Fundamental limitations of case isolation

Jonas Hansson* Alain Govaert* Richard Pates*
Emma Tegling* Kristian Soltesz*,**

* Lund University, Dept. of Automatic Control, Lund, Sweden

** Correspondence: kristian@control.lth.se

Abstract: Case isolation is a strategy with the potential to curb infectious disease epidemics. Expressions for the stability boundary of a case isolation scheme defined through the proportion of the infectious population that it isolates with a given delay have recently been established. Here we quantify how this stability boundary moves when heterogeneity of the inter-individual contact network increases, and explain the underlying mechanism through insightful examples.

Keywords: Fundamental limitations, network dynamics, case isolation, COVID-19

1. INTRODUCTION

1.1 Background

Early in an epidemic of a previously unknown disease, transmitted through a previously unknown pathogen, it can happen that existing pharmaceutical treatments are not applicable or efficient. This was the case with COVID-19, caused by the pathogen SARS-CoV-2. In early 2020 countries across the world experienced exponential increases in detected SARS-CoV cases, hospitalizations and ICU admittance due to COVID-19, and ultimately related deaths. While statistics are still debated one year later, as exemplified through Oliver (2021), it stood clear early on that pharmaceutical treatment and clinical care would not alone suffice to curb the epidemic outbreak.

The exponential increase of detected SARS-CoV-2 cases, and particularly of COVID-19 related deaths, during early 2020 resulted in governments implementing schemes to reduce the number of probable transmission events. Such schemes are within infectious disease epidemiology referred to as non-pharmaceutical interventions (NPIs for short), and can be partitioned into two categories:

- Recommendations or legislation aimed at decreasing inter-individual contact rates;
- Schemes for case isolation through testing, and possibly contact tracing.

The former includes, for example, discouraging unnecessary in-person interaction, banning large gatherings, closing schools or imposing societal lockdowns. The effectiveness of such interventions is still poorly understood. In particular, we have illustrated a lack of practical identifiability of individual NPI effectiveness from data obtained during the first European COVID-19 wave in Soltesz et al. (2020). This was the case despite a simplistic model formulation and optimistic assumptions on uncertainty in data.

In light of the above, case isolation schemes have gained attention. In contrast to NPIs aimed at reducing contact rates in general, the case isolation schemes target reduction of contacts involving infectious individuals, as schematically illustrated in Fig. 1. If successfully implemented, they

Fig. 1. Schematic illustration of a case isolation scheme. The infectious proportion of the population I is partitioned into one isolated (quarantined) proportion Q , and one proportion I_{mix} . The proportion I_{mix} interacts with the susceptible proportion S , and this interaction is governed by the mixing parameter β . The transition rate from the infectious I to the removed (recovered or deceased) proportion R is γ . We will focus on analysing how the design of the case isolation scheme that partitions I into I_{mix} and Q affects the evolution of the epidemic.

could be the key to an open society where lockdowns are replaced by regular testing. Under such schemes, a positive test results in case isolation. The effectiveness of the scheme can then be further increased by complementing systematic testing with contact tracing.

Notably, such a scheme has been demonstrated effective in, for example, New Zealand, as described by Robert (2020). At times when the case isolation scheme was deemed insufficient, New Zealand authorities instead enforced strict societal lockdowns.

The question we consider here is under what circumstances a case isolation scheme alone is sufficient to stabilize the epidemiological trajectory. We first derive criteria for an idealized case of a homogeneously interacting population, and then investigate how interaction heterogeneities affect these criteria.

1.2 Preliminaries

We consider a closed population. This is an optimistic approximation, since there will be no import cases within the model. The population is partitioned into susceptibles, infectious and removed (recovered and permanently immune or deceased) proportions. Here we will not distinguish between being infected and being infectious. However, the

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

results we are about to present can easily be extended to cater for this distinction.

If the proportion of infectious goes to zero, the strategy has been successful. The question we consider is how large the infectious proportion can grow, relative to the susceptible, in order for a case isolation strategy to remain effective.

To answer this, we consider the setting early in an epidemic, where the proportion of susceptibles is much larger than the infectious and removed, respectively. This allows us to disregard herd immunity effects. Another way to put it is that in an interaction between two individuals where one is infectious, the other is susceptible (with probability one).

Furthermore, we realistically assume the considered time window for our model to be short compared to the typical duration of immunity. This means that we neglect any flow from the removed to the susceptible sub-populations. As with distinguishing between infected and infectious individuals, introducing such reflux dynamics into the model we propose is straightforward, should the considered disease differ in this aspect from SARS-CoV-2.

We will focus on criteria for stabilization of the epidemic trajectory. We will refer to a case isolation scheme as stabilizing if it eventually empties the infectious sub-population, without being aided by the herd immunity effect as introduced in Topley and Wilson (1923).

Note that this stability condition does not specify *performance* in the sense that it allows for an arbitrarily large infectious sub-population in the transient during which the infectious sub-population is emptied. As such, this stability condition constitutes a bare minimum. Any practically feasible strategy would need to fulfill it with some performance margin, as further elaborated in Pates et al. (2021).

Finally, we need to formalize our case isolation scheme. Here we will utilize a simple yet versatile model that quantifies the proportion of those infected on a given day that is subsequently isolated by the scheme T_{delay} days later.

This can be readily re-parameterized into, for example, the frequency with which individuals are tested, compliance to the scheme, logistic delays and specificity of the employed test for infection.

2. CASE ISOLATION IN HOMOGENEOUS POPULATIONS

2.1 Fundamental limitations

Assuming that the population is large, so that quantization effects become negligible, we can model the trajectory of the epidemiological state through

$$\frac{d}{dt} \begin{bmatrix} S \\ I \\ R \end{bmatrix} = \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} \beta S(I-Q) + \begin{bmatrix} 0 \\ -1 \\ 1 \end{bmatrix} \gamma I, \quad S+I+R=1. \quad (1)$$

This is a slight variant of the traditional SIR model of Kermack and McKendrick (1927), where as usual S , I and R denote the proportions of the population that are susceptible, infectious and removed. The transitions

Fig. 2. Illustration of the stability boundary for the model in (1)–(2). The model is stable if and only if $(R_0, \gamma T_{\text{delay}})$ lies below the corresponding α curve. That is, at least a proportion α of persons becoming infectious on a particular day need to be isolated T_{delay} days later, given a natural recovery rate is $\gamma \text{ days}^{-1}$.

between the states are governed by two rates that model the effect of disease spread and recovery, with mixing and recovery parameters β and γ , respectively. Note in particular that the mixing rate has been adjusted to account for the effect of isolating infectious individuals.

More specifically, Q denotes the proportion of the population that is both infectious and isolated, and the rate that describes the spread of the disease has been modified so that the spread is only driven by interactions between the remaining infectious population and the susceptible population (the $\beta S(I-Q)$ term).

While simple, the formulation in (1) has a strength in that it allows various feedback strategies for regulating disease spread to be stated and analysed within a standard framework. Suppose, for example, that we wish to analyse the effectiveness of isolation schemes, in which infectious individuals are isolated as soon as they are identified as infectious. Arguably the simplest way to model this is through the feedback law

$$Q(t) = \alpha e^{-\gamma T_{\text{delay}}} I(t - T_{\text{delay}}). \quad (2)$$

In words this equation states that the isolated population at time t is equal to some proportion $0 \leq \alpha \leq 1$ of the infectious population T_{delay} days in the past¹. This captures to a first approximation the two most important features of a case isolation scheme, namely, how quickly infectious individuals are identified and isolated (T_{delay}), and what proportion of cases are found (α).

It is important to note here that cases in this context is equivalent to infectious individuals. This should not be confused with *detected* cases. For example, a detection rate of 80 % imposes an upper bound of $\alpha = 0.8$. Particularly, $\alpha = 1$ could only be achieved in a population where *all* cases are detected and *every* individual complies to the scheme.

¹ The $e^{-\gamma T_{\text{delay}}}$ term is correcting for the fact that some of the infectious population will have recovered over these T_{delay} days (recall that Q is the sub-population that is both infectious and isolated).

By analysing (1)–(2) together we can begin to build our intuition for how the isolation scheme parameters α and T_{delay} affect disease spread for a disease with parameters β and γ . For example (as is well known) if no isolation occurs ($\alpha = 0$), then at the start of an epidemic the infectious population will initially grow exponentially according to

$$I(t) \approx e^{(\beta-\gamma)t} I_0,$$

where I_0 is the size of the initial outbreak. A natural question is then how large α must be (or how short T_{delay}) before the period of exponential growth can be averted. This was the question studied in (Pates et al., 2021, §2.1.2), where it was shown that the equations (1)–(2) are locally asymptotically stable in a population in which $S \gg I$ if and only if

$$\gamma T_{\text{delay}} < \ln \left(\frac{\alpha}{1 - R_0^{-1}} \right). \quad (3)$$

In the above $R_0 = \beta/\gamma$ is the basic reproduction number of the disease in the absence of isolation measures.

The specific trade-off between parameters and delay implied by (3) is shown in Fig. 2. This figure can be used to quickly assess the amount of delay that can be tolerated before instability, and hence exponential growth, occurs.

For example, with $\alpha = 0.8$, $R_0 = 3$ and $\gamma = 0.1$, parameters chosen to be representative for SARS-CoV-2, the stability condition becomes

$$T_{\text{delay}} \gamma < 0.18, \implies T_{\text{delay}} < 1.8 \text{ days}.$$

Clearly, short isolation times are essential when dealing with an infectious disease! We also see the importance of identifying a significant proportion of cases. By the time $\alpha \approx 0.7$ (that is, the scheme detects and isolates 70 % of the cases) exponential growth will occur even with $T_{\text{delay}} = 0$.

2.2 Notes on the reproduction number

The reproduction number is a common source to confusion, and therefore deserves explicit attention: In general, the reproduction number R describes the expected number of secondary infections caused by one primary infection. It holds that

$$R \propto \underbrace{\left(\frac{\text{infection}}{\text{contact}} \right)}_{\rho} \cdot \underbrace{\left(\frac{\text{contact}}{\text{time}} \right)}_c \cdot \underbrace{\left(\frac{\text{time}}{\text{infection}} \right)}_d, \quad (4)$$

where

- $\rho \in [0, 1]$: probability of transmission given contact between a susceptible and infectious individual;
- $c \in \mathbb{R}_+$: average contact rate between susceptible and infectious individuals;
- $d \in \mathbb{R}_+$: average duration of infectiousness.

For the homogeneous population model (1) we have that $\beta = \rho c$ and $d = 1/\gamma$. It is also evident from (4) that the reproduction number reveals how *much* an epidemic grows, and not how *fast*. In order to quantify the latter, the serial interval between infections is needed in addition. Yet, the reproduction number is more commonly used than the corresponding growth rate $r = \beta - \gamma$ among infectious disease epidemiologists, which is why we have chosen to use the former in our parameterization of fundamental limitations.

The basic reproduction number R_0 is the reproduction in absence of a *considered* intervention. There is a common misconception that for a particular pathogen R_0 is a universal constant (that can be looked up in the literature). Instead it depends on, among other time-varying parameters, the virulence of the pathogen and the societal structure under consideration. For instance, a particular virus would typically result in different R_0 in a two countries, or in a city versus a village.

If an intervention is enacted, it is instead common to talk about the resulting effective reproduction number R_e (sometimes referred to as the time-varying reproduction number R_t), and it really only makes sense to consider R_0 in relation to a particular intervention: $R = R_0$ in absence of the intervention; $R = R_e$ if the intervention is enacted.

It was shown in Pates et al. (2021) that for the case isolation scheme (2) applied to the model (1) the relation between R_0 and R_e is given by

$$R_e = R_0 (1 - \alpha e^{-\gamma T_{\text{delay}}}), \quad R_0 = \frac{\beta}{\gamma}$$

which is less than 1 if and only if (3) holds.

3. CASE ISOLATION IN HETEROGENEOUS POPULATIONS

3.1 Fundamental limitations

The model introduced in the previous section assumes a homogeneously interacting population. This is clearly a simplification. In reality, the interactions that may lead to infection are much more complex and hard to model. The probability to interact with someone at your workplace is for example much higher than to interact with someone from a remote country. There is also a time-varying aspect. For instance, if your workplace issues a work-from-home guideline to reduce disease transmission, your probability to interact with someone from your workplace will typically drop.

Numerous works have been dedicated to modelling the associated contact network dynamics, with May (2006); Salathé et al. (2010) and survey Nowzari et al. (2016) constituting representative examples. The validity of such network models is hard to verify, and the time-varying aspects further increases the uncertainty surrounding their accuracy. We therefore delimit ourselves to a simple but important question: how does introduction of interaction heterogeneity alter the requirements on our case isolation scheme (2), expressed in terms of the isolation proportion parameter α and associated time delay T_{delay} ?

To account for contact rate heterogeneity, early models of infectious diseases (particularly STDs) in May and Anderson (1987); May et al. (2001) incorporated a distribution of contact rates $p_k = N_k/N$ being the proportion of the population of size N , who on average have k contacts per time unit and in all other regards are homogeneous. As in Anderson et al. (1986) the corresponding disease dynamics of the homogeneous case generalize to

$$\frac{d}{dt} \begin{bmatrix} X_k \\ Y_k \end{bmatrix} = \begin{bmatrix} -k\lambda & 0 \\ k\lambda & -\gamma \end{bmatrix} \begin{bmatrix} X_k \\ Y_k \end{bmatrix}, \quad (5)$$

where X_k and Y_k denote the number of susceptible and infectious individuals in partitions $k = 1, 2, \dots, n$, respec-

tively. The parameter $\lambda = \beta \sum_k k Y_k / \sum_k k N_k$ is the probability of an infection acquired from any one randomly-chosen contact—now more likely to come from the partition with higher contact rates.

Early on in the epidemic it holds $X_k = N_k$, which allows the collective epidemic dynamics to be written as

$$\frac{d}{dt} \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{bmatrix} = \underbrace{\left[\frac{\beta}{\sum_k k N_k} \begin{bmatrix} N_1 \cdot 1 \\ N_2 \cdot 2 \\ \vdots \\ N_n \cdot n \end{bmatrix} [1 \ 2 \ \dots \ n] - \gamma \mathbf{I}_n \right]}_{\mathbf{A}} \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{bmatrix},$$

where $\mathbf{I}_n$ is the $(n \times n)$ identity matrix. The eigenvalues of $\mathbf{A}$ now determine the trajectory of the epidemic early on. They are given by $-\gamma$ and

$$\frac{\beta}{\sum_k k N_k} \sum_k N_k k^2 - \gamma = \beta \left(\frac{\sigma^2}{\mu} + \mu \right) - \gamma,$$

where $\mu = \sum_k k p_k$ and $\sigma^2 = \sum_k k^2 p_k - \mu^2$ denote the mean and variance of the contact rates respectively. It follows that the linearised system about the disease-free equilibrium ($X_k = N_k$ for all k) is stable if and only if

$$\underbrace{\frac{\beta}{\gamma} \left(\frac{\sigma^2}{\mu} + \mu \right)}_h < 1. \quad (6)$$

As originally described in May and Anderson (1987) and May et al. (2001), the adjustment factor $h \geq 1$ in (6) thus increased the mixing parameter $\beta = \rho c$ of (4) relative to the homogenous case, where $\{c = \mu, \sigma = 0\} \Rightarrow h = 1$. Since $R \propto \rho c = \beta$, an equivalent interpretation is that the reproduction number is increased by h .

The adjustment by the factor h remains valid under the case isolation scheme (2), since the probability to isolate an infectious individual is independent of the contact degree of that individual. Consequently, the stability condition of Theorem 2.1 in Pates et al. (2021) becomes

$$T_{\text{delay}} < \frac{1}{\gamma} \ln \left(\frac{\alpha \beta h}{\beta h - \gamma} \right). \quad (7)$$

The derivative of the upper bound on the delay with respect to the coefficient of variation, or relative standard deviation, $c_v = \sigma/\mu$ is

$$\frac{d}{dh} \left[\frac{1}{\gamma} \ln \left(\frac{\alpha \beta h}{\beta h - \gamma} \right) \right] \frac{dh}{dc_v} = - \frac{2c_v}{(c_v + 1)(\beta(c_v^2 + 1) - \gamma)}.$$

This implies that heterogeneity ($c_v > 0$) *increases* the upper bound on T_{delay} if and only if the reproduction number of the corresponding homogeneous system ($c = \mu$) in absence of control (2) fulfils $R_0 < 1/h$. Since $c_v > 0 \Leftrightarrow h > 1$ this requirement implies that heterogeneity allows for a longer T_{delay} in the control law (2) only if the uncontrolled system is already stable in the sense that $R_0 < 1$. On the other hand, and of larger practical importance, the upper bound on admissible T_{delay} *decreases* when heterogeneity (c_v) is increased whenever $R_0 > 1$.

In the latter case, one may ask how large the coefficient of variation can be to allow for a positive delay and a stable equilibrium $(I, R, Q) = (0, 0, 0)$ of the linearised model with the scaled parameter β . That is,

Fig. 3. Illustration of maximum admissible coefficient of variance c_v for stabilization using the case isolation scheme (2) early in an epidemic (when $S \gg I$). In this example, we assume that a proportion α of the infectious population are isolated without delay ($T_{\text{delay}} = 0$). The figure illustrates, for different α , the relation between the basic reproduction number R_0 of a corresponding homogeneously interacting population (in absence of case isolation), and the admissible c_v . The shaded curve shows what combinations of R_0 and c_v cannot be stabilized for $\alpha = 0.8$. The horizontal lines at $c_v = 0.37$ and $c_v = 0.67$ represent bounds on c_v for human infectious disease transmission networks reported in Salathé et al. (2010).

$$\frac{1}{\gamma} \ln \left(\frac{\alpha \beta h}{\beta h - \gamma} \right) > 0 \xrightarrow{R_0 > 1, \alpha \in [0, 1]} c_v^2 < \frac{1}{R_0(1 - \alpha)} - 1, \quad (8)$$

where, again, R_0 is the basic reproduction number for a corresponding homogeneously interacting population.

Fig. 3 illustrates this maximum allowed coefficient of variance c_v of the contact graph, as a function of the basic reproduction number and proportion of cases α that are isolated according to (2).

Revisiting the example from Sec. 2 with $\gamma = 0.1$ and $R_0 = 3$, the condition (8) reveals that the maximum coefficient of variation for which a positive isolation delay $T_{\text{delay}} > 0$ is allowed evaluates as $c_v \approx 0.82$. As a reference, Salathé et al. (2010) reports lower and upper bounds of 0.37 and 0.63 for c_v in human infectious disease transmission contact networks. The corresponding lines have been added to Fig. 3 and indicate that if $R_0 = 3$ was computed or estimated for a homogeneous population, delay-free ($T_{\text{delay}} = 0$) isolation of a proportion somewhere in the range $0.6 < \alpha < 0.8$ would be required to stabilize the epidemic under (2).

3.2 Notes on contact heterogeneity

The phenomenon of the effective mixing parameter β increasing with increased heterogeneity (quantified by the coefficient of variation c_v) is also the answer to the question “Why your friends have more friends than you do”, explained with mathematical insight in Feld (1991). The intuition is that individuals with many contacts are more likely to get infected, and therefore the infectious proportion of the population will comprise individuals who

Fig. 4. Simulation of 100 epidemics (grey) over an Erdős Rényi graph where the mean number of connections to the infectious (black) are shown each day. At the beginning of the epidemic this amounts to the mean contact degree μ of the graph, but soon stabilizes at the expectation $\mu + \sigma^2/\mu$ of the excess degree distribution of the network, with σ^2 being the contact degree variance of the network.

have more social contacts (higher degree) than those in the susceptible proportion of the population.

As with friendships, the contacts considered here are mutual, translating into an undirected network graph. Thus, the early spread of the disease is proportional to the mean number of contacts of a randomly chosen contact of an infectious individual (node), given by h in (6). Treating the network as homogeneous ($c_v = 0 \Leftrightarrow h = 1$) can therefore lead to an under-estimation of the reproduction number by disregarding the growth of the infectious subpopulation, which is fueled by highly connected individuals that are both more likely to acquire and spread infection.

In Fig. 4 we have illustrated this effect through simulations of the early stage of 100 epidemics on an undirected random graph representing a population. The nodes of the graph represent individuals; the edges contacts.

For the example in Fig. 4, we chose to generate Erdős Rényi graphs of $n = 2 \cdot 10^6$ nodes and $M = 5 \cdot 10^6$ undirected edges due to their simplicity, but other types of graph generation methods could also be used to illustrate the phenomenon.

The epidemics were each seeded by randomly assigning 10 infectious individuals on day one. This corresponds to $I = 5 \cdot 10^{-6} \ll S = 1 - I$. The transmission probability from (4), being the per-day probability that an edge between one infectious and one susceptible node leads to disease transmission, was set to $\rho = 0.05$. The recovery rate from (1) was set to $\gamma = 0.1$. Assuming independent interactions, the probability for a susceptible to become infectious the next day, given that it has m infectious contacts, is thus $1 - (1 - \rho)^m$.

The cumulative infectious proportion was retrospectively evaluated to satisfy $I + R \ll 0.01$ for each of the 100 simulations, validating the assumption of a fully susceptible

population (or early epidemic stage) on which the stability analysis has been based.

In Fig. 4 the results from 100 epidemics simulations are shown. In particular, we see the average number of contacts of the infectious population throughout the early stages of the epidemic. From this figure we can clearly identify how the infectious individuals have a higher degree of contacts than the general population of the network. Furthermore, we can see that average number of contacts of the infectious population tends towards the expected value of the excess degree distribution, being the distribution of edges from a node reached by following an edge.

4. DISCUSSION

In Sec. 2 we characterized the stability boundary (3) associated with the case isolation scheme (2) as a function of the basic reproduction number R_0 in absence of case isolation, and a recovery rate γ defined through (1).

In Sec. 3 we then characterized how the stability boundary is moved if R_0 of a homogeneous population is used in (3), when in fact the population is heterogeneous in the sense that the coefficient of variation of the contact network c_v is positive, while the mean degree remains unchanged.

Introducing heterogeneity in this way corresponds to multiplication of the mixing parameter β in (1) with the scaling factor $h = c_v^2 + 1$. Since β is directly proportional to R_0 , the heterogeneity can be interpreted as increasing the reproduction number by the corresponding factor. Note that the same shifting applies for a population partitioned based on heterogeneity of the infectiousness parameter ρ . This makes it possible to apply the same modelling to account for a variability in infectiousness across variants of the considered pathogen.

Whether to multiply by the factor h before using R_0 in the analysis comes down to how R_0 was estimated from data. If it was estimated taking the heterogeneity into account it should not be adjusted, otherwise it should.

In the heterogeneous model (5) proposed in Anderson et al. (1986), the interaction between individuals with many and few contacts respectively are stochastic, and therefore do not correspond to a fixed (time-invariant) contact graph. For the case of a fixed contact graph, the stability bound (7) obtained for (5) would be conservative because of the assumption that the disease can spread to any contact of an infectious individual. In particular, an infectious individual cannot re-infect its own infector. Therefore, when considering epidemics on configurator model graphs, the shift of the epidemic threshold needs to be corrected by -1 , resulting in $h = \frac{\sigma^2}{\mu} + \mu - 1$, as further discussed in Newman (2018); Meyers (2007); Kiss et al. (2017).

Relatedly, clustering of the contact network result in the bound (7) becoming conservative. When many neighbours in the network are shared, (e.g. as in small-world networks) local clustering coefficients in the contact network are high and voids the assumption that all neighbors of an infectious node are susceptible. To analytically quantify how local and global clustering effects affect the stability

bound of the studies case isolation schema along the lines of Trapman (2007) requires more complex network models. It falls outside the scope of this contribution, but constitutes a worthwhile direction for future work.

5. CONCLUSION

While obtaining detailed graph models of (the time-varying) human contact graphs relevant for infectious disease transmissions is generally not tractable, control theoretic analyses can provide qualitative, and to some extent quantitative criteria for the feasibility of strategies aimed at halting disease spread. This has been illustrated here by expressing stability conditions for case isolation schemes in homogeneous and heterogeneous populations, as functions of fundamental epidemiological parameters.

6. ACKNOWLEDGMENTS

This work was partially funded by the Swedish Research Council (grants 2017-04989, 2019-00691) and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The authors are members of the ELLIIT Strategic Research Area at Lund University.

REFERENCES

- Anderson, R., Medley, G., May, R., and Johnson, A. (1986). A preliminary study of the transmission dynamics of the human immunodeficiency virus (HIV), the causative agent of AIDS. *Mathematical Medicine and Biology: a Journal of the IMA*, 3(4), 229–263.
- Feld, S. (1991). Why your friends have more friends than you do. *American Journal of Sociology*, 96(6), 1464–1477. doi:10.1086/229693.
- Kermack, W.O. and McKendrick, A.G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London A: mathematical, physical and engineering sciences*, 115(772), 700–721.
- Kiss, I.Z., Miller, J.C., Simon, P.L., et al. (2017). *Mathematics of epidemics on networks*. Springer.
- May, R. (2006). Network structure and the biology of populations. *Trends in Ecology & Evolution*, 21(7), 394–399. doi:10.1016/j.tree.2006.03.013.
- May, R., Gupta, S., and McLean, A. (2001). Infectious disease dynamics: what characterizes a successful invader? *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 356(1410), 901–910.
- May, R.M. and Anderson, R.M. (1987). Transmission dynamics of HIV infection. *Nature*, 326(137), 10–1038.
- Meyers, L. (2007). Contact network epidemiology: Bond percolation applied to infectious disease prediction and control. *Bulletin of the American Mathematical Society*, 44(1), 63–86.
- Newman, M. (2018). *Networks*. Oxford University Press.
- Nowzari, C., Preciado, V.M., and Pappas, G.J. (2016). Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems Magazine*, 36(1), 26–46.
- Oliver, D. (2021). Mistruths and misunderstandings about covid-19 death numbers. *British Medical Journal*, 372. doi:10.1136/bmj.n352.
- Pates, R., Ferragut, A., Pivo, E., You, P., Paganini, F., and Mal, E. (2021). Respect the unstable: delays and saturation in contact tracing for disease control. *SIAM Journal on Control and Optimization (SICON)*. Submitted.
- Robert, A. (2020). Lessons from New Zealand’s COVID-19 outbreak response. *The Lancet Public Health*, 5(11), E569–E570. doi:10.1016/S2468-2667(20)30237-1.
- Salathé, M., Kazandjieva, M., Lee, J., Levis, P., Feldman, M., and Jones, J. (2010). A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences*, 107(51), 22020–22025. doi:10.1073/pnas.1009094108.
- Soltesz, K., Gustafsson, F., Timpka, T., Jaldén, J., Jidling, C., Heimerson, A., Schön, T., Spreco, A., Ekberg, J., Dahlström, Ö., Bagge Carlson, F., Jöud, A., and Bernhardsson, B. (2020). The effect of interventions on COVID-19. *Nature*, 588, E26–E28. doi:10.1038/s41586-020-2405-7.
- Topley, W. and Wilson, G. (1923). The spread of bacterial infection. the problem of herd-immunity. *The Journal of Hygiene*, 21(3), 243–249. doi:10.1017/s0022172400031478.
- Trapman, P. (2007). On analytical approaches to epidemics on networks. *Theoretical population biology*, 71(2), 160–173.