Title: Quantifying the impact of individual and collective compliance with infection control measures for ethical public health policy

Short Title: Modeling for ethically sound public health policy

Authors Daniel Roberts[1], Euzebiusz Jamrozik[2,3,4], George S. Heriot[4], Michael J. Selgelid[5], and Joel C. Miller[1]*

Affiliations

[1]La Trobe University, School of Engineering and Mathematical Sciences.
[2]Monash University.
[3]University of Oxford.
[4]University of Melbourne.
[5]McMaster University

*joel.c.miller.research@gmail.com

Abstract

Compliance with infectious disease control measures can benefit public health but be burdensome for individuals. This raises ethical questions regarding the value of the public health benefit created by individual and collective compliance. Answering such questions requires estimating the total benefit from an individual's compliance, and how much of that benefit is experienced by others. This is complicated by "overdetermination" in infectious disease transmission: each susceptible person may have contact with more than one infectious individual, such that preventing one transmission may have no net effect if the same susceptible person is infected later. This article explores mathematical techniques enabling quantification of the impacts of individuals and groups complying with three types of public health measures: quarantine of arrivals, isolation of infected individuals, and vaccination/prophylaxis. The models presented suggest that these interventions all exhibit synergy: each intervention becomes more effective on a per-individual basis as the number complying increases, because overdetermination of outcomes is reduced, Thus additional compliance reduces transmission to a greater degree.

Introduction

Infectious diseases and the public health policies used to control them can both impose significant burdens on societies and individuals. Notably, the ongoing COVID-19 pandemic and response have had severe worldwide health and economic impacts. Collective action plays a significant role in public health interventions against infectious disease. This is because the actions one individual takes can have significant impact on whether others are exposed to infection, and thus they affect outcomes that others experience. Further, the magnitude of the impact of one individual's actions on others may be increased or undermined by the actions of others.

Increased population compliance with effective measures against infectious disease can lead to larger health benefits. However, many people may be unwilling to comply with public health measures, especially where they suffer a large personal cost but gain very little, or

perhaps no, individual benefit from their own compliance. Examples include cases where the individual is expected to have only mild disease (if infected) but severity is high for others (who the individual might end up infecting), or where the individual is unlikely to be infected but if infected would be likely to cause many additional infections. The potential conflicts highlighted here lead to challenging questions about the potential ethical acceptability (or lack thereof) of policy involving coercion of compliance with public health measures. There has, to date, been relatively little mathematically-informed *ethical analysis* of public health policy (*1*, *2*).

Ethical evaluations of infectious disease control policies should partly depend on the anticipated benefits and burdens of compliance with public health measures. Among other considerations, burdensome public health policies are arguably more ethically justifiable to the extent that the policies are expected to produce net public health benefits (outweighing the burdens or harms of the interventions). Further, policies that involve limitations on individual freedom might be more justifiable to the extent that the behavior restricted by policy is likely to result in harm to others that otherwise would not have occurred.

Significant mathematical modeling efforts have focused on health and economic impacts of SARS, pandemic influenza, and COVID-19 as well as the impacts of public health interventions(*3*, *4*). However, although numerous ethical considerations are directly relevant to the justification of epidemic control policies (*5*), investigations of the ethical implications of pandemic response in general and the COVID-19 response in particular have typically focused on allocation of scarce resources (*5*–*9*), disparities in health outcomes(*10*), and issues of research ethics (*11*, *12*), as opposed to quantitatively-informed ethical analysis of the benefits and harms of control policies and individual compliance with public health measures.

This article examines mathematical modeling techniques we have developed to explore how individual and collective behavior changes can affect two specific outcomes:

- The probability an epidemic becomes established in a population.
- The total number of infections that occur once an epidemic is established.

Our modeling approaches allow us to quantify the impact of a single individual's behavior on these population outcomes. The models are designed to measure the impact of ethically salient aspects of transmission dynamics, including overdetermination and superspreading.

- **Overdetermination** occurs when a given outcome has more than one sufficient cause. For example, whether (i) an epidemic becomes established in a population or whether (ii) a specific individual becomes infected in an epidemic may be "overdetermined" when, respectively, (i) there are multiple introductions into a population (each of which would have been sufficient to cause an epidemic) and (ii) an individual is exposed to multiple infectious people (where each exposure would have been sufficient to infect the individual in question). This might be ethically salient because where overdetermination is significant (*e.g.*, where there are multiple introductions into a population or in high transmission settings where each susceptible person experiences multiple exposures to infection), one person changing their behavior to reduce their risk of infecting others might make less difference to harmful outcomes because these outcomes will be more likely to occur in any case, due to the risk imposition of others .
- **Superspreading** diseases, including COVID-19, are characterized by the tendency for a small fraction of infected individuals to cause a large proportion of all

transmissions, while most cause few or even no transmissions (*13–17*). Among other things, this affects the probability that a single introduction leads to established transmission in the population(*17*, *18*). Because most individuals cause very little transmission, the disease typically only becomes established if there is an early superspreading event. We therefore explore the impact of superspreading (or `dispersion' of the offspring distribution) by comparing the expected spread of an epidemic where superspreading is uncommon with an epidemic where superspreading accounts for a high proportion of transmissions, but holding the average number of transmissions caused (the reproduction number, $\mathcal{R}_0$) constant.

The primary goal in this article is to examine the development and application of mathematical models to investigate the impact of individuals' compliance with infectious disease control measures. We measure the effectiveness of relevant behaviors in terms of their impact on (i) the probability of an epidemic occurring and (ii) the total number of infections caused in an epidemic. These are both affected by overdetermination and superspreading, both of which are closely connected to random (stochastic) events. The methods we develop allow us to understand how stochasticity influences outcomes. The specific behaviors we investigate are isolating to prevent an epidemic from starting, and—if an epidemic is established—behavior changes to avoid infection or onwards transmission.

We begin this paper by describing the major results of our model, showing that the benefit of interventions can be significant, but much of the benefit can be experienced by people other than those who are changing their behavior. Then we discuss the ethical implications of the results, with a focus on the fact that as more individuals comply with an intervention, the population benefit is increased, though much of the benefit may be experienced by those who are not experiencing the cost of the intervention. Finally, we end the paper with a derivation of the mathematical tools that allow us to quantify the benefit from behavior change at the individual scale.

Results

We focus on the impact of three behaviors:

- *Border Quarantine*: Some individuals may quarantine upon arriving at a population to prevent an epidemic from occurring
- *Post-exposure isolation*: In an ongoing epidemic, some newly infected individuals may isolate to prevent transmitting to others.
- *Vaccination or Prophylaxis* : In an ongoing epidemic, some individuals may take extra actions such as vaccination or prophylaxis to reduce their own risk of becoming infected and thereby reduce their risk of transmitting to others.

As a general rule, we find that the per-individual impact of multiple individuals changing behavior generally increases as more individuals change behavior. Motivated by this, for each of these three interventions we are interested in answering three questions:

- What is the expected impact if a single individual adopts the behavior while everyone else continues as normal?
- What is the average impact of individuals if a fraction of the population adopts the behavior?
- What is the marginal impact if one more individual were to adopt (or abandon) the behavior after a fraction has adopted it?

Quarantine to Prevent an Epidemic

Many communities have historically had restrictions to prevent the introduction of individuals with infections from one geographic area or (sub-)population to another. Recently quarantines have been widely used in response to COVID-19. If an infected individual enters the population either by skipping quarantine or due to ineffective quarantine, there is a chance that an epidemic may result. However, by random chance, the infected individual perhaps would not cause any transmissions or only start a small chain of infections that dies out quickly.

We start by analyzing how an individual isolating impacts the probability that an epidemic occurs. Mathematically, this requires only some well-known results of probability generating functions and some basic results from probability. The results allow us to investigate the role of overdetermination on the start of an epidemic, and to introduce the framework by which we will evaluate individual actions to reduce epidemic spread.

To quantify the impact of quarantining arrivals to a location, we begin by considering the introduction of a single infected individual into a large completely susceptible population. Based on assumed knowledge about the offspring distribution we calculate the probability that this single introduction results in an epidemic. From this it is a simple calculation to look at what happens if multiple introductions occur.

Given a known offspring distribution, the probability $\mathcal{P}$ that a single introduction into a large completely susceptible population results in an epidemic can be calculated using probability generating functions (*17, 18*). For $\mathcal{R}_c < 1$ we find $\mathcal{P} = 0$, while for $\mathcal{R}_c > 1$ we find $\mathcal{P} > 0$. The epidemic probability tends to increase with $\mathcal{R}_c$ and decrease as the offspring distribution becomes more heterogeneous. It should be noted however, that in the case of higher heterogeneity, those outbreaks that do become established epidemics typically have an initially

larger growth than expected on average because they are likely to be seeded by superspreading events (*18*).

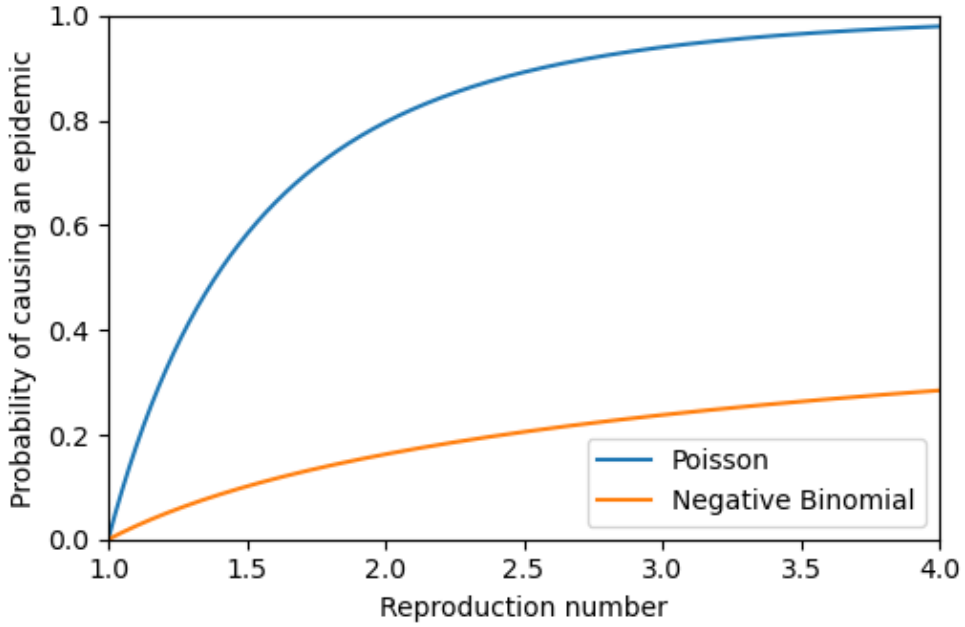

Figure 1: The probability a single infected individual entering a completely susceptible population causes an epidemic, assuming no other introductions. The offspring distribution follows either a Poisson distribution or a negative binomial distribution with a dispersion parameter of 0.16.

Figure 1 shows the probability that a single infected individual entering a completely susceptible well-mixed population would cause an epidemic, as a function of $\mathcal{R}_c$ for our two offspring distributions. If this single infected individual undergoes effective isolation, and no other infected individuals arrive, then Figure 1 would show the reduction in epidemic probability. However, if other infected individuals are not isolated, then they can trigger an epidemic instead. If an epidemic occurs, the first individual's isolation would be futile. In other words, transmission chains beginning from multiple introduced infections may each be separately sufficient to spark an epidemic. That is, the establishment of an epidemic may be overdetermined.

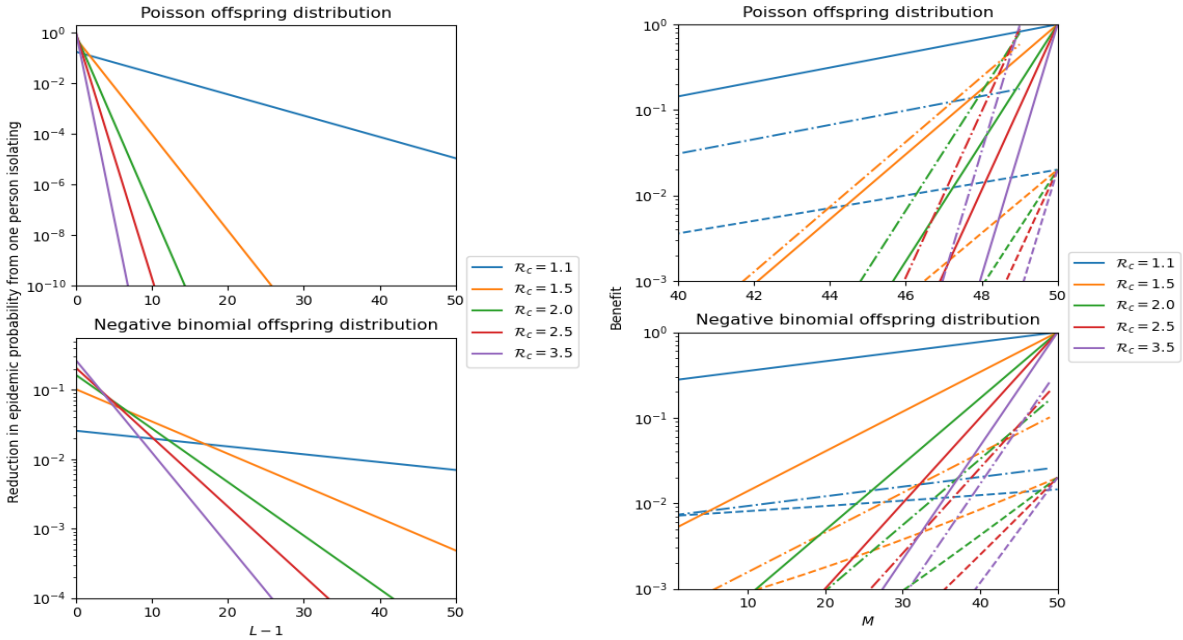

Figure 2 The benefit measured as the reduction in probability of an epidemic. We use a Poisson offspring distribution and a Negative Binomial offspring distribution with dispersion 0.16. We assume that L infected individuals enter the population of whom M isolate. ***Left*** *(note different vertical scales): The reduction in epidemic probability if $M = 1$ (only one isolates) as a function of the number ($L-1$) who do not isolate. Due to overdetermination, this decreases as L increases.* ***Right*** *(note different horizontal scales): The reduction in probability if M isolate out of $L = 50$ introduced cases. This shows the total reduction in epidemic probability (solid line), the reduction averaged across the M isolating individuals (dashed line), and the marginal benefit of one more isolating individual, (*i.e.*, the increased reduction in probability if one more were to isolate) (dash-dot line). As a consequence of overdetermination each additional individual who isolates increases the average impact of all isolating individuals.*

The impact of a single isolating infectious individual on the probability of an epidemic depends on the number of total infectious individuals (M) and how many of them isolate (L) as shown in Figure *2*. Overdetermination plays a large role in the trends in these figures. Each individual who does not isolate has an independent chance of triggering an epidemic. If even a single individual does trigger an epidemic, then those who isolate have no impact on the epidemic probability. Thus the effectiveness of those who do isolate is increased as the number that isolate increases. We see that if the probability that a single infected individual would trigger an epidemic is larger (larger $\mathcal{R}_c$), only a small number of nonisolating infected individuals are needed to effectively eliminate the impact of all those who isolate. Thus, isolating to prevent an epidemic is unlikely to be effective unless almost all infected individuals isolate (it should be noted that although it is a low probability event the societal benefit would be large).

Conversely, if almost every infection is isolating, the impact of one individual not isolating is significantly larger.

The burdens of quarantine largely fall on the quarantined population. As many of these may not be infected, in calculating the benefit achieved per quarantined individual it must be scaled by the probability of being infected. However, the general trend remains consistent: to significantly reduce the probability of an epidemic, almost all infected individuals must isolate, and so in the absence of an effective screening test, the quarantine must apply broadly, significantly increasing the burdens. It might be said that, as more individuals quarantine, the impact is synergistic: each additional quarantining individual increases the effectiveness of the others who quarantine. The impact of quarantine grows faster than linearly with increasing numbers of individuals quarantining. If even a few individuals fail to quarantine, however, this may significantly undermine the efforts of those who quarantine. As greater numbers of individuals quarantine, therefore, the stakes regarding individual behavior with regard to additional compliance and/or continuing lack thereof become higher: the benefits of an additional individual's compliance and the risks attached to one's continuing noncompliance both become greater.

Isolation following infection

We now analyze the impact of behavior changes on the total number of infections in a population assuming an epidemic is established. We consider two different actions an individual could take:

- actions taken after exposure to avoid potentially transmitting further (such as isolation after a rapid test or after being identified through contact tracing) which does not affect one's own probability of becoming infected but does protect others or
- actions taken in advance to avoid infection and hence also onwards transmission (such as vaccine or prophylactic medication).

We begin here focusing on the mathematically simpler of these two, isolation of a newly-infected individual. As before, we look first at the effect of a single individual and then the effect of a collective behavior change. We assume the control measures remain constant for the duration of the epidemic.

We now consider the expected (average) number of averted infections from a single individual isolating *after* she is exposed, but before she becomes infectious. So we assume that her probability of infection is unaffected, but that if infected she causes 0 additional infections. This model is appropriate for a setting with strong contact tracing or with good access to rapid testing.

We denote the expected total number of averted infections from a single newly infected individual isolating by $\mathcal{F}(\mathcal{R}_c)$. We account for overdetermination when calculating $\mathcal{F}(\mathcal{R}_c)$. That is, in calculating $\mathcal{F}(\mathcal{R}_c)$ we exclude those who would be infected anyways through another transmission chain (overdetermination), but otherwise consider all "descendants". For

$\mathcal{R}_c > 1$, overdetermination plays a role and the impact of overdetermination grows with $\mathcal{R}_c$. For $\mathcal{R}_c \leq 1$, overdetermination is negligible (for large populations). Perhaps surprisingly, the value of $\mathcal{F}$ is independent of the details of the offspring distribution, it depends only on the average, $\mathcal{R}_c$. It turns out [see methods for derivation] that if $\mathcal{R}_c < 1$ then $\mathcal{F}(\mathcal{R}_c) = \frac{\mathcal{R}_c}{1-\mathcal{R}_c}$. However if $\mathcal{R}_c > 1$, then $\mathcal{F}(\mathcal{R}_c) = \frac{\mathcal{R}_\infty}{1-\mathcal{R}_\infty}$ where $\mathcal{R}_\infty = (1-\mathcal{A})\mathcal{R}_c$ is the effective reproduction number at the end of the epidemic.

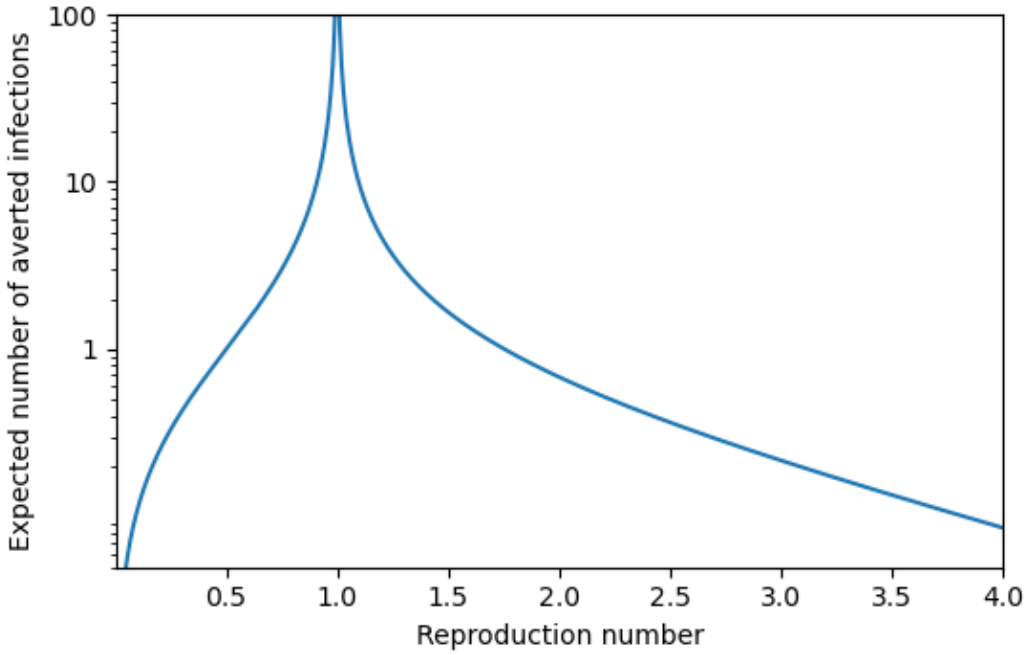

Figure 3 The expected number of averted infections, $\mathcal{F}(\mathcal{R}_c)$, due to a single infected individual isolating after infection and not transmitting, after accounting for overdetermination. The result is independent of the specific offspring distribution. This depends only on the average, $\mathcal{R}_c$. Note the divergence to ∞ as the reproduction number $\mathcal{R}_c \to 1$, and the decay to 0 as $\mathcal{R}_c$ decreases towards 0 or increases towards ∞.

Figure 3 shows a striking effect. Close to $\mathcal{R}_c = 1$ the value of $\mathcal{F}(\mathcal{R}_c)$ is very large, approaching infinity in the limit. So the expected number of infections averted if one newly-infected individual successfully isolates and avoids causing any infections is very large if $\mathcal{R}_c$ is close to 1. However, it becomes quite small if $\mathcal{R}_c$ is either large (because of overdetermination) or small (because there is little transmission). In practical terms, this analysis will be valid for large or small $\mathcal{R}_c$, but when $\mathcal{R}_c$ is close to 1, this analysis will break down if $\mathcal{F}(\mathcal{R}_c)$ is comparable to the population size. If $\mathcal{F}(\mathcal{R}_c)$ is less than 1% of the population size, we expect this to be a very good approximation.

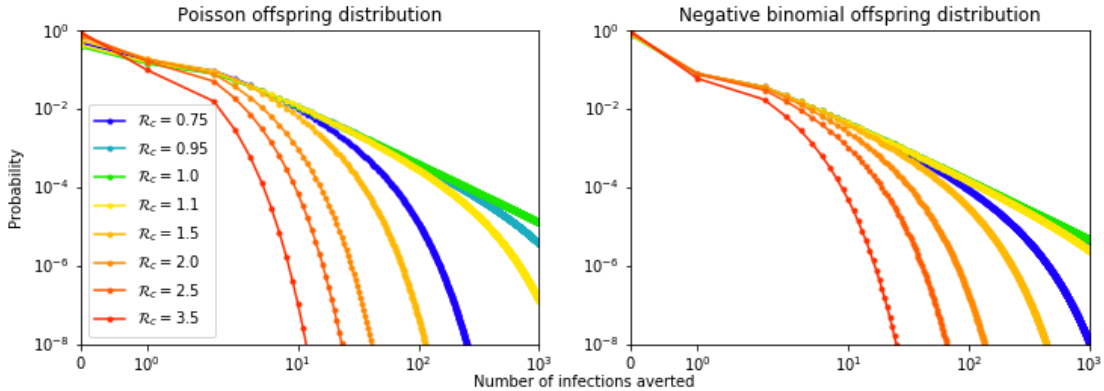

Figure 4: The distribution of the number of infections averted if a single individual isolates after infection. Calculated for Poisson and Negative binomial distribution (with dispersion parameter 0.16). For $\mathcal{R}_c = 1$ the result is a power law distribution. If $\mathcal{R}_c$ is larger or smaller than 1, the distribution falls off quicker. For a given $\mathcal{R}_c$, both distributions have the same mean, but the Negative binomial distribution results in a more heterogeneous outcome.

In addition to the expected number of infections averted, $\mathcal{F}(\mathcal{R}_c)$, it is possible to calculate the distribution of the number of infections averted. This is shown in Figure 4 for the Poisson and Negative binomial distributions. Although the averages $\mathcal{F}(\mathcal{R}_c)$ are the same, the distributions are different. In the Negative binomial case it is more common that either 0 infections are averted or a large number of infections are averted. As is often seen in systems at the critical threshold, when $\mathcal{R}_c = 1$ the distributions are given by powerlaws. This figure shows that although it is rare for the number of averted infections to be large, the large events are frequent enough to produce a large average, $\mathcal{F}(\mathcal{R}_c)$, near $\mathcal{R}_c = 1$.

If the individual does not isolate immediately and so is infectious for a short time, then obviously the benefit is smaller. If the expected number of transmissions is reduced to $\phi\mathcal{R}_c$ where $0 < \phi < 1$, then the expected number of infections averted is $(1-\phi)\mathcal{F}(\mathcal{R}_c)$.

Having analyzed the impact of a single individual who isolates following infection, let us now consider what happens if some nonzero fraction of the infected individuals isolate. If the number is not large enough to materially affect $\mathcal{R}_c$, then each isolation is effectively independent, and to find the expected benefit, we can simply multiply $\mathcal{F}(\mathcal{R}_c)$ by the number who isolate. However if a nonvanishing fraction isolate this will alter $\mathcal{R}_c$. Then we see an increasing marginal benefit of collective compliance until $\mathcal{R}_c$ reaches 1. Each additional

individual who isolates increases the effectiveness of those who have already isolated by reducing the level of overdetermination.

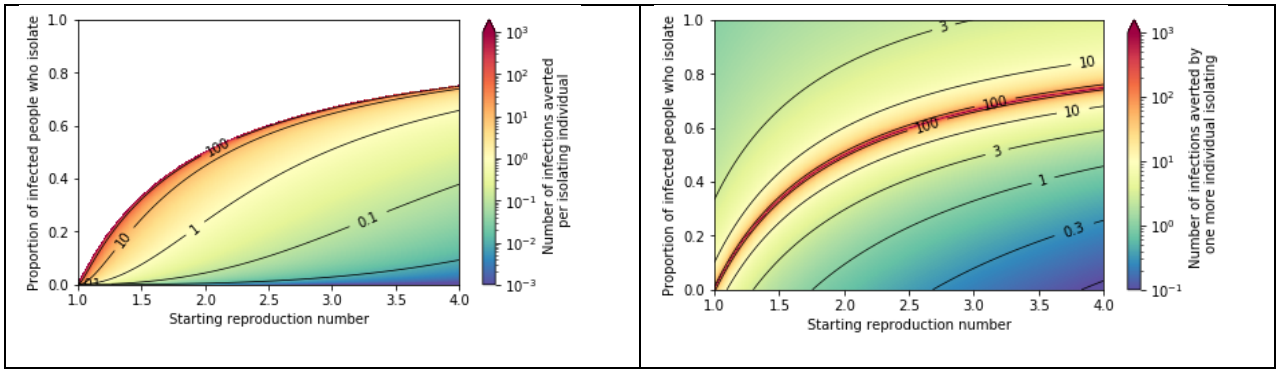

Figure 5: Measuring the effectiveness of many individuals isolating. ***Left*** *The average impact assuming that others interventions set a starting value of $\mathcal{R}_c$, but in addition to that an additional fraction of the infected individuals isolate. Calculations are done in the limit of an infinite population. As the fraction of infections isolating increases, the size of the epidemic decreases. Eventually a larger fraction isolating translates to a smaller absolute number isolating due to the smaller epidemic. When enough isolate to drive the resulting $\mathcal{R}_c$ close to 1 almost all infections are averted while only a small fraction of the total population isolate, so the number of infections averted per isolating individual diverges. Beyond this threshold (white region), epidemics are impossible and the number of infections averted per infected individual is infinite.* ***Right*** *The marginal benefit of one more infected individual isolating. That is, given the fraction isolating and the initial $\mathcal{R}_c$, this gives the number of infections averted if one additional infected individual who would not isolate is successfully identified and isolated.*

Figure 5 shows a similar trend to Figure 2. Namely, if a larger fraction of infected individuals isolate, the expected number of infections averted per isolating individual grows. So their combined impact grows faster than linearly in the fraction isolating. That is, we again see a synergistic effect. If $\mathcal{R}_c$ is large and very few infected individuals are isolating, then those who do have little impact. However as the fraction of infectious individuals who isolate increases, the average impact grows. As the system nears the epidemic threshold the average impact grows large, and the impact of each additional infected individual isolating diverges. Past the threshold where epidemics become impossible, the average impact is infinite (in the large-population limit) because only a small number of individuals isolate, but in doing so they eliminate the epidemic. In this region (where epidemics are not possible) if one infected individual who would not have isolated changes to isolating the marginal benefit begins to shrink as a larger fraction isolate (but very few infections occur).

Action taken to avoid infection

Above we considered individuals acting to prevent onwards transmission only after their infection. This has no direct benefit to the individual. Now we instead consider an individual acting to escape infection. This protects the individual and prevents onwards transmission. This would require the individual to either receive a vaccine or some other one-off treatment which prevents infection or to be taking long term actions such as a prophylactic medication to

avoid infection. The latter case requires sustained action, and so may have significant cost to the individual.

We first consider the impact of a single individual taking proactive actions to prevent her infection. We assume these actions are fully effective. The probability that she would be infected if she did not take those actions is equal to the fraction of the population infected in the epidemic (often called the *attack rate*). Assuming that the population is well-mixed and susceptibility is uniform across the population, the attack rate $\mathcal{A}$ depends only on the mean of the offspring distribution, not on any other details of the distribution. It can be calculated from the implicit *final size relation* (*19–21*)

$$\mathcal{A} = 1 - e^{-\mathcal{R}_c \mathcal{A}}$$

We can think of the attack rate as a function of $\mathcal{R}_c$, that is $\mathcal{A} = \mathcal{A}(\mathcal{R}_c)$. If a single individual takes perfectly effective measures to eliminate her probability of being infected, then with probability $\mathcal{A}(\mathcal{R}_c)$ she avoids her own infection. This also prevents transmissions from her to those who would otherwise have been her "descendants". So if she would have been infected, this reduces the expected number of additional cases by an amount $\mathcal{F}(\mathcal{R}_c)$. The total reduction in expected infections is $\mathcal{A}(\mathcal{R}_c)[1+\mathcal{F}(\mathcal{R}_c)]$. The first factor $\mathcal{A}(\mathcal{R}_c)$ represents the probability that the vaccine prevents an infection that would have happened. In the second factor, $1+\mathcal{F}(\mathcal{R}_c)$, the 1 represents the individual preventing her own infection, while $\mathcal{F}(\mathcal{R}_c)$ represents the infections averted due to the lack of onwards transmission.

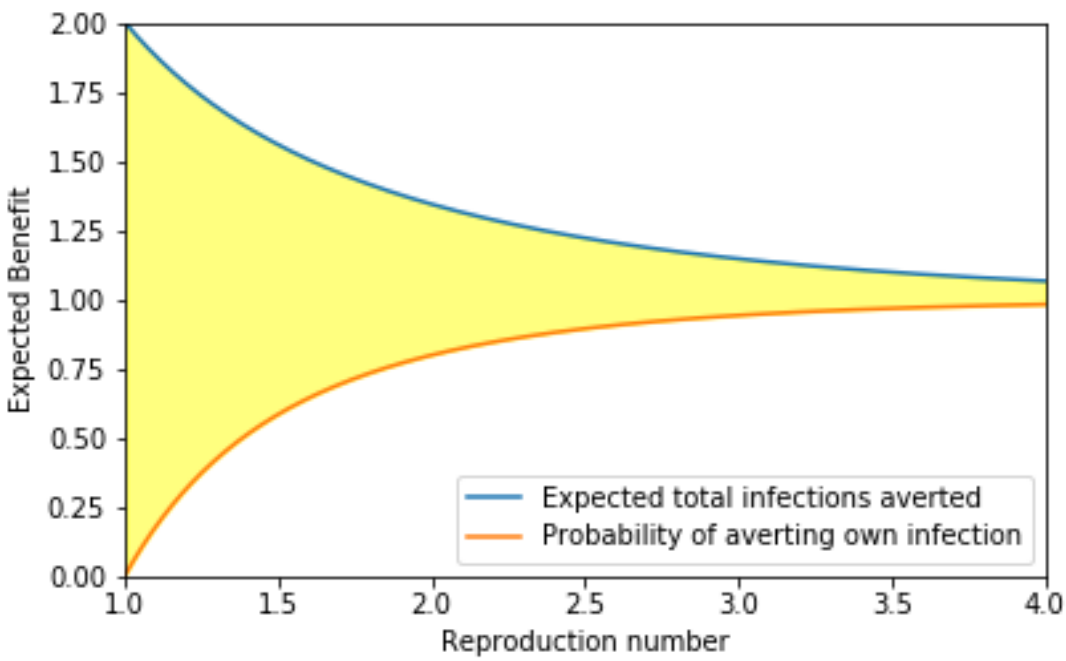

Figure 6: The expected number of infections averted by an individual taking sufficient action (such as a vaccine or prophylactic medication) to prevent her own infection (top blue curve). This is partly from the probability that she prevents her own infection [bottom orange curve, equal to $\mathcal{A}(\mathcal{R}_c)$*]. The difference between the two curves (shaded yellow) is* $\mathcal{F}(\mathcal{R}_c)\mathcal{A}(\mathcal{R}_c)$*,*

the expected number of descendants who would avoid infection due to her actions to avoid her own infection.

Figure 6 tells a surprising story. We consider an individual in the population who avoids infection through vaccine, prophylactic medication, or some other method. When she takes this action, we have no prior knowledge of whether she would be infected or otherwise. The probability her action prevents her own infection is equal to the probability she would have been infected without the action, in other words it is the attack rate $\mathcal{A}(\mathcal{R}_c)$, shown in the orange curve. This increases from 0 at $\mathcal{R}_c = 1$ to near 1 at $\mathcal{R}_c = 4$ (and approaches 1 as $\mathcal{R}_c$ increases further). Because $\mathcal{F}(\mathcal{R}_c)$ is small for large $\mathcal{R}_c$, even though the individual is very likely to prevent her own infection, the expected number of additional infections this avoids is very small. So the expected combined number of infections averted (including her own) is just above 1, with her own benefit constituting almost all of that.

For smaller $\mathcal{R}_c$ however, the probability that the action protects herself drops. At the same time, the expected number of additional infections averted if infected grows. The combined effect is larger for smaller $\mathcal{R}_c$. Remarkably, for values of $\mathcal{R}_c$ near 1, the expected number of additional infections averted approaches 2, even as the probability that any infections are averted goes to 0. A key additional observation is that if $\mathcal{R}_c > 1$, then on the expected number of infections prevented by each vaccine administered is greater than 1.

To make this clearer, consider the example of $\mathcal{R}_c = 1.01$. The probability a given vaccination directly protects the recipient is small, equal to $\mathcal{A} \approx 0.0197$. However, in those rare cases in which the recipient does become infected, the number of additional infections comes from a distribution as in Figure 4. From Figure 3, the expected number of additional infections would be $\mathcal{F}(\mathcal{R}_c) \approx 99.66$. The expected number of individuals whose infections are averted due to indirect protection from the vaccine is $\mathcal{A}\mathcal{F} \approx 1.9670$ (shaded part of Figure 6). The combined benefit is about $0.0197 + 1.9670 \approx 1.9867$ (top line of Figure 6).

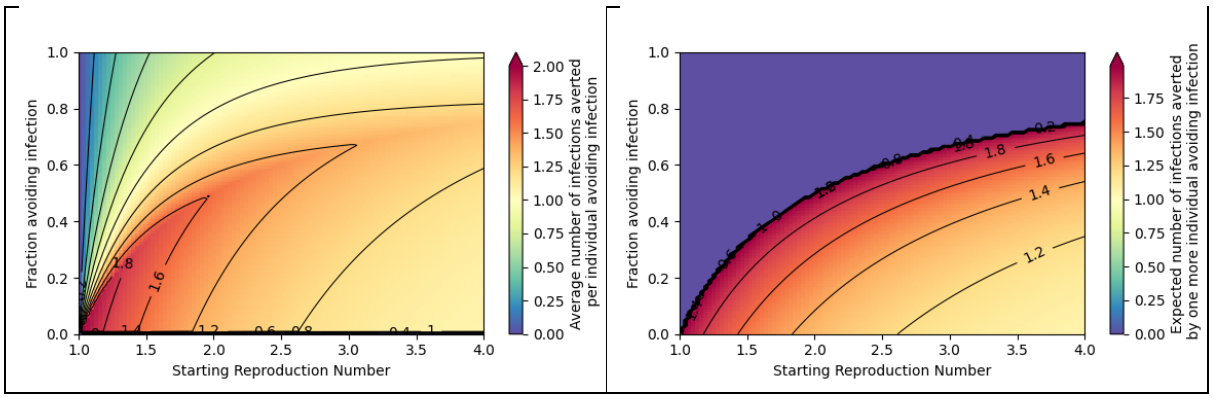

Figure 7: Measuring the impact of many individuals taking actions (such as a vaccine or prophylactic medication) to avoid infection. ***Left****: The number of infections averted over the entire population taking protective action. Until enough individuals take action to prevent an epidemic, the average is greater than 1, and the average increases until the threshold is*

reached after which it begins to fall. ***Right****: The marginal reduction in infections from one more individual changing behavior. The marginal benefit increases until the epidemic threshold is reached. Near the epidemic threshold, on average each additional individual taking action prevents about 2 infections. Once the epidemic threshold is crossed each additional individual has no net impact.*

In Figure 7 we see the impact of multiple individuals taking a vaccine or prophylactic medication or otherwise avoiding infection. A key observation is that as more individuals take self-protective actions, the marginal benefit of the next individual to act increases, until the epidemic threshold is reached and epidemics are eliminated. So the average benefit increases as more individuals act until the threshold is crossed. After this, the marginal benefit is 0, and so having more individuals act reduces the benefit averaged over those who act. Remarkably, until the epidemic threshold is reached, each individual that acts prevents (on average) greater than 1 infection.

Discussion

We have analyzed the impact of individual and collective behavior modifications related to several important public health measures to control spread of infectious disease, namely quarantine of arrivals, isolation of infected individuals, and the use of vaccines/prophylaxis to prevent infections. A common theme of our results is that there is a synergistic impact: as a larger proportion of the population adopt a protective behavior, the benefit created per individual changing behavior increases faster than linearly.

We now discuss some of the implications our observations have for the ethical policy decision-making.

Quarantine to prevent epidemic

Quarantines have shown that they can prevent the introduction of infection into a population. However, they are rarely perfectly effective and leaks have occurred in many. If we expect that the benefits of quarantine are almost entirely lost if the disease manages to successfully establish itself within the community, the strategy cannot allow even one single successful incursion (i.e., an introduction that results in sustained transmission).

With this standard, our analysis shows that any single individual's compliance with quarantine will only have a large impact where there is near perfect compliance amongst others. This is because, if an epidemic will happen anyways, then the impact of an additional introduction from a quarantine breach is generally minimal. The impact of a quarantine leak once disease spread is already established in the population is equivalent to having an infected individual who could isolate failing to isolate, which is discussed below.

It should be noted that if we assume that the community will be able to introduce burdensome interventions and would successfully eliminate the introduction, then the balance changes somewhat. In this case the cost of introductions (provided they occur during a period of elimination) is additive, and so each individual who does not comply with the intervention poses a separate risk to the population.

Isolation to avoid transmission
We find that isolation after infection can be a highly efficient intervention. The high efficiency is for two reasons.

- First, as a larger fraction of the infected individuals isolate, the total number infected drops, which tends to limit the number that will isolate. Once the *fraction* isolating becomes large enough, the drop in total cases is sufficient that the *number* who isolate also drops.
- Second the intervention only targets those who are infected, and only when they are infected. This means that those who are never infected are not burdened by the intervention (although they benefit from the compliance of infected individuals).

Isolation is most efficient on a per-individual basis when the reproduction number is close to 1, or if a large enough fraction is already isolating to reduce the reproduction to near 1. Near the epidemic threshold each individual who isolates prevents, on average, the infection of a large number of people (though still small compared to the population size). At the epidemic threshold, the distribution of number of infections prevented is a power-law (a straight line on a logarithmic plot). In these situations, large outbreaks are rare, but they are not so uncommon that we can ignore them. Their large size outweighs their rarity. In fact it turns out that the average outbreak size at the epidemic threshold is infinite.

Isolation of newly infected individals requires the ability to quickly identify infected individuals, perhaps through the use of rapid tests or effective contact tracing. In cases where there is asymptomatic or presymptomatic transmission, this will generally be more difficult (*22*, *23*)

The costs of this intervention are borne by individuals who do not experience the benefits of reduced transmission. So while it is highly efficient, the individuals who are protected from infection do not incur any costs.

Actions taken to avoid individual infection
An individual may avoid infection in several ways. If an effective vaccine or prophylactic medication is available, then he may be able to use these to avoid infection. Alternately, the use of personal protective equipment (PPE), or behavior changes can reduce the probability of infection. If the measures taken are less than 100% effective, the impact will be scaled by the corresponding factor.

These options to avoid infection may impose costs on the individual even though we do not know in advance whether he would ever be infected without them. This cost can be mitigated in part by changing the level of protection based on local prevalence, *e.g.*, by wearing PPE only when the risk of infection is high.

A key observation is that for an intervention such as a perfect vaccine, unless the control threshold is crossed and the epidemic is eliminated, every vaccine prevents on average more than one infection. Our analysis allows us to quantify how that benefit is distributed. Most of the benefit goes to the vaccine recipient if $\mathcal{R}_c$ is large. In contrast, if $\mathcal{R}_c$ is near 1, the protective measure has almost no impact on the individual's probability of infection, but it has a large indirect effect on others in those rare cases where an individual does become infected. This is because overdetermination is reduced in low transmission settings, so that the net expected number of infections averted approaches 2. For the example of $\mathcal{R}_c = 1.01$ shown in the results section, the probability that vaccine prevents infection of the recipient is about 2%. But when this happens, on average about 100 subsequent infections are prevented. So on average about $2\% \times 100 = 2$ infections are averted, but about 99% of the benefit is due to indirect protection of those not vaccinated.

Ethical Implications

Some of our results are relatively unsurprising: if compliance with a quarantine is low, then the burden on those who do comply is hard to justify if the goal is to prevent an epidemic. However, some of our other results are more unexpected. Specifically, when $\mathcal{R}_c$ is just a little above 1, individual actions have more impact than when $\mathcal{R}_c$ is well above 1.

In the case of isolating after infection, the expected benefit can be very large if $\mathcal{R}_c$ is close to 1. However that benefit is enjoyed by someone other than the one enduring the burden of compliance. In the case of perfect vaccination/prophylaxis the expected number of infections averted is always at least 1 if $\mathcal{R}_c > 1$. However, when $\mathcal{R}_c$ is close to 1 the benefit is almost entirely experienced by those not receiving the intervention. These observations provide quantitative data to inform ethical policy design for interventions in cases where the burden of compliance is borne by individuals who receive little or no benefit from compliance.

As a general rule, we see that these interventions are most effective on a per individual basis when the reproduction number is close to 1. When more individuals comply with an intervention, this acts to reduce the overall reproduction number, and the public health benefits per individual increase faster than linearly – each additional individual complying increases the average effectiveness of those already complying due to reductions in overdetermination.

This leads to interesting conceptual questions about how to attribute "credit" for the benefit that comes from compliance. Do all individuals who comply contribute equally to the total benefit? Or do the later ones play a larger role? This conceptual question becomes more practical when incentives and/or coercive measures are used to increase compliance. If incentives are given to those who comply, should they share equally? Should the last ones to comply (who might be considered to create the largest marginal benefit) be given a larger incentive? If coercive measures are used, should the punishment increase as compliance increases? Should it be calibrated to the average benefit of those who are complying or expected to comply?

Promoting collective compliance

Our observations lend support to the argument that for a highly infectious disease, if compliance is not high, then there is little ethical rationale for compliance on the basis that each individual makes little difference to overall population-level harm (although in this scenario vaccination/prophylaxis provides a significant direct individual benefit and we are ignoring health system capacity constraints in this analysis). Population benefits of collective compliance increase with greater compliance. One implication might be that public health messaging should promote high levels of compliance as a good to society rather than stigmatize those who fail to comply. It is perhaps not widely recognized by the public that large numbers of people acting together to reduce transmission of an infectious disease result in *synergistic* public health benefits. To the extent that individuals are aware of such patterns and act on this knowledge, this might reinforce behavior that improves epidemic control through higher compliance.

Our modeling illustrates that the public health consequences directly related to an individual's noncompliance with interventions is largest when $\mathcal{R}_c$ is near 1, for which overdetermination is rare and the epidemic might seem small and/or manageable. Communicating the benefits of collective action might help to improve public cooperation with control measures, especially by making people aware that in a highly susceptible population it is all the more important that people (continue to) contribute to control measures as they drive $\mathcal{R}_c$ close to 1, where someone might reasonably conclude his individual incentive to comply is low (as his risk of infection is negligible) (Figure 3 and Figure 6).

A key consequence of our results is that there is a disparity in who receives the benefits from individual actions. Near $\mathcal{R}_c = 1$ an individual's action produces much more public benefit than private benefit. This suggests that policies that mitigate the individual costs of compliance are likely to be an important part of ethical interventions.

Limitations and future work:

There are a number of limitations to this work that should be addressed in future analysis.

Mathematically, we have studied a relatively simple model of disease – using an SIR model without any incubating or asymptomatic stages. We treat the population as fully mixed, with all individuals equally susceptible and ignore variable risks of severe outcomes. We have focused our attention on measuring the benefit of an intervention – the costs should be explicitly measured as well to balance against those benefits. A more nuanced model will be needed to investigate issues related to more complex disease and population structure. Additionally, our focus has been on basic measures of impact, namely the prevention of an epidemic and reduction in infections. We have not considered the fact that infected individuals may occupy health care resources that prevent others from accessing care. Nor have we considered issues such as the importance of minimizing the epidemic peak as a separate issue from minimizing total infections.

We have assumed that the background conditions are constant. The interventions are assumed to be unchanging, while in reality they may change in response to the epidemic dynamics. In a real-world scenario we might expect that treatment methods might improve

over time, thus a delayed infection might be less severe than an earlier infection. We might also anticipate that a vaccine may be developed, thus a delayed infection may actually end up being a prevented infection or a less-severe infection. On the other hand, we have also assumed that the disease remains unchanged. In reality we expect a more transmissible variant may emerge, and so delaying early transmission may come at the cost of a significantly larger later epidemic. Similarly if the disease spreads best in winter, efforts to reduce transmission in summer may lead to an enhanced winter wave.

Materials and Methods

In this section we derive our mathematical approaches and perform some of the technical analysis of the model. We begin by briefly describing the model assumptions and providing two small examples demonstrating key features of stochastic infection spread. We then derive the mathematical approaches. Finally we perform some of the analysis, ending with a rigorous derivation of the observation that as $\mathcal{R}_c \to 1^+$ the expected number of infections prevented by a vaccine approaches 2.

Model Assumptions

We assume a stochastic model of infection spread, based on the standard Susceptible-Infected-Recovered (SIR) model (*24*, *25*). Each infected individual potentially transmits to k others, chosen uniformly at random from the population, where k is chosen from some distribution. If the recipient of a transmission is susceptible, an infection event occurs. If not, the transmission has no effect. The distribution of the number of transmissions k is known as the *offspring distribution*, with p_k denoting the probability of k transmissions. After transmitting to its *offspring* the infected individual recovers with immunity.

If there are no interventions in place, then the average of the offspring distribution is the *basic reproduction number* $\mathcal{R}_0$. If interventions are in place, then the average of the offspring distribution after accounting for those interventions is the *reproduction number under control* $\mathcal{R}_c$. We assume that $\mathcal{R}_c$ remains constant for the duration of the epidemic. Because transmissions have no effect when the recipient has already been infected, we also introduce the *effective reproduction number* $\mathcal{R}_{\text{eff}}(t)$ which measures the average number of successful transmissions from an infected individual. This equals $\mathcal{R}_c$ times the proportion of the population that is susceptible. $\mathcal{R}_{\text{eff}}$ accounts for both the intervention and the immunity that the population has developed.

It is known that the probability an epidemic becomes established is sensitive to the frequency of superspreading, the tendency that a small fraction of the infected individuals cause a large fraction of the transmissions (*14*, *16–18*). To investigate the significance of superspreading in our analysis, we consider the impact of a Poisson offspring distribution and a Negative Binomial distribution. The Poisson distribution is parametrized by a single variable, the mean $\mathcal{R}_c$, and individuals who transmit significantly more than average are negligibly rare. The Negative Binomial distribution is parametrized by the mean $\mathcal{R}_c$ and a dispersion

parameter, which we take to be 0.16, which includes significant superspreading. This is consistent with estimates for SARS (*17*) and allows us to clearly show what impact (if any) superspreading has. More recent estimates for COVID-19 suggest that the dispersion parameter should be larger (*16*); however our goal is to investigate the qualitative impact of superspreading rather than making exact predictions for a particular disease.

Example Outbreaks

In Figure 8, we show transmission chains in two outbreaks in populations of 50 individuals. The two populations have $\mathcal{R}_c = 1.5$ and $\mathcal{R}_c = 2.5$. Each infected individual transmits independently to each of the other 49 individuals with probability $\mathcal{R}_c / 49$, resulting in a distribution that is approximately Poisson with mean $\mathcal{R}_c$. In both cases, the outbreak successfully establishes and only terminates because eventually a significant number of transmissions go to previously infected individuals. In a larger population, those blocked transmissions would have gone to other susceptible individuals, resulting in a large-scale outbreak, that is, an *epidemic*.

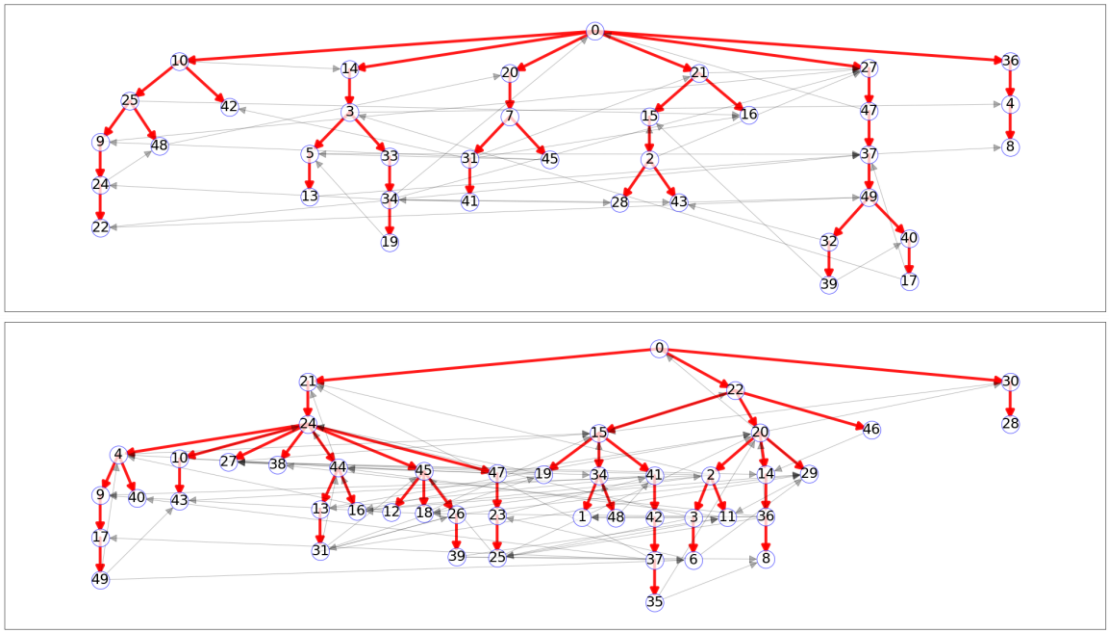

Figure 8 Two sample outbreaks with $\mathcal{R}_c = 1.5$ (top) and $\mathcal{R}_c = 2.5$ (bottom), each starting from a single infection in a population of 50 individuals. Only the eventually infected nodes are shown. The red lines denote successful transmissions while the gray lines denote transmissions from an infectious individual to one who had already been infected. Note that given the transmissions in the $\mathcal{R}_c = 1.5$ case, had individual 36 (second row on right) been the initial infection, only individuals 36, 4, and 8 would have been infected. Additionally, had individual 49 (fifth row on right) isolated after infection (or been effectively protected from infection), this would have prevented the infections of 32, 39, 40, and 17. In the $\mathcal{R}_c = 2.5$ case, there are many more (potential) transmission chains: more infections occur,

transmission chains tend to be longer, and removing one individual tends to protect fewer others.

Figure 8 shows that although many epidemics spread far and are limited only by the population size, outbreaks starting from some individuals would not spread far. For example, in the top plot if the infection introduced in individual 42 (second node in third level) rather than individual 0 or if in the bottom plot infection started with 28 (last node in third level), these would not lead to long transmission chains. However, many of the other individuals would spark large-scale transmission through the population. So, if multiple introductions occur, there is a likelihood that more than one of them is sufficient to spark an epidemic. This is more likely for larger $\mathcal{R}_c$. This illustrates how *epidemics* can be overdetermined, which plays an important role in our analysis.

Additionally, we see that individual *cases of infection* may or may not be overdetermined. If we blocked transmission from some individuals it would provide effective protection to others. For example, in the top plot preventing transmissions from individual 49 (lower right) would be sufficient to prevent the infection of 32, 39, 40, and 17. However, many infections are overdetermined because there are alternate transmission routes. For example, in the top plot preventing transmissions from 10 (first node in second level) would prevent the infection of 25, but all other descendants would eventually be infected through alternate chains of transmission. The existence of multiple transmission chains to the same individual becomes more likely for larger $\mathcal{R}_c$.

In a practical setting we do not know *a priori* which transmissions would occur. Based on the offspring distribution, we can calculate the probability that an introduced infection results in an epidemic. We can also calculate the distribution of the number of infections averted by one individual's behavior change. However, due to the stochasticity inherent in the system, for a specific infected individual introduced to a population, we cannot know in advance whether he would spark an epidemic, or in an ongoing epidemic, we do not know exactly how many infections another individual might avert by changing her behavior. Thus, our analysis will focus on the expected (*i.e.*, the average) impact over many realizations.

The key quantities we focus on are:

- The expected impact of a single infected individual acting alone or multiple infected individuals acting together to isolate prior to entering a community to prevent an epidemic.
- The expected impact of a single individual acting alone or multiple individuals acting together to reduce the total number of infections occurring in an epidemic.

Mathematical Methods

We now build up the mathematical methods used to analyze the sort of outbreaks that can occur. We will assume throughout that the population size is very large. Under this assumption two typical outcomes occur in large populations: either an outbreak remains small

and dies out quickly or it becomes an epidemic that grows until it is limited by the population size. If $\mathcal{R}_c < 1$ only small outbreaks occur. If $\mathcal{R}_c > 1$ large epidemics can occur but small outbreaks are still possible. The size distribution of the *number* infected in small outbreaks is independent of the population size but does depend on the offspring distribution. In contrast, in a large epidemic the *proportion* infected is independent of the population size (so the number infected is proportional to the population size) and the proportion is independent of the offspring distribution.

Probability Generating Functions

In this section we briefly introduce some properties of probability generating functions (PGFs). More complete details can be found in (*18*, *26*).

The number of outgoing edges from a particular node is chosen from the offspring distribution. If p_0, p_1, … represent the probability of zero, one, … offspring, then the PGF of the distribution (*26*) is defined to be

$$\mu(x) = \sum_{i=0}^{\infty} p_i x^i$$

We have considered a Poisson distribution with mean $\mathcal{R}_c$ for which $\mu(x) = \exp(-\mathcal{R}_c(1-x))$ and a negative binomial distribution with a dispersion parameter of $\kappa = 0.16$ for which $\mu(x) = \left(\frac{\kappa}{\kappa + \mathcal{R}_c - \mathcal{R}_c x}\right)^{\kappa}$, but any probability distribution can be chosen based on the characteristics of the disease. The mean of the offspring distribution satisfies $\mathcal{R}_c = \sum_{i=0}^{\infty} i p_i = \mu'(1)$.

In many cases it may be difficult to find the coefficients of a PGF $\mu(x)$ directly, but it is possible to calculate the values of $\mu(x)$ to high precision on the unit circle in the complex plane. Then we can use a Cauchy integral to find arbitrary coefficients of $\mu(x)$ [see section A.1 of (*18*)] (in fact if we parametrize the unit circle by the angle θ, then $\mu(x)$ becomes a complex-valued Fourier Series in θ, and the Cauchy integral becomes the formula for the coefficients of a Fourier Series). In the context of disease spread, if $\Omega_\infty(x)$ is the PGF for the final size distribution when the reproduction number is less than 1, using approaches shown below, it is possible to calculate $\Omega_\infty(x)$ at arbitrary values of x. Then this approach can be used to find the coefficients of the series expansion of $\Omega_\infty(x)$ (*18*).

The final size relation

We have assumed homogeneous susceptibility in a well-mixed population. It is well-known that the final size of epidemics under these assumptions is independent of the offspring distribution in the large population limit (it depends only on the average of the distribution). Here we briefly derive the final size relation following (*19*, *21*) [see also (*20*)]. We let $\mathcal{A}(\mathcal{R}_c)$ denote the expected proportion of the population infected in an epidemic. The total

number of transmissions that occur in the epidemic is well-approximated by $\mathcal{R}_c \cdot \mathcal{A}(\mathcal{R}_c) \cdot N$ where N is the population size. On average each member of the population thus receives $\mathcal{R}_c \cdot \mathcal{A}(\mathcal{R}_c)$ transmissions. In the large population limit, we can reasonably assume independence of these events. This means that the number of transmissions received is Poisson-distributed with mean $\mathcal{R}_c$. The probability of receiving no transmissions is $e^{-\mathcal{R}_c \cdot \mathcal{A}(\mathcal{R}_c)}$. However, because this is the probability of not being infected it also equals the proportion of the population that remains susceptible. We arrive at the final size relation

$$\mathcal{A}(\mathcal{R}_c) = 1 - e^{-\mathcal{R}_c \mathcal{A}(\mathcal{R}_c)}$$

This can be solved iteratively by setting $\mathcal{A}_{g+1} = 1 - e^{-\mathcal{R}_c \mathcal{A}_g}$ with $\mathcal{A}_0 = 1$.

Calculating epidemic probability

To calculate the epidemic probability, we consider a process known as the Galton-Watson process (or birth-death process). Let α denote the probability that a given individual in a branching process has a finite number of descendants. Then $1-\alpha$ is the probability of an infinite number of descendants. The number of descendants is finite exactly when every single offspring has a finite number of descendants. Since each offspring has a finite number of descendants also with probability α we find

$$\alpha = \sum_i p_i \alpha^i = \mu(\alpha)$$

Thus, the probability of an infinite number of descendants can be calculated by finding the roots of $\alpha = \mu(\alpha)$ where $\mu(x)$ is the PGF. One root is always $\alpha = 1$ corresponding to never having an infinite number of descendants, but if there is another root, it will lie between 0 and 1, and it is the correct root to choose. This other root exists when $\mathcal{R}_c > 1$.

In practice we can find α by setting $\alpha_0 = 0$ and letting $\alpha_{g+1} = \mu(\alpha_g)$, until the values found for α_g converge. In this approach, α_g can be interpreted as the probability that the outbreak terminates by generation g.

Calculating the impact of isolation after infection.

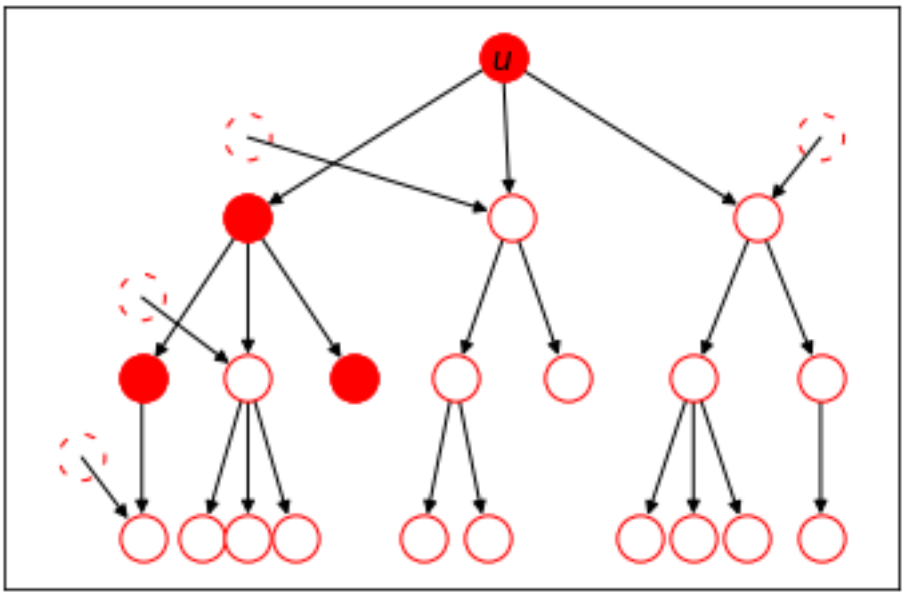

Figure 9: A schematic showing the descendants of u with filled circles showing those whose infections would be averted if u's transmissions were prevented by u isolating. Dashed hollow circles denote infected individuals that are not descendants of u who provide additional transmission paths to some descendants of u. Hollow circles with solid edges denote the individuals whose infection is overdetermined.

To determine the expected impact of isolation of u after becoming infected, we need to calculate the expected number of descendants an individual who would not be infected through some transmission path not through u, as shown in Figure 9. To determine this, we consider the *residual offspring distribution*, that is the distribution of the number of direct offspring who would not be reachable along any other transmission chain. If an individual would cause i transmissions, a fraction $\mathcal{A}(\mathcal{R}_c)$ of them go to individuals who would otherwise be infected through a different chain, as would all their descendants.

The average of the residual offspring distribution is $\mathcal{R}_\infty = [1-\mathcal{A}(\mathcal{R}_c)]\mathcal{R}_c$, which is the initial reproduction number multiplied by the fraction who remain susceptible at the end. This is the reproduction number at the end of the epidemic. If $\mathcal{R}_c \le 1$, then $\mathcal{A}(\mathcal{R}_c) = 0$ and we find $\mathcal{R}_\infty = \mathcal{R}_c$, while if $\mathcal{R}_c > 1$, then $\mathcal{R}_\infty$ is less than 1. The expected number of infections averted among those who are reachable from a path of g generations from u ends up being $\mathcal{R}_\infty^g$. Summing this over all g, we find that the initial infected individual's isolation prevents

$$\mathcal{F}(\mathcal{R}_c) = \sum_{g=1}^{\infty} \mathcal{R}_\infty^g = \frac{\mathcal{R}_\infty}{1-\mathcal{R}_\infty}$$

infections, where $\mathcal{R}_\infty = [1-\mathcal{A}(\mathcal{R}_c)]\mathcal{R}_c$ is the effective reproduction number at the end of transmission in a population having reproduction number under control of $\mathcal{R}_c$.Note that when $\mathcal{R}_\infty$ is close to 1, this is a large number. This happens when $\mathcal{R}_c$ is close to 1.

To go further, we can calculate the distribution of the total number of infections averted. The PGF for the residual distribution is $\mu(\mathcal{R}_\infty x + 1 - \mathcal{R}_\infty)$. Following methods derived in (*18*), the PGF for the size distribution of the number of infections averted is $\Omega_\infty(x) = x\mu(1-\mathcal{R}_\infty + \mathcal{R}_\infty \Omega_\infty(x))$. This can be calculated for arbitrary values of x by iterating $\Omega_{g+1}(x) = x\mu(1-\mathcal{R}_\infty + \mathcal{R}_\infty \Omega_g(x))$, starting with $\Omega_0(x) = 1$. Iterating for values of x on the complex unit circle until the results converge allows us to calculate the individual coefficients of $\Omega_\infty(x)$ through a Cauchy integral, from which we know the distribution of the total number of infections averted.

To calculate the impact of multiple individuals isolating, we note that if a fraction ρ of the population isolate after infection, then the offspring distribution is modified. With probability ρ an infected individual isolates and causes no infections, while with probability $1-\rho$ they cause a number of infections chosen from the original distribution. This means that $\mathcal{R}_c$ is effectively multiplied by $1-\rho$. We can redo the calculations for final size and individual impact above using $(1-\rho)\mathcal{R}_c$ instead of $\mathcal{R}_c$.

Calculating the impact of avoiding infection

If an individual either gets a 100% effective vaccine or prophylactic medication, she can avoid infection. Doing so, she prevents her own infection with probability $\mathcal{A}(\mathcal{R}_c)$, as it has no effect if she would have avoided infection anyways. When she prevents her own infection, the total number of infections averted due to this is $1+\mathcal{F}(\mathcal{R}_c)$ where the 1 accounts for protection of herself and $\mathcal{F}(\mathcal{R}_c)$ accounts for protection of others. Combining this reduction with the probability that the reduction occurs, we find that the expected reduction in infections de to a single individual taking measures to prevent her own infection is $\mathcal{A}(\mathcal{R}_c)[1+\mathcal{F}(\mathcal{R}_c)]$.

To calculate the impact of multiple individuals getting vaccinated or otherwise avoiding infection, we note that if a fraction ρ of the population is immune to infection, then the offspring distribution is modified. With probability ρ a random transmission is blocked by the protection to the recipient. This means that $\mathcal{R}_c$ is effectively multiplied by $1-\rho$, though for a different reason than before. We can redo the calculations for final size and individual impact above using $(1-\rho)\mathcal{R}_c$ instead of $\mathcal{R}_c$.

We now derive the number of infections averted for the limit approaching $\mathcal{R}_c = 1$ from above, showing that it matches the apparent limit shown by the numerics. We need to calculate $\lim_{\mathcal{R}_c \to 1^+} \mathcal{A}(\mathcal{R}_c)(1+\mathcal{F}(\mathcal{R}_c))$ where $\mathcal{A}(\mathcal{R}_c) = 1 - e^{-\mathcal{R}_c \mathcal{A}}$ and $\mathcal{F}(\mathcal{R}_c) = \mathcal{R}_\infty / (1-\mathcal{R}_\infty) = [1-\mathcal{A}(\mathcal{R}_c)]\mathcal{R}_c / (1-[1-\mathcal{A}(\mathcal{R}_c)]\mathcal{R}_c)$. To find this value, we invert the relationship between $\mathcal{A}$ and $\mathcal{R}_c$, writing $\mathcal{R}_c(\mathcal{A}) = -\frac{1}{\mathcal{A}}\ln(1-\mathcal{A})$. Then $\mathcal{F}(\mathcal{A}) = [1-\mathcal{A}]\mathcal{R}_c(\mathcal{A}) / (1-[1-\mathcal{A}]\mathcal{R}_c(\mathcal{A}))$. We now take the limit

$\lim_{\mathcal{A}\to 0^+} \mathcal{A}\left(1+\frac{(1-\mathcal{A})\mathcal{R}_c(\mathcal{A})}{1-(1-\mathcal{A})\mathcal{R}_c(\mathcal{A})}\right)$. Substituting for $\mathcal{R}_c(\mathcal{A})$ and noting that the 1 inside the parenthesis will not matter in the limit, we get

$$\lim_{\mathcal{A}\to 0^+} \mathcal{A}\left(1+\frac{(1-\mathcal{A})\mathcal{R}_c(\mathcal{A})}{1-(1-\mathcal{A})\mathcal{R}_c(\mathcal{A})}\right) = \lim_{\mathcal{A}\to 0^+} \frac{\mathcal{A}(1-\mathcal{A})\ln(1-\mathcal{A})}{-\mathcal{A}-(1-\mathcal{A})\ln(1-\mathcal{A})}$$

Taking one round of L'Hopital's rule yields

$$\lim_{\mathcal{A}\to 0^+} \frac{\mathcal{A}(1-\mathcal{A})\ln(1-\mathcal{A})}{-\mathcal{A}-(1-\mathcal{A})\ln(1-\mathcal{A})} = \lim_{\mathcal{A}\to 0^+} \frac{(1-2\mathcal{A})\ln(1-\mathcal{A})-\mathcal{A}}{\ln(1-\mathcal{A})} = 1-\lim_{\mathcal{A}\to 0^+} \frac{\mathcal{A}}{\ln(1-\mathcal{A})}$$

Applying L'Hopital's rule one again yields

$$1-\lim_{\mathcal{A}\to 0^+} \frac{\mathcal{A}}{\ln(1-\mathcal{A})} = 1-\lim_{\mathcal{A}\to 0^+} \frac{1}{-\frac{1}{1-\mathcal{A}}} = 1+\lim_{\mathcal{A}\to 0^+}(1-\mathcal{A})$$

which is 2.

References

1. N. Arinaminpathy, J. Savulescu, A. R. Mclean, Effective use of a limited antiviral stockpile for pandemic influenza. *J. Bioeth. Inq.* **6** (2009), doi:10.1007/s11673-009-9164-3.
2. J. A. Lewnard, N. C. Lo, Scientific and ethical basis for social-distancing interventions against COVID-19. *Lancet Infect. Dis.* **20** (2020), , doi:10.1016/S1473-3099(20)30190-0.
3. A. Brodeur, D. Gray, A. Islam, S. Bhuiyan, A literature review of the economics of COVID-19. *J. Econ. Surv.* (2021), doi:10.1111/joes.12423.
4. M. I. Meltzer, N. J. Cox, K. Fukuda, The economic impact of pandemic influenza in the United States: Priorities for intervention. *Emerg. Infect. Dis.* **5** (1999), doi:10.3201/eid0505.990507.
5. World Health Organisation, Ethical considerations in developing a public health response to pandemic influenza. *World Heal. Organ.* (2007) (available at www.who.int/csr/resources/publications/).
6. L. Rosenbaum, Facing Covid-19 in Italy — Ethics, Logistics, and Therapeutics on the Epidemic's Front Line. *N. Engl. J. Med.* **382** (2020), doi:10.1056/nejmp2005492.
7. P. Angelos, Surgeons, Ethics, and COVID-19: Early Lessons Learned. *J. Am. Coll. Surg.* **230** (2020), doi:10.1016/j.jamcollsurg.2020.03.028.
8. E. J. Emanuel, G. Persad, R. Upshur, B. Thome, M. Parker, A. Glickman, C. Zhang, C. Boyle, M. Smith, J. P. Phillips, Fair Allocation of Scarce Medical Resources in the Time of Covid-19. *N. Engl. J. Med.* **382** (2020), doi:10.1056/nejmsb2005114.
9. J. M. Marron, S. Joffe, R. Jagsi, R. A. Spence, F. J. Hlubocky, Ethics and Resource Scarcity: ASCO Recommendations for the Oncology Community during the COVID-19 Pandemic. *J. Clin. Oncol.* **38** (2020), doi:10.1200/JCO.20.00960.
10. D. S. Silva, M. J. Smith, Social distancing, social justice, and risk during the COVID-19 pandemic. *Can. J. Public Heal.* **111** (2020), doi:10.17269/s41997-020-00354-x.
11. A. J. London, J. Kimmelman, Against pandemic research exceptionalism. *Science (80-*

.). **368** (2020), , doi:10.1126/SCIENCE.ABC1731.
12. E. Jamrozik, M. J. Selgelid, COVID-19 human challenge studies: ethical issues. *Lancet Infect. Dis.* **20** (2020), , doi:10.1016/S1473-3099(20)30438-2.
13. D. C. Adam, P. Wu, J. Y. Wong, E. H. Y. Lau, T. K. Tsang, S. Cauchemez, G. M. Leung, B. J. Cowling, Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong. *Nat. Med.* **26** (2020), doi:10.1038/s41591-020-1092-0.
14. B. M. Althouse, E. A. Wenger, J. C. Miller, S. V. Scarpino, A. Allard, L. Hebert-Dufresne, H. Hu, Superspreading events in the transmission dynamics of SARS-CoV-2: Opportunities for interventions and control. *PLoS Biol.* **18** (2020), doi:10.1371/journal.pbio.3000897.
15. Severe Acute Respiratory Syndrome - Singapore 2003. *Morb. Mortal. Wkly. Rep.* **52**, 405–411 (2003).
16. Z. Susswein, S. Bansal, Characterizing superspreading of SARS-CoV-2: From mechanism to measurement. *medRxiv* (2020).
17. J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, W. M. Getz, Superspreading and the effect of individual variation on disease emergence. *Nature*. **438** (2005), doi:10.1038/nature04153.
18. J. C. Miller, A primer on the use of probability generating functions in infectious disease modeling. *Infect. Dis. Model.* **3**, 192–248 (2018).
19. Diekmann, J. Heesterbeek, *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation - O. Diekmann, J. A. P. Heesterbeek* (2000).
20. J. Ma, D. J. D. Earn, Generality of the final size formula for an epidemic of a newly invading infectious disease. *Bull. Math. Biol.* **68** (2006), , doi:10.1007/s11538-005-9047-7.
21. J. C. Miller, A Note on the Derivation of Epidemic Final Sizes. *Bull. Math. Biol.* **74**, 2125–2141 (2012).
22. M. E. Kretzschmar, G. Rozhnova, M. C. J. Bootsma, M. van Boven, J. H. H. M. van de Wijgert, M. J. M. Bonten, Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. *Lancet Public Heal.* **5** (2020), doi:10.1016/S2468-2667(20)30157-2.
23. C. Fraser, S. Riley, R. M. Anderson, N. M. Ferguson, Factors that make an infectious disease outbreak controllable. *Proc. Natl. Acad. Sci. U. S. A.* **101** (2004), doi:10.1073/pnas.0307506101.
24. R. M. Anderson, R. M. May, *Infectious diseases of humans: dynamics and control* (Oxford University Press, 1991).
25. W. O. Kermack, A. G. McKendrick, A contribution to the mathematical theory of epidemics. *Proc. R. Soc. London. Ser. A, Contain. Pap. a Math. Phys. Character*. **115**, 700–721 (1927).
26. H. S. Wilf, *generatingfunctionology, third edition* (2005).
27. D. Ludwig, Final size distribution for epidemics. *Math. Biosci.* **23** (1975), doi:10.1016/0025-5564(75)90119-4.

Acknowledgments

We thank Anja C Slim for sorting out the limit $\lim_{\mathcal{R}_c \to 1^+} \mathcal{A}(1+\mathcal{F})$.

Funding:

- Australian Mathematical Sciences Institute summer research fellowship (DR, JCM).
- CaRE grant scheme from La Trobe University's School of Engineering and Mathematical Sciences (DR, JCM)
- The Wellcome Trust [203132] and [221719] (EJ)

Author contributions:

Conceptualization: DR, GH, JCM, MS, EJ
Methodology: JCM & DR
Investigation: DR & JCM
Figure Generation: DR & JCM
Writing—original draft: DR & JCM
Writing—review & editing: EJ, DR, GH, MS, JCM