passing, we note that $\alpha(i)$ is a monotonically increasing function and that the cost per infection at $i = 0$ is not necessarily exactly α_0 , but $\alpha(0) = \frac{1}{2}(\alpha_0 + \alpha_1 - (\alpha_1 - \alpha_0)\tanh(i_{hc}\sigma)) > \alpha_0$. For the healthcare thresholds that we studied, the difference can be completely neglected for $i_{hc} = 0.1$ and 0.03 , whereas for $i_{hc} = 0.01$ one

obtains a correction of $\alpha(0) \approx \alpha_0 + 2.5 \times 10^{-3}(\alpha_1 - \alpha_0)$ and for $i_{hc} = 0.003$, $\alpha(0) \approx \alpha_0 + 0.14(\alpha_1 - \alpha_0)$.

Varying the maximum infection cost α_1 at a given threshold i_{hc} , we find in general two qualitatively different Nash equilibrium strategies, see Fig 3B–3G. For instance, let us focus on $i_{hc} = 0.1$ for now, see the yellow lines.

(1) Low infection cost strategy: For low α_1 , it is rational to enact stronger social distancing than for the case of a constant high cost of infection, $\alpha = \alpha_1$. As an illustrative example, we show the equilibrium behaviour in the situation where the infection cost rises from $\alpha_0 = 100$ to $\alpha_1 = 1.75\alpha_0$ at $i_{hc} = 0.1$ in Fig 3B1 and 3B2 and compare that with the limiting cases of having constant infection cost $\alpha = \alpha_0$ (grey lines) and $\alpha = \alpha_1$ (black lines). We find that social distancing in the presence of the threshold is stronger than for both constant cost cases. It is obvious that social distancing would be more extreme when there is a healthcare threshold at which the cost increases from $\alpha = \alpha_0$ to α_1 than if $\alpha = \alpha_0$ always. But it is perhaps surprising that the situation with a healthcare threshold would call for stronger social distancing as compared to the case where $\alpha = \alpha_1$ always, given that the time averaged infection cost in the presence of the threshold is lower without any additional social distancing. However, the additional investment in social distancing is more than offset by the reduction in infection cost. Still, the peak of infection generally exceeds the health care threshold, see Fig 3D, if only slightly in the example. Strategy (1) is found to the left of the constant infection cost line for $\alpha = \alpha_1$ (black) in Fig 3C. Strategy (1) is also characterised by lowered case numbers as compared to the constant infection cost case, Fig 3E and slightly longer epidemic durations, Fig 3F. The total epidemic cost $-U$ is only slightly lower than for constant infection cost, in fact it is almost imperceivable on the scale of Fig 3G. Focusing again on Fig 3C, we see that the higher the infection cost, the lower the infection peak becomes, until it approximately meets the health care threshold at $\alpha_1 \approx 2\alpha_0 - 3\alpha_0$ for $i_{hc} = 0.1$, where the situation crosses over into:

(2) High infection cost strategy: If α_1 exceeds a critical value which depends on i_{hc} (and to a lesser extent on α_0 and σ), the rational strategy is not to exceed the health care threshold but to remain close to it. An illustrative example for the equilibrium behaviour in the situation where the infection cost rises from $\alpha_0 = 100$ to $\alpha_1 = 4\alpha_0$ at $i_{hc} = 0.1$ is shown as yellow line in Fig 3D1 and 3D2, comparing to the limiting cases of having constant infection cost $\alpha = \alpha_0$ (grey lines) and $\alpha = \alpha_1 = 4\alpha_0$ (black lines). This strategy yield less severe social distancing than the constant $\alpha = \alpha_1$ case. As a result we observe higher peaks of infection, Fig 3C, with a higher total of cases, Fig 3E, and shorter duration of the epidemic, Fig 3F, when compared to the constant infection cost case. However, since the healthcare threshold is generally not exceeded, the total epidemic cost is much lower than in the constant infection cost case, Fig 3G.

For the lower healthcare thresholds (darker colours), we find qualitatively similar behaviour. However, the more slowly $\alpha(i)$ varies at the threshold, Fig 3A, the more gradual is the transition between strategies (1) and (2). For the lower values of i_{hc} , e.g. $i_{hc} = 0.003$ (purple lines), the peak of infection keeps decreasing with increasing α_1 across the whole studied range, Fig 3C. This is due to the fact that the infection cost $\alpha(i)$ does not reach the constant value α_0 for finite i and thus any reduction in $\max(i)$ can yield a lower infection cost. For larger i_{hc} , $\max(i)$ becomes practically independent of α_1 at large α_1 because the infection cost at the peak has already reached $\alpha(\max(i)) = \alpha_0$. Corresponding tendencies are found for the total cases, epidemic duration, as well as total epidemic cost.

Fig 3G shows, that if α_1 is held constant, the total epidemic cost strongly decreases with increasing i_{hc} . This underlines the potentially significant benefit of investing in healthcare infrastructure in order to raise i_{hc} .

Results with government intervention

If government and individuals share the same preferences, $\alpha_g(i) = \alpha(i)$, $\beta_g = \beta = 1$, and if government intervention is free of cost for the government, $\gamma_g = 0$, the optimal government strategy $\varepsilon(t)$ gives rise to the utilitarian maximum for the population, see the gold lines in Fig 2 for an example where the infection cost is constant. To achieve the utilitarian maximum, the ε field is used to bias the individuals' equilibrium strategy in the presence of government intervention away from the unperturbed Nash equilibrium to coincide with the utilitarian maximum. If the government wishes to encourage more cautious behaviour, it selects $\varepsilon < 0$, which rewards behaviour $\kappa < \kappa^*$ and taxes $\kappa > \kappa^*$. Owing to the level of control the government has over the population with its intervention strategy, the government is able to achieve a course of the epidemic that is shorter while resulting in fewer infections in total. It achieves this by initially incentivising social activity and later on incentivising social distancing in a precisely controlled manner. It is very encouraging that this closely resembles the strategy of the Japanese government, with its "Go To campaign" from July 2020 onwards. This was designed to increase demand for domestic tourism. This was eventually phased out and replaced with policies to promote social distancing.

We note that when $\gamma_g = 0$, individual preferences are irrelevant for the course of the epidemic: The government will always be able to find an intervention strategy ε that makes the population's equilibrium behaviour align with the government's preferences. However, the greater the difference in preferences, the greater the amplitude of ε necessary to achieve this.

Next, we consider the case when intervention is costly for the government, see the cyan lines in Fig 2 for an example where the infection cost is constant and $\gamma_g = 0.5$. then we find that the government selects an intervention strategy, Fig 2B, which incentivises population behaviour that is less strongly varying over time, Fig 2A. The government does not necessarily intervene less, but it chooses to incentivise social distancing earlier in time so that the peak of social distancing can be less extreme, Fig 2A. Note that this policy yields fewer uninfected at the end of the epidemic, $s(t \rightarrow \infty)$, Fig 2C, as the socially optimal strategy. This occurs even though the peak of the epidemic is lower Fig 2D.

For constant infection cost, the government strategy only weakly depends on the infection cost, regardless of whether the intervention is costly or not: The peak of infection is relatively insensitive to α for both cost-free and costly intervention, see gold and cyan lines in Fig 4A respectively. However, the total number of cases approaches the herd immunity threshold more rapidly with intervention than without (black line), Fig 4B. Cost-free intervention enables this at lower infection cost than costly intervention. Government intervention also manages to keep the duration of the epidemic much shorter than without incentives, Fig 4C, at a lower total epidemic cost, Fig 4D. The inset, in which the total epidemic cost is normalised by the maximum infection cost, shows this more clearly. The intervention policy, Fig 5A, and its effect on population behaviour, Fig 5D, and the course of the epidemic, Fig 5G, vary only subtly with rising infection cost. The larger the infection cost, the longer social distancing is practised and the more gradually it is relaxed over extended periods of time.

If the capacity of the healthcare system is limited according to Eq 6, see Fig 3A, government intervention leads to a markedly different course of the epidemic as compared to the no-intervention equilibrium, see green and purple lines in Fig 4 and compare with the red lines. (We show data with $i_{hc} = 0.01$ but expect the scenario to be qualitatively the same for other thresholds.) Instead of a continuous reduction of the infection peak without government intervention, Fig 4A, incentives lead to a sharp switch between policies that favour high peak infection and those that track the health care threshold as the maximum infection cost α_{1g} increases. At low α_{1g} the government targets a solution with a higher peak of infections, Fig 4A, without

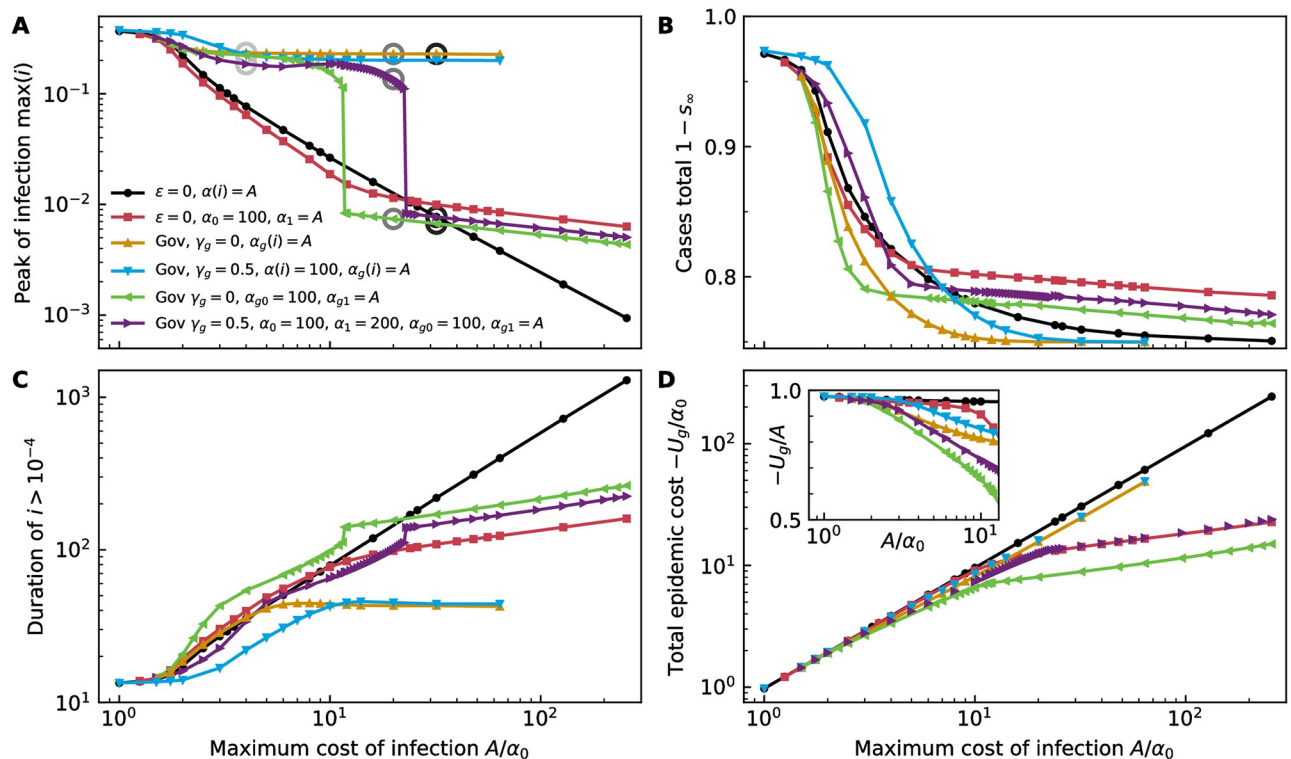


Fig 4. Optimal government policy. (A) Peak of the epidemic as a function of the maximum cost of infection for a range of scenarios where the (maximum) infection cost for either population or government is varied: Nash equilibrium behaviour of the population, without government intervention for a constant infection cost (black line, replotted from Fig 3) and with a healthcare threshold at $i_{hc} = 0.01$ (red, replotted from Fig 3); with government intervention for a constant infection cost α_g (cost-free $\gamma_g = 0$: gold, costly $\gamma_g = 0.5$: cyan) and with a healthcare threshold at $i_{hc} = 0.01$ (cost-free $\gamma_g = 0$: green, costly $\gamma_g = 0.5$: purple). The circles mark the scenarios shown in Fig 5. For these scenarios, we also show (B) the total number of cases, (C) the duration of the epidemic, as well as (D) the total cost of the epidemic in units of the minimal infection cost α_0 . In the cases without government intervention, the total cost is calculated as $-U$, whereas in the cases with government intervention, we report $-U_g$. In the inset, the epidemic cost is shown in units of the maximum infection cost. Lines serve as guides to the eye.

<https://doi.org/10.1371/journal.pcbi.1011533.g004>

necessarily increasing the total number of cases, Fig 4B, at the expense of a longer duration of the epidemic, Fig 4C. This strategy resembles the government strategy for constant infection cost. When the maximum infection cost is high, the government targets the healthcare threshold, albeit at a lower peak of infections than the population would be able to reach on its own. The crossover between the regimes depends on the direct cost for the interventions, controlled by the parameter γ_g . When the intervention is cost-free $\gamma_g = 0$, the crossover occurs at a markedly smaller maximum infection cost than for the case without government intervention.

Cost-free intervention enables a significantly lower total epidemic cost than no intervention, Fig 4D, as it targets the utilitarian optimum. As compared to no intervention, the costly intervention scenario also results in lowered total epidemic cost at low maximum infection cost but roughly the same total cost at high maximum infection costs. However, it achieves this by lowering the total case numbers which is offset by the cost of intervening.

For reference, we show the government intervention, the behaviour of the population in response to it, and the course of the epidemic for a range of maximum infection costs α_{1g} in Fig 5.

Regarding the switch in the government strategy, which leads to the sharp jump in the infection peak observed in Fig 4A, and examples of which we show in Fig 5B and 5C: As stated earlier, we always find two locally optimal solutions. These form branches, with one being

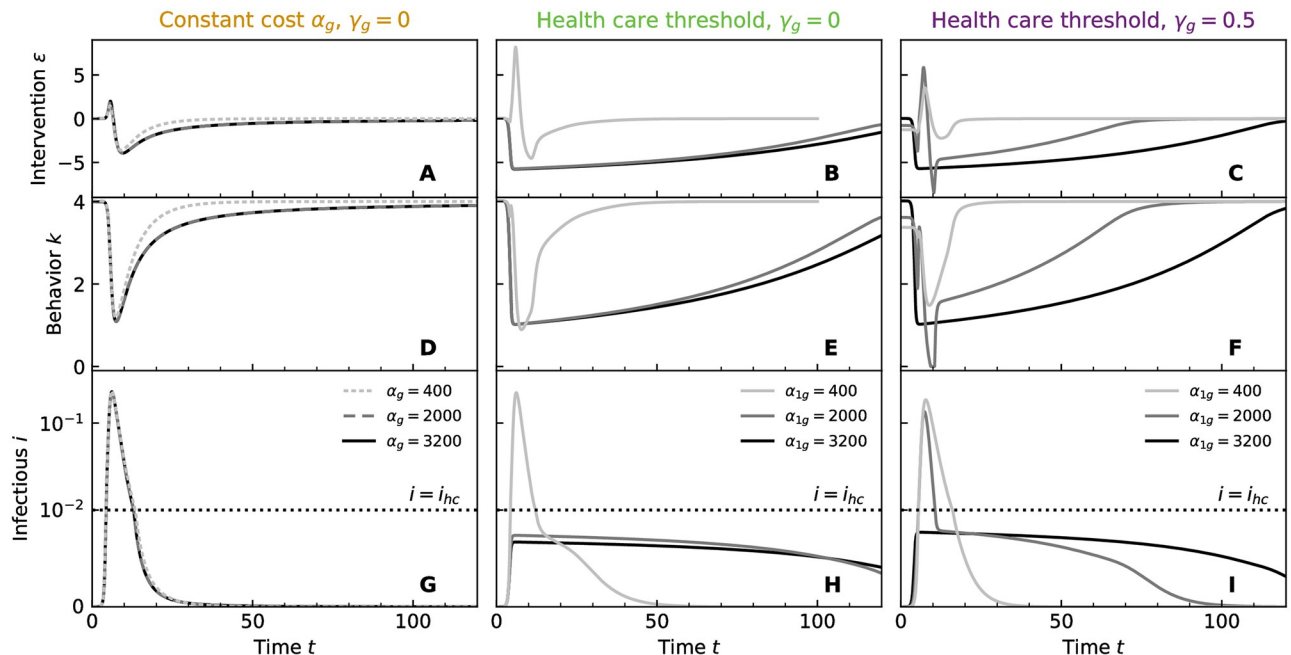


Fig 5. Course of the epidemic with government intervention. Government intervention ε assuming (A) constant infection cost α_g and cost-free intervention $\gamma_g = 0$, (B) a healthcare threshold (HT) and cost-free intervention $\gamma_g = 0$, and (C) a healthcare threshold and costly intervention $\gamma_g = 0.5$; all for a range of α_g or α_{1g} , respectively, as marked by circles in Fig 4A and listed in the legends of (G-I), with individuals assuming that $\alpha_0 = 100$. (D-F) Equilibrium population behaviour k in response to ε of (A-C), respectively. (G-I) Infectious i over time, corresponding to the behaviour shown in (D-F), respectively. Here, the y-axis has linear scale between 0 and 10^{-2} and logarithmic scale above that.

<https://doi.org/10.1371/journal.pcbi.1011533.g005>

globally optimal for low maximum infection cost and one being optimal for large maximum infection cost. We only show the solutions that are globally optimal, but the two branches both appear as linear on the log-log plot of Fig 4D, one at low infection cost and one at large infection cost. The switch in the government strategy occurs at the α_{g1} at which these branches yield the same value of the objective function. The policy that is optimal at low maximum infection cost α_{g1} is characterised by a high infection peak, Fig 4A, and shorter epidemic duration, Fig 4C; it therefore tolerates higher infection numbers in order to reduce costs incurred from social distancing. The policy that is optimal at large maximum infection cost α_{g1} favours an investment in stronger social distancing to avoid infections. The policy under high/low infection costs results in a greater/lower s_{∞} , Fig 4B. While the crossover between these policies is continuous in the maximum values of the objective function, Fig 4D, it results in very different disease trajectories, in particular in a discontinuous change in the peak infections, Fig 4A. In contrast, the Nash solution in the absence of government control is smooth because the ability of individuals to defect from an optimal consensus strategy leads to a smoothing out of the switch, Fig 4A.

Conclusion and discussion

Here, we have shown how costly interventions, such as taxes or subsidies on behaviour, can be used to exactly align individuals' decision making with government preferences even when these are not aligned. In order to achieve this, we developed a nested optimisation algorithm of both the government intervention strategy and the resulting equilibrium behaviour of

individuals. Healthcare systems in general, and intensive care facilities in particular, have limited capacity. For instance, intensive care units in Japan, the UK, and Germany had approximately 5, 7, and 34 beds per 100,000 people, respectively, in April 2020 [63, 64], with most of them regularly occupied. Assuming a healthcare threshold above which costs rise as a result of the rationing of scarce (intensive) care resources among patients, we find that it can be rational to adjust behaviour so that infections remain close to this threshold. This is a generic response when *either* the above-threshold costs α_1 or α_{1g} , for the individuals or government respectively, are high enough. However, the disease dynamics can be very different under government intervention than without it, see e.g. Fig 4A. We find that optimal government intervention strategies undergo a sharp “switch” from high peak infection numbers to a lower level, around the healthcare threshold. Furthermore, we find that both the maximum infection cost at which this switch occurs and the form of the intervention adopted are sensitive to how costly the intervention is to the government. For diseases that have infection costs around the value at which this policy switch occurs we anticipate that it would be very difficult for policymakers to know whether to adopt a high- or low-peak infection approach, particularly in the face of uncertainties. In the context of the COVID-19 epidemic it may be that the costs were such that the system was located close to this switch. This might help to explain why government policies to tackle COVID-19 differed so markedly between countries. Crude back-of-the-envelope analysis indicates that this may indeed be the case, although we are reluctant to assign values, this being fundamentally a political decision. In particular, if the infection cost were an order of magnitude higher/lower, policy determination would be straightforward.

Our results also show that a dramatic reduction in total epidemic cost can be achieved by increasing the healthcare threshold, implying the policy recommendation to do so.

Future work could include expanding our formalism to noisy dynamics, noisy control [25, 65–68], imperfect information or to study the robustness of the control, similar to [40, 69, 70]. There is also the intriguing possibility of allowing individuals to directly influence government [71] in the same way that ε allows the government to influence individuals. One approach might be to model political contentment, controlled by individuals, that would appear in the government objective function. This could give rise to a formalism with significant game theoretic complexity.

Supporting information

S1 Text. This text provides a background on the calculus of variations, optimal control theory, the forward-backward sweep for solving optimal control problems, and how to calculate our utility salvage term.

(PDF)

Acknowledgments

We would like to dedicate this work to the memory of Prof. George Rowlands, who passed away in early 2021 and who was involved in many of the early discussions leading to this work. We would like to thank Andrew J. Oswald, Steven Satchell, Michael Tildesley, Giorgos Galanis, Patrick Mellacher, and Holly C. Turner for careful reading of the manuscript and helpful comments. We also thank Tetsuya J. Kobayashi and Shuhei A. Horiguchi for helpful discussions. Matthew S. Turner acknowledges the kind hospitality of the Yamamoto group. The numerical code was written in Python, in particular using the numpy [72] and scipy [57] packages. Figures were generated using the Matplotlib Python library [73].

Author Contributions

Conceptualization: Simon K. Schnyder, John J. Molina, Ryoichi Yamamoto, Matthew S. Turner.

Formal analysis: Simon K. Schnyder, Matthew S. Turner.

Funding acquisition: Simon K. Schnyder, John J. Molina, Ryoichi Yamamoto, Matthew S. Turner.

Investigation: Simon K. Schnyder, Matthew S. Turner.

Methodology: Simon K. Schnyder, John J. Molina, Matthew S. Turner.

Software: Simon K. Schnyder, John J. Molina.

Supervision: Ryoichi Yamamoto, Matthew S. Turner.

Validation: John J. Molina.

Visualization: Simon K. Schnyder.

Writing – original draft: Simon K. Schnyder, John J. Molina, Matthew S. Turner.

Writing – review & editing: Simon K. Schnyder, John J. Molina, Ryoichi Yamamoto, Matthew S. Turner.

References

1. Reluga TC. Game Theory of Social Distancing in Response to an Epidemic. *PLoS Comput Biol*. 2010; 6(5):e1000793. <https://doi.org/10.1371/journal.pcbi.1000793> PMID: 20523740
2. Fenichel EP, Castillo-Chavez C, Coddia MG, Chowell G, Parra PAG, Hickling GJ, et al. Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences of the United States of America*. 2011; 108(15):6306–6311. <https://doi.org/10.1073/pnas.1011250108> PMID: 21444809
3. Wang Z, Bauch CT, Bhattacharyya S, D'Onofrio A, Manfredi P, Perc M, et al. Statistical physics of vaccination. *Physics Reports*. 2016; 664:1–113. <https://doi.org/10.1016/j.physrep.2016.10.006>
4. Chang SL, Piraveenan M, Pattison P, Prokopenko M. Game theoretic modelling of infectious disease dynamics and intervention methods: a review. *Journal of Biological Dynamics*. 2020; 14(1):57–89. <https://doi.org/10.1080/17513758.2020.1720322> PMID: 31996099
5. Verelst F, Willem L, Beutels P. Behavioural change models for infectious disease transmission: A systematic review (2010–2015). *Journal of the Royal Society Interface*. 2016; 13(125). <https://doi.org/10.1098/rsif.2016.0820> PMID: 28003528
6. Bhattacharyya S, Reluga T. Game dynamic model of social distancing while cost of infection varies with epidemic burden. *IMA J Appl Math (Institute Math Its Appl)*. 2019; 84(1):23–43.
7. Makris M, Toxvaerd F. Great Expectations: Social Distancing in Anticipation of Pharmaceutical Innovations; 2020. 2097.
8. Yan Y, Malik AA, Bayham J, Fenichel EP, Couzens C, Omer SB. Measuring voluntary and policy-induced social distancing behavior during the COVID-19 pandemic. *Proceedings of the National Academy of Sciences of the United States of America*. 2021; 118(16):1–9. <https://doi.org/10.1073/pnas.2008814118> PMID: 33820846
9. Reluga TC, Galvani AP. A general approach for population games with application to vaccination. *Math Biosci*. 2011; 230(2):67–78. <https://doi.org/10.1016/j.mbs.2011.01.003> PMID: 21277314
10. Acemoglu D, Chernozhukov V, Werning I, Whinston MD. Optimal targeted lockdowns in a multi-group SIR model. *National Bureau of Economic Research*; 2020. 27102.
11. Prem K, Cook AR, Jit M. Projecting social contact matrices in 152 countries using contact surveys and demographic data. *PLOS Computational Biology*. 2017; 13(9):e1005697. <https://doi.org/10.1371/journal.pcbi.1005697> PMID: 28898249
12. Huang CI, Crump RE, Brown PE, Spencer SEF, Miaka EM, Shampa C, et al. Identifying regions for enhanced control of gambiense sleeping sickness in the Democratic Republic of Congo. *Nature Communications*. 2022; 13(1):1–11. <https://doi.org/10.1038/s41467-022-29192-w> PMID: 35304479

13. Tildesley MJ, Vassall A, Riley S, Jit M, Sandmann F, Hill EM, et al. Optimal health and economic impact of non-pharmaceutical intervention measures prior and post vaccination in England: a mathematical modelling study. *Royal Society Open Science*. 2022; 9(8). <https://doi.org/10.1098/rsos.211746> PMID: 35958089
14. Chandrasekhar AG, Goldsmith-Pinkham P, Jackson MO, Thau S. Interacting regional policies in containing a disease. *Proceedings of the National Academy of Sciences of the United States of America*. 2021; 118(19):1–7. <https://doi.org/10.1073/pnas.2021520118> PMID: 33906950
15. Holme P, Saramäki J. Temporal networks. *Physics Reports*. 2012; 519(3):97–125. <https://doi.org/10.1016/j.physrep.2012.03.001>
16. Holme P, Masuda N. The basic reproduction number as a predictor for epidemic outbreaks in temporal networks. *PLoS ONE*. 2015; 10(3):1–15. <https://doi.org/10.1371/journal.pone.0120567> PMID: 25793764
17. He D, Dushoff J, Day T, Ma J, Earn DJD. Inferring the causes of the three waves of the 1918 influenza pandemic in England and Wales. *Proceedings of the Royal Society B: Biological Sciences*. 2013; 280(1766). <https://doi.org/10.1098/rspb.2013.1345> PMID: 23843396
18. Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, et al. Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Medicine*. 2008; 5(3):0381–0391. <https://doi.org/10.1371/journal.pmed.0050074>
19. Tildesley MJ, House TA, Bruhn MC, Curry RJ, O'Neil M, Allpress JLE, et al. Impact of spatial clustering on disease transmission and optimal control. *Proceedings of the National Academy of Sciences of the United States of America*. 2010; 107(3):1041–1046. <https://doi.org/10.1073/pnas.0909047107> PMID: 19955428
20. Sun K, Wang W, Gao L, Wang Y, Luo K, Ren L, et al. Transmission heterogeneities, kinetics, and controllability of SARS-CoV-2. *Science*. 2021; 371(6526):eabe2424. <https://doi.org/10.1126/science.abe2424> PMID: 33234698
21. Ferguson NM, Cummings DAT, Fraser C, Cajka JC, Cooley PC, Burke DS. Strategies for mitigating an influenza pandemic. *Nature*. 2006; 442(7101):448–452. <https://doi.org/10.1038/nature04795> PMID: 16642006
22. Tanimoto J. *Social Dilemma Analysis of the Spread of Infectious Disease*. Springer Nature Singapore; 2018.
23. Mellacher P. *COVID-Town: An Integrated Economic-Epidemiological Agent-Based Model*. Graz Schumpeter Centre; 2020. 23.
24. Grauer J, Löwen H, Liebchen B. Strategic spatiotemporal vaccine distribution increases the survival rate in an infectious disease like Covid-19. *Scientific Reports*. 2020; 10(1):1–10. <https://doi.org/10.1038/s41598-020-78447-3>
25. Yong J, Zhou XY. *Stochastic controls: Hamiltonian systems and HJB equations*. vol. 43. Springer Science & Business Media; 1999.
26. Molina JJ, Schnyder SK, Turner MS, Yamamoto R. *Nash Neural Networks: Inferring Utilities from Optimal Behaviour*. 2022;. arXiv:2203.13432.
27. Kahneman D. Maps of bounded rationality: Psychology for behavioral economics. *American economic review*. 2003; 93(5):1449–1475. <https://doi.org/10.1257/000282803322655392>
28. McAdams D. *Nash SIR: An Economic-Epidemiological Model of Strategic Behavior During a Viral Epidemic*. Covid Economics. 2020.
29. Eichenbaum MS, Rebelo S, Trabandt M. The Macroeconomics of Epidemics. *Review of Financial Studies*. 2021; 34(11):5149–5187. <https://doi.org/10.1093/rfs/hhab040>
30. Rowthorn R, Toxvaerd F. *The optimal control of infectious diseases via prevention and treatment*. University of Cambridge; 2020. 2027.
31. Toxvaerd F, Rowthorn R. *On the management of population immunity*. University of Cambridge; 2020. 2080.
32. Li J, Lindberg DV, Smith RA, Reluga TC. Provisioning of Public Health Can Be Designed to Anticipate Public Policy Responses. *Bull Math Biol*. 2017; 79(1):163–190. <https://doi.org/10.1007/s11538-016-0231-8> PMID: 27924411
33. Bethune ZA, Korinek A. *COVID-19 infection externalities: trading off lives vs. livelihoods*. Cambridge, MA: National Bureau of Economic Research; 2020. 27009.
34. Aurell A, Carmona R, Dayanikli G, Laurière M. Optimal Incentives to Mitigate Epidemics: A Stackelberg Mean Field Game Approach. *SIAM Journal on Control and Optimization*. 2022; 60(2):S294–S322. <https://doi.org/10.1137/20M1377862>
35. Althouse BM, Bergstrom TC, Bergstrom CT. A public choice framework for controlling transmissible and evolving diseases. *Proceedings of the National Academy of Sciences of the United States of*

- America. 2010; 107(SUPPL. 1):1696–1701. <https://doi.org/10.1073/pnas.0906078107> PMID: 20018681
36. Stiglitz J, Dasgupta P. Differential Taxation, Public Goods, and Economic Efficiency. *Review of Economic Studies*. 1971; 38(2):151–174. <https://doi.org/10.2307/2296778>
 37. Palmer AZ, Zabinsky ZB, Liu S. Optimal control of COVID-19 infection rate considering social costs. 2020; p. 1–27.
 38. Hayhoe M, Barreras F, Preciado VM. Data-Driven Control of the COVID-19 Outbreak via Non-Pharmaceutical Interventions: A Geometric Programming Approach. *Annual reviews in control*. 2020; 52 (NA):495–507. <https://doi.org/10.1016/j.arcontrol.2021.04.014>
 39. Piguillem F, Shi L. Optimal COVID-19 Quarantine and Testing Policies; 2020. 20/04.
 40. Kantner M, Koprucki T. Beyond just “flattening the curve”: Optimal control of epidemics with purely non-pharmaceutical interventions. *Journal of Mathematics in Industry*. 2020; 10(1). <https://doi.org/10.1186/s13362-020-00091-3> PMID: 32834921
 41. Ketcheson DI. Optimal control of an SIR epidemic through finite-time non-pharmaceutical intervention. *Journal of mathematical biology*. 2021; 83(1):7. <https://doi.org/10.1007/s00285-021-01628-9> PMID: 34176029
 42. Schwarzendahl FJ, Grauer J, Liebchen B, Löwen H. Mutation induced infection waves in diseases like COVID-19. *Scientific Reports*. 2022; 12(1):1–11. <https://doi.org/10.1038/s41598-022-13137-w> PMID: 35688998
 43. Giannitsarou C, Kissler S, Toxvaerd F. Waning Immunity and the Second Wave: Some Projections for SARS-CoV-2. *American Economic Review: Insights*. 2021; 3(3):321–338.
 44. Lux T. The social dynamics of COVID-19. *Physica A: Statistical Mechanics and its Applications*. 2021; 567:125710. <https://doi.org/10.1016/j.physa.2020.125710> PMID: 33879957
 45. Bauch CT, Galvani AP, Earn DJD. Group interest versus self-interest in smallpox vaccination policy. *Proceedings of the National Academy of Sciences of the United States of America*. 2003; 100 (18):10564–10567. <https://doi.org/10.1073/pnas.1731324100> PMID: 12920181
 46. Bauch CT, Earn DJD. Vaccination and the theory of games. *Proceedings of the National Academy of Sciences of the United States of America*. 2004; 101(36):13391–13394. <https://doi.org/10.1073/pnas.0403823101> PMID: 15329411
 47. Reluga TC, Bauch CT, Galvani AP. Evolving public perceptions and stability in vaccine uptake. *Mathematical Biosciences*. 2006; 204(2):185–198. <https://doi.org/10.1016/j.mbs.2006.08.015> PMID: 17056073
 48. Tildesley MJ, Savill NJ, Shaw DJ, Deardon R, Brooks SP, Woolhouse MEJ, et al. Optimal reactive vaccination strategies for a foot-and-mouth outbreak in the UK. *Nature*. 2006; 440(7080):83–86. <https://doi.org/10.1038/nature04324> PMID: 16511494
 49. Chen F, Toxvaerd F. The economics of vaccination. *Journal of Theoretical Biology*. 2014; 363:105–117. <https://doi.org/10.1016/j.jtbi.2014.08.003> PMID: 25111844
 50. Moore S, Hill EM, Dyson L, Tildesley MJ, Keeling MJ. Modelling optimal vaccination strategy for SARS-CoV-2 in the UK. *PLoS Computational Biology*. 2021; 17(5):1–20. <https://doi.org/10.1371/journal.pcbi.1008849> PMID: 33956791
 51. Moore S, Hill EM, Dyson L, Tildesley MJ, Keeling MJ. Retrospectively modeling the effects of increased global vaccine sharing on the COVID-19 pandemic. *Nature Medicine*. 2022; 28(11):2416–2423. <https://doi.org/10.1038/s41591-022-02064-y> PMID: 36302894
 52. Hill EM, Prosser NS, Ferguson E, Kaler J, Green MJ, Keeling MJ, et al. Modelling livestock infectious disease control policy under differing social perspectives on vaccination behaviour. *PLOS Computational Biology*. 2022; 18(7):e1010235. <https://doi.org/10.1371/journal.pcbi.1010235> PMID: 35834473
 53. Keeling MJ, Moore S, Penman BS, Hill EM. The impacts of SARS-CoV-2 vaccine dose separation and targeting on the COVID-19 epidemic in England. *Nature Communications*. 2023; 14(1):1–10. <https://doi.org/10.1038/s41467-023-35943-0> PMID: 36765050
 54. Kucharski AJ, Klepac P, Conlan AJK, Kissler SM, Tang ML, Fry H, et al. Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of SARS-CoV-2 in different settings: a mathematical modelling study. *The Lancet Infectious Diseases*. 2020; 20(10):1151–1160. [https://doi.org/10.1016/S1473-3099\(20\)30457-6](https://doi.org/10.1016/S1473-3099(20)30457-6) PMID: 32559451
 55. Schnyder SK, Molina JJ, Yamamoto R, Turner MS. Rational social distancing in epidemics with uncertain vaccination timing. *PLOS ONE*. 2023; 18(7):e0288963. <https://doi.org/10.1371/journal.pone.0288963> PMID: 37478107
 56. Kermack WO, McKendrick AG. 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*. 1927; 115(772):700–721.

57. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*. 2020; 17(3):261–272. <https://doi.org/10.1038/s41592-019-0686-2> PMID: 32015543
58. Bensoussan A, Frehse J, Yam P. *Mean Field Games and Mean Field Type Control Theory*. Springer; 2013.
59. Carmona R, Delarue F. *Probabilistic Theory of Mean Field Games with Applications I*. Springer; 2018.
60. Lenhart S, Workman J. *Optimal Control Applied to Biological Models*. Chapman and Hall/CRC; 2007.
61. Pontryagin LS, Boltyanskii VG, Gamkrelidze RV, Mishchenko EF. *The Mathematical Theory of Optimal Processes*. Gordon and Breach Science Publishers; 1986.
62. Schnyder SK, Molina JJ, Yamamoto R, Turner MS. In preparation.
63. Maemura A. Japan ranks below Italy and Spain in ICU bed capacity; 2020. <https://asia.nikkei.com/Spotlight/Coronavirus/Japan-ranks-below-Italy-and-Spain-in-ICU-bed-capacity>.
64. Campbell D. NHS hospital bosses urge ministers to increase ICU beds in England; 2021. <https://www.theguardian.com/society/2021/feb/28/uk-government-must-increase-number-of-nhs-beds-hospital-bosses-warn>.
65. Lorch L, De A, Bhatt S, Trouleau W, Upadhyay U, Gomez-Rodriguez M. *Stochastic Optimal Control of Epidemic Processes in Networks*. 2018;.
66. Tottori T, Kobayashi TJ. Memory-Limited Partially Observable Stochastic Control and Its Mean-Field Control Approach. *Entropy*. 2022; 24(11):1–27. <https://doi.org/10.3390/e24111599> PMID: 36359688
67. Tottori T, Kobayashi TJ. Forward-Backward Sweep Method for the System of HJB-FP Equations in Memory-Limited Partially Observable Stochastic Control. *Entropy*. 2023; 25(2). <https://doi.org/10.3390/e25020208> PMID: 36832575
68. Tottori T, Kobayashi TJ. Decentralized Stochastic Control with Finite-Dimensional Memories: A Memory Limitation Approach. *Entropy*. 2023; 25(5):1–26. <https://doi.org/10.3390/e25050791> PMID: 37238546
69. Köhler J, Schwenkel L, Koch A, Berberich J, Pauli P, Allgöwer F. Robust and optimal predictive control of the COVID-19 outbreak. *Annual Reviews in Control*. 2021; 51(November 2020):525–539. <https://doi.org/10.1016/j.arcontrol.2020.11.002> PMID: 33362428
70. Morris DH, Rossine FW, Plotkin JB, Levin SA. Optimal, near-optimal, and robust epidemic control. *Communications Physics*. 2021; 4(1):1–8. <https://doi.org/10.1038/s42005-021-00570-y>
71. Mellacher P. The impact of corona populism: Empirical evidence from Austria and theory. *Journal of Economic Behavior & Organization*. 2023; 209:113–140. <https://doi.org/10.1016/j.jebo.2023.02.021> PMID: 36941842
72. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, et al. Array programming with NumPy. *Nature*. 2020; 585(7825):357–362. <https://doi.org/10.1038/s41586-020-2649-2> PMID: 32939066
73. Hunter JD. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*. 2007; 9(3):90–95. <https://doi.org/10.1109/MCSE.2007.55>