Increasing efficacy of contact-tracing applications by user referrals and stricter quarantining

Leslie Ann Goldberg[*], Joost Jorritsma[†], Júlia Komjáthy[†], and John Lapinskas[‡]

November 30, 2020

Abstract

We study the effects of two mechanisms which increase the efficacy of contact-tracing applications (CTAs) such as the mobile phone contact-tracing applications that have been used during the COVID-19 epidemic. The first mechanism is the introduction of user referrals. We compare four scenarios for the uptake of CTAs — (1) the $p\%$ of individuals that use the CTA are chosen randomly, (2) a smaller initial set of randomly-chosen users each refer a contact to use the CTA, achieving $p\%$ in total, (3) a small initial set of randomly-chosen users each refer around half of their contacts to use the CTA, achieving $p\%$ in total, and (4) for comparison, an idealised scenario in which the $p\%$ of the population that uses the CTA is the $p\%$ with the most contacts. Using agent-based epidemiological models incorporating a geometric space, we find that, even when the uptake percentage $p\%$ is small, CTAs are an effective tool for mitigating the spread of the epidemic in all scenarios. Moreover, user referrals significantly improve efficacy. In addition, it turns out that user referrals reduce the yearly quarantine load. The second mechanism for increasing the efficacy of CTAs is tuning the severity of quarantine measures. Our modelling shows that using CTAs with mild quarantine measures is effective in reducing the maximum hospital load and the number of people who become ill, but leads to a relatively high quarantine load, which may cause economic disruption. Fortunately, under stricter quarantine measures, the advantages are maintained but the quarantine load is reduced. Our models incorporate geometric inhomogeneous random graphs to study the effects of the presence of super-spreaders and of the absence of long-distant contacts (e.g., through travel restrictions) on our conclusions.

1 Introduction

Agent-based epidemiological models allow a population to be embedded in a geometric space to capture the effect of the "local" and "long-distance" contacts that arise in real populations. We use such a model to study the efficacy of contact-tracing applications such as the mobile phone contact-tracing applications that have been introduced during the COVID-19 epidemic. Our simulation study compares the efficacy of four different uptake scenarios for the contact-tracing application (CTA), combined with the effect of quarantining measures of various strengths. We compare four uptake scenarios: (i) the users of the CTA are a random subset of the population, the users are selected by one of two different recommender scenarios, in which randomly-chosen initial users recommend the CTA to some of their contacts ((ii) either to just one of them, or (iii) to roughly half of them)

[*]Department of Computer Science, University of Oxford, UK.
[†]Department of Mathematics and Computer Science, Eindhoven University of Technology, NL.
[‡]Department of Computer Science, University of Bristol, UK.

who also use the CTA, and finally (iv) a (hypothetical) degree-targetted uptake scenario, in which CTA-users are individuals with the highest number of connections. In each of these CTA-uptake scenarios, we assume that individuals in the population will quarantine themselves as soon as their symptoms are sufficiently strong (exceeding a certain threshold). At this point, users of the CTA also notify their CTA-using contacts, who also quarantine themselves. We examine the combined effect of differing CTA-uptake scenarios and different symptom thresholds for quarantine. We decouple the effects of CTA-uptake scenarios and symptom thresholds from the effects of testing delays, which have been studied in [1]. In order to focus on the former, we work in an idealised hypothetical scenario where testing is immediate upon symptom onset (or where there is only a single virus present in the population that could cause the symptoms) and where the symptom-severity does not influence the probability of a virus transmission between two individuals. We emphasise that our study is qualitative, rather than quantitative, and we aim for universally valid observations in terms of the performance of CTA-uptake scenarios.

Main observations. In our first experiment we qualitatively compare the four uptake scenarios. In all uptake scenarios, we find that contact-tracing applications (CTAs) are effective in decreasing the volume of an epidemic (the total number of people who become ill) and in decreasing the maximum number of people who are simultaneously hospitalised, and typically this effectiveness increases linearly or super-linearly with the percentage of the population that use CTA. In these respects, we show that CTAs are effective even with low uptake rates. This confirms the results of a recent study modelling CTAs in Washington state [2]. The novelty of our study is to compare recommendation-based uptake with random uptake, where we find that recommending can significantly improve the efficacy of the CTA. In brief, we find that in scenarios where the CTA is recommended to acquaintances the volume and maximum hospital load decrease at a much higher rate than the rate that would be achieved by randomly selecting the same number of CTA users. Even better, in recommender scenarios also the yearly quarantine load on society is typically reduced. It cannot be left unmentioned, however, that recommender scenarios might elongate the epidemic for some types of underlying contact networks of individuals, especially at low uptake rates. Finally, we emphasise that the uptake percentage necessary to completely eradicate the epidemic from the population is generally very high, as was also found in [3], though this is not the focus of this article.

In our second experiment we study the impact of the severity of quarantine measures, both on the epidemic itself and on the economic disruption that it may cause (measured in terms of the time that people spend in quarantine over the course of the epidemic). Already under "mild" quarantine measures, the maximum hospital load is drastically reduced and the volume of the epidemic is somewhat reduced. However, the economic disruption caused by quarantine is very high. Perhaps surprisingly, under stricter quarantine measures (i.e., quarantining individuals already with less severe symptoms), the epidemic is not only better contained, but also the economic disruption is lower (under strict enough social distancing measures, travel restrictions, and sufficiently high app uptake). The critical point on the 'quarantine-strictness scale' above which the social disruption starts decreasing comes earlier in recommender scenarios and degree-targetted uptake than in random uptake, and even earlier if individuals are more restricted in their movements, i.e., the underlying contact network does not have many long-range connection, ranging over large spatial distances.

As a disclaimer, we emphasise that we studied a case where a single epidemic is present in the population and individuals who are not infected do not show any symptoms similar to those that are caused by the virus under investigation. This might be unrealistic for certain types of diseases, but might be a better approximation of reality for others. Finally, we emphasise that we aim for qualitative, comparative results that are robust against changes in the dynamics and the underlying contact network. We refrain from making numerical predictions.

Organisation. In the remainder of this section we describe the underlying contact network, the CTA-uptake scenarios and varying quarantine measures that we study. We then explain and interpret our main findings regarding each of these. We also discuss the effect of social distancing measures. In Sections 2 and 3 we define the agent-based model that we use and the underlying contact networks that we use to model the population.

1.1 Networks to model the population.

We use so called agent-based models to model the spread of the epidemic. We assume that a population is embedded in a geometric space, and the connections between individuals are modelled using a random network model called a Geometric Inhomogenous Random Graph (GIRG) [4]. This is a state-of-the-art model for real-world social and technological networks, embedded in a geometric space. We think of the nodes of the network as individuals, having a fixed location in space. The neighbours of a node u are nodes with a direct connection (also called a link or an edge) to u. A connection may correspond to a friendship or an acquaintance, or simply a contact event, and is correlated with, but does not necessarily coincide with spatial proximity.

There are two robust parameters of GIRGs that control the qualitative features of the network: a parameter τ controls the variability in the number of neighbours that individual nodes have (the number of neighbours of a node is called its "degree"). Smaller values of τ correspond to more variability, while keeping the average the same. A second parameter α controls the number of long-range edges: long-range edges would tend to be present in populations without travel restrictions, (as observed in real-life contact networks, see e.g. [5, 6, 7]).

A value of α close to 1 corresponds to having many long-range edges, in this case the network resembles an a-geometric network model (for more on this see Section 3.1 below); while increasing α reduces the number of long-range contacts, thus transforming the network into one where the underlying geometry is also more apparent intrinsically. The following summary will help the reader to understand the experimental results reported below.

- $\tau = 2.3$: a GIRG with plenty of "super-spreaders" (nodes with many connections) due to degree variability,
- $\tau = 3.3$: a GIRG with fewer super-spreaders,
- $\alpha = 1.3$: a less geometric GIRG with many long-range edges,
- $\alpha = 2.3$ a more geometric GIRG, i.e., a network that resembles a lattice better, it does not have too many long-range connections.

Details are given in Section 3. For comparison, we also used a common uncorrelated network model, the *configuration model*, corresponding to a network with no underlying geometry. See Section 3.2 — configuration models do not have an associated α parameter expressing the strength of geometry, though there is still a parameter τ indicating the presence of many (or fewer) super-spreaders. Once again, we take $\tau = 2.3$ for the former and $\tau = 3.3$ for the latter.

1.2 Key performance indicators (KPIs).

We study four KPIs as a function of the parameters:

- Vol: The total volume, meaning the total number of individuals ever infected during the whole course of the epidemic.

- HMax: The (approximated) maximum hospital load. In our approximation, we define the hospital load at any time to be 5% of the infected (symptomatic or asymptomatic) individuals. Of course, the 5% is just a scaling, and changing this factor does not change the shape of the curve.
- Quar: The average number of days spend in quarantine per person per year, over the course of the epidemic.
- Time: The total duration of the epidemic, in years.

1.3 Experiment 1: Application-uptake scenarios

An important aspect of our study is a comparison of the efficacy of four different CTA-uptake scenarios. Suppose that $p\%$ of the population use and also comply with a CTA, in the sense of reporting symptoms and also going into self-quarantine when instructed to do so. We compare the effectiveness of the CTA in terms of reducing the total volume Vol of the epidemic and the maximum hospital load HMax, when this $p\%$ uptake is achieved in four different ways:

(i) randomly chosen: $p\%$ of the population uses the CTA, and this $p\%$ is chosen uniformly at random,
(ii) via basic recommendation: a set of "initial users" is chosen uniformly at random. Each initial user recommends the CTA to a single uniformly-chosen neighbour who then also uses it,
(iii) via ring recommendation: each initial user recommends the CTA to its neighbours, each of them will use the app with probability 0.5 independently of the rest,
(iv) via degree targetting: the $p\%$ of the population with the highest number of connections use the CTA.

In terms of controlling the epidemic, one would expect that it is desirable if the $p\%$ most influential members of the population use the CTA, rather than a randomly-chosen $p\%$. However, *targetting* the application in this way incurs cost (it is necessary to determine which members of the population have the most influence, and to persuade them to use the application). The recommender scenarios provide a less-costly, easier-to-implement solution. Each time a (randomly-chosen) member of the population begins using the CTA, this user is asked to persuade either a single, randomly-chosen contact or all of its contacts to also use the CTA. Informally, these contacts are friends or acquaintances. In our network model, they correspond to neighbours in the network. See Section 1.1 and Definition 4 for details. In the case where all contacts are asked to use the CTA, our model assumes that roughly half of these contacts do begin using the CTA. This ring recommendation method has some connections with a certain vaccination protocol, called ring vaccination, hence the name. Ring vaccination has been successful in the past to eradicate small pox [8] and has also been used against Ebola [9]. Such recommender scenarios have previously been shown to be effective in the context of vaccines [10] on synthetic models. Note that, to achieve the same uptake percentage $p\%$, the initial percentage of users for the recommender-based scenarios is *smaller* than $p\%$, see Fig 1.

Results. We varied the CTA-uptake percentage for the four uptake scenarios on six networks, each with $500,000$ individuals and average degree 13 (average degree 13 roughly corresponds to empirical findings concerning the number of contacts per individual [11]), using the network models from Section 1.1. In each case, we studied the effect of the CTA-uptake percentage and the uptake scenario (randomly chosen, degree-targetted, basic recommender or ring recommender) on the four KPIs from Section 1.2. We explain our results using the GIRGs from Section 1.1 with $\tau \in \{2.3, 3.3\}$ and $\alpha = \{1.3, 2.3\}$. The last two networks mentioned in Section 1.1, modelled by the configuration model, are merely for comparison and are included in our later figures.

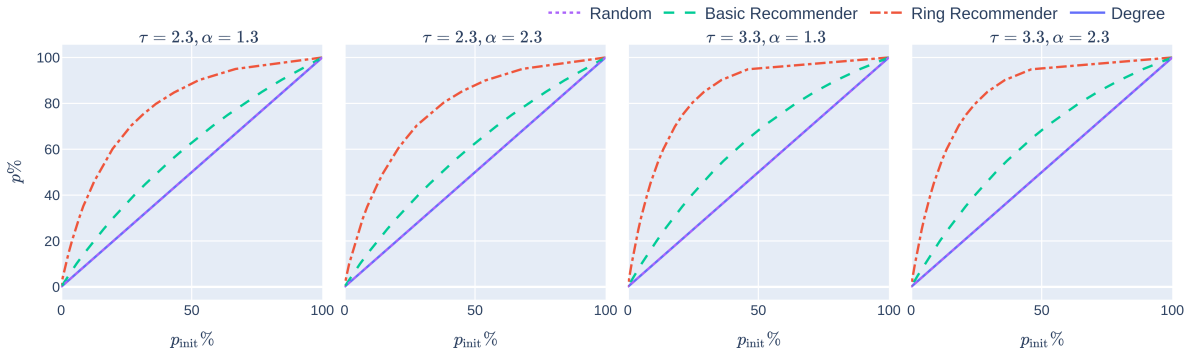

Figure 1: The initial app uptake $p_{\text{init}}\%$ in relation to the final app uptake $p\%$ on four networks and the different uptake scenarios. The curves for random uptake degree-targetted uptake coincide and are straight lines, since it holds that $p\% = p_{\text{init}}\%$, while for recommender scenarios $p\% > p_{\text{init}}\%$. Increasing $p\%$ increases the probability that a node is recommended by multiple persons, explaining the concave curves corresponding to the recommender scenarios.

1.3.1 Hospital load and Volume: Recommenders perform very well, ring recommender best.

The main message indicated by our simulations (row 2 of Fig 2) is that introducing recommender scenarios strongly increases the efficacy of a CTA in terms of reducing hospital load: Having achieved uptake $p\%$ via any recommender scenario is significantly more desirable (in terms of reducing the maximum hospital load) than having achieved $p\%$ uptake by randomly chosen members of the population. Furthermore, the degree-targetted scenario works best, and the ring-recommender scenario is almost as good as this. Adding the subscripts rand, deg, basic, and ring to indicate the uptake scenarios, we find typically that

$$\text{HMax}_{\text{deg}}(p) < \text{HMax}_{\text{ring}}(p) < \text{HMax}_{\text{basic}}(p) < \text{Hmax}_{\text{rand}}(p). \tag{1}$$

The same inequalities hold true for the volume of the epidemic (first row of Fig 2):

$$\text{Vol}_{\text{deg}}(p) < \text{Vol}_{\text{ring}}(p) < \text{Vol}_{\text{basic}}(p) < \text{Vol}_{\text{rand}}(p). \tag{2}$$

For a different depiction of this phenomenon at $p = 5\%$ see Fig 3, where the course of the epidemic is plotted against *time*. The epidemic curve on the left represents the total number of infected individuals (during the course of the epidemic). The epidemic curve on the right represents the number of individuals that are currently infected. The figure demonstrates how the random-uptake scenario is outperformed by the other uptake scenarios.

The reason that recommender scenarios always perform better than the random uptake scenario is explained by the so-called friendship-paradox [14]: a friend of a random user is likely to have a higher number of contacts than the random user itself, so a recommender scenario intrinsically finds users with high numbers of connections. This effect is more pronounced when the node-degrees are highly varying ($\tau = 2.3$) and it is even more pronounced in the ring-recommendation scenario. The reason for this is that ring recommendation introduces highly connected clusters of CTA users, and it is hard for the epidemic to thrive in such an environment: the CTA clusters block its spreading.

1.3.2 Improving CTA efficacy by slightly increasing uptake rates.

We study the shape of the curves $\text{Vol}_{\square}(p)$ and $\text{HMax}_{\square}(p)$ as p varies and $\square = \text{rand}, \text{basic}, \text{ring}, \text{deg}$. Intuitively, we examine what happens when we slightly increase the CTA-uptake percentage, espe-

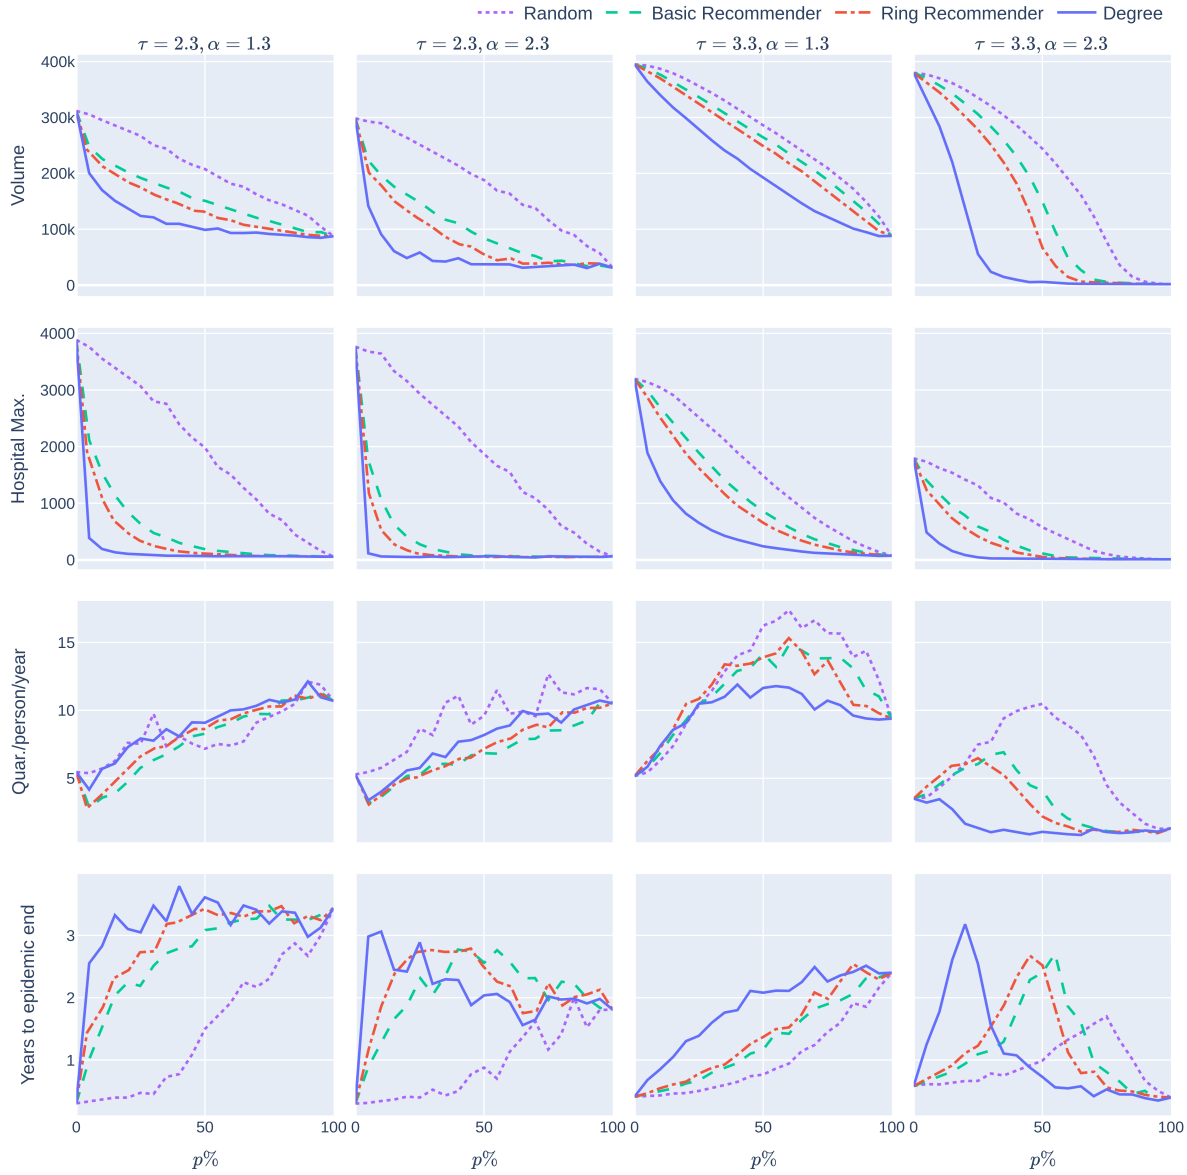

Figure 2: The KPIs (Vol, HMax, Quar, Time) as a function of the CTA-uptake percentage $p\%$ for the four CTA-uptake scenarios can be found in the rows. The population is modelled using four GIRG networks where $\tau = 2.3$ has more super-spreaders and $\tau = 3.3$ has fewer super-spreaders, $\alpha = 1.3$ is less geometric and $\alpha = 2.3$ is more geometric, arranged in the four columns. The x-axes shows the uptake percentage $p\%$ varying from 0 to 100% at step sizes 5%, the y-axes shows the corresponding value of the KPI. The four uptake scenarios correspond to the four curves on each figure. For each parameter value, the results are the median over 10 runs.

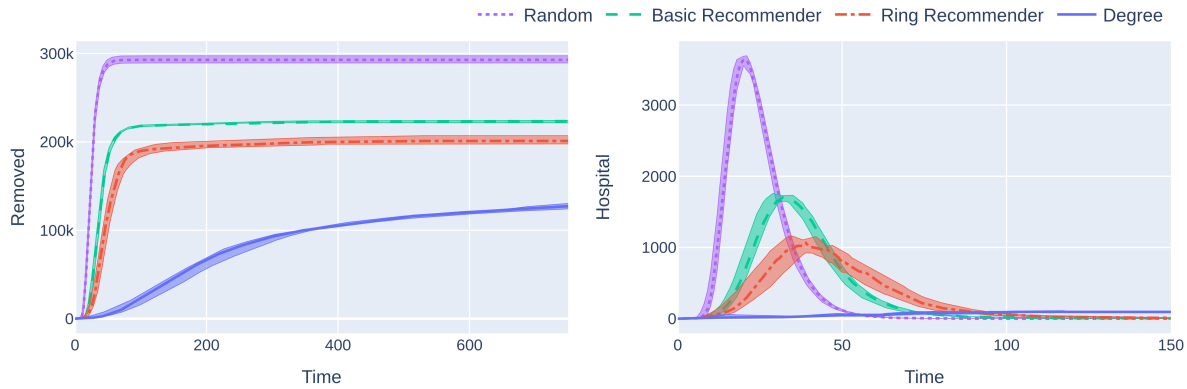

Figure 3: Left: The number of removed individuals over the course of an epidemic in a network with many super-spreaders and strong geometry, with $p\% = 5\%$. The random uptake scenario is strongly outperformed by all other scenarios and has a steep incline, leading to a high volume of people who get infected. It is remarkable that already at such a low uptake percentage, even a basic-recommendation scenario can improve the efficacy of the CTA compared to random uptake: the total volume shrinks to roughly $2/3$ of that by random uptake (left), and the maximum hospital load shrinks to roughly $1/3$ of that by random uptake. The (dashed/dotted/dashdotted) curves correspond to the median over 20 runs; the shaded area covers the 25th-75th percentile of these runs over time. Right: The number of hospitalised individuals over time in the same epidemic simulation. Once again, recommender scenarios strongly outperform the random uptake scenario.

cially in the (at least in Europe) more realistic low-uptake regimes ($< 50\%$).

Our first finding (see the second row of Fig 2) is that random uptake reduces HMax roughly *linearly* in the uptake percentage $p\%$. For some slope $z > 0$ that depends on the underlying network and the infectiousness of the disease (which is fixed in this section, but varied in our experiments, see Fig 9 below).

$$\mathrm{HMax}_{\mathrm{rand}}(p) \approx (1 - zp)\mathrm{HMax}(0), \tag{3}$$

The same holds for the volume $\mathrm{Vol}_{\mathrm{rand}}(p)$ (see the first row of Fig 2), when $\tau = 2.3$ so there are many super-spreaders in the network, while the curve is more concave when $\tau = 3.3$ (fewer super-spreaders). See also Fig 8, where the level of infectiousness (the infection-rate) is varied. Observe that $\mathrm{Vol}(0)$ and $\mathrm{Hmax}(0)$ do not depend on the uptake scenario, since no individual is using the CTA when $p = 0$. We explain why the approximately linear decline is intuitively surprising. For an infection to be prevented via the CTA, both the infector and the infected individual have to use the CTA, and the chance of this by random uptake is $(p/100)^2 \ll p/100$. Thus one's first guess would be that $\mathrm{HMax}_{\mathrm{rand}}(p), \mathrm{Vol}_{\mathrm{rand}}(p)$ will be concave in p, not linear, over all infectiousness parameters. Yet, we see that the curve is concave only in scenarios with high infectiousness and in contact networks with strong underlying geometry (see Figs 8 and 9, where the infectiousness parameter (β) is varied). The conclusion is that, even with random uptake, the CTA works better than might be expected, decreasing Vol and HMax linearly for a large range of parameter values (of infectiousness, and of the underlying network), rather than concavely.

Even better, we find (see the first row of 2) that the basic-recommender scenario and especially the ring-recommender scenario reacts more strongly (than the random-uptake scenario) to a slight increase in the uptake percentage $p\%$: i.e., the volume curves decrease at a faster rate under basic and ring recommendation, as compared to under the random uptake scenario. That is, we find that for

$p\% < 50\%$ and a small positive number $\mathrm{d}p$,

$$\begin{aligned}\mathrm{Vol}_{\mathrm{ring}}(p) - \mathrm{Vol}_{\mathrm{ring}}(p+\mathrm{d}p) &\gg \mathrm{Vol}_{\mathrm{rand}}(p) - \mathrm{Vol}_{\mathrm{rand}}(p+\mathrm{d}p) > 0,\\ \mathrm{Vol}_{\mathrm{basic}}(p) - \mathrm{Vol}_{\mathrm{basic}}(p+\mathrm{d}p) &\gg \mathrm{Vol}_{\mathrm{rand}}(p) - \mathrm{Vol}_{\mathrm{rand}}(p+\mathrm{d}p) > 0,\end{aligned} \tag{4}$$

meaning that (apart from perhaps for $\tau = 3.3, \alpha = 1.3$), the total volume as a function of p decreases substantially more steeply for ring-recommender and basic-recommender uptake than for the random uptake. Furthermore, in the situation where there are many super-spreaders ($\tau \in (2,3)$), the curves are convex:

$$\mathrm{Vol}_{\mathrm{ring}}(p) - \mathrm{Vol}_{\mathrm{ring}}(p+\mathrm{d}p) > \mathrm{Vol}_{\mathrm{basic}}(p) - \mathrm{Vol}_{\mathrm{basic}}(p+\mathrm{d}p) > \mathrm{Vol}(0)z\mathrm{d}p,$$

meaning that the volume decreases the most in the ring-recommender scenario. For the maximum hospital load HMax, the effect is similar, but stronger (see row 2 of 2), and for all underlying networks, HMax is a convex curve for ring-recommender and basic-recommender scenarios:

$$\begin{aligned}\mathrm{HMax}_{\mathrm{ring}}(p) - \mathrm{HMax}_{\mathrm{ring}}(p+\mathrm{d}p) > \mathrm{HMax}_{\mathrm{basic}}(p) - \mathrm{HMax}_{\mathrm{basic}}(p+\mathrm{d}p)\\ \gg z\mathrm{d}p\mathrm{HMax}(0) > 0.\end{aligned} \tag{5}$$

Mathematically, we find that for many underlying contact networks, the functions $\mathrm{Vol}_{\mathrm{ring}}(p)$, $\mathrm{Vol}_{\mathrm{basic}}(p)$, $\mathrm{HMax}_{\mathrm{ring}}(p)$, and $\mathrm{HMax}_{\mathrm{basic}}(p)$ are convex, having a steep decline when $p\% < 50\%$, while $\mathrm{Vol}_{\mathrm{rand}}(p)$ and $\mathrm{HMax}_{\mathrm{rand}}(p)$ are linear or concave functions. This is illustrated further in s 8 and 9 where also the level of infectiousness is varied: the phenomenon is robust in the 'supercritical' regime where there is a large outbreak, apart from in graphs with few super-spreaders in situations with very high infection rates. In this latter case the curves might coincide or are more concave.

1.3.3 Geometry helps the CTA.

A less obvious observation that we make is that (considering networks with the same average number of connections per node), the more geometric the underlying contact network is, the better the performance of the CTA in terms of reducing volume and the maximum hospital load, across the four uptake scenarios. In other words, lack of long-range connections helps CTAs.

Moreover, the positive effect of recommending becomes more exaggerated in geometric networks, see Fig 4. Emphasising in notation by adding a superscript 'geo' and 'ageo' to emphasise whether the underlying network has strong geometry ($\alpha = 2.3$) or less geometry ($\alpha = 1.3$), we typically find that (for fixed τ and reasonable uptake percentages)

$$\begin{aligned}\mathrm{Vol}^{\mathrm{geo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{geo}}_{\mathrm{basic}}(p) &\gg \mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{basic}}(p),\\ \mathrm{Vol}^{\mathrm{geo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{geo}}_{\mathrm{ring}}(p) &\gg \mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{ring}}(p).\end{aligned} \tag{6}$$

This is illustrated on Fig 4, where the volume difference (corresponding to the number of infections that are prevented by the recommendation) is plotted. Observe that the dashed red curve, corresponding to the volume difference in geometric networks stays above the blue curve, corresponding to the volume difference in ageometric networks. We emphasise that both curves are always positive, i.e., recommendation always performs better than random uptake.

The intuitive explanation for Equation (6) is the same as what we gave in Section 1.3.1: recommending has two main effects. The first effect is finding the super-spreaders via the friendship-paradox effect. The second effect of recommending is only present in geometric networks: recommending also

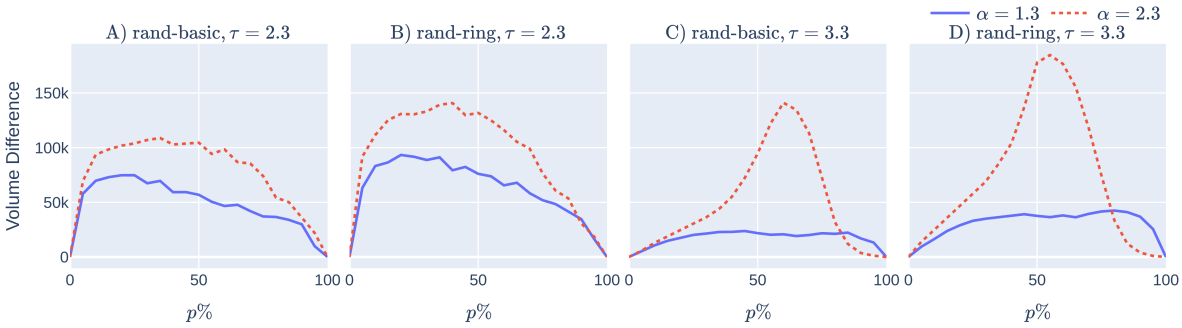

Figure 4: The reduction in epidemic volume achieved by CTAs with user recommendation improves when the network has stronger geometry (for all values of $p\%$ when $\tau = 2.3$, and for $p\% < 85\%$ when $\tau = 3.3$). Subfigures (A) and (C): The red dashed curve shows $\mathrm{Vol}^{\mathrm{geo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{geo}}_{\mathrm{basic}}(p)$, while the blue continuous curve shows $\mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{basic}}(p)$ as a function of the uptake percentage $p\%$. In subfigures (B), and (D), the red dashed curve shows $\mathrm{Vol}^{\mathrm{geo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{geo}}_{\mathrm{ring}}(p)$, while the blue continuous curve shows $\mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{rand}}(p) - \mathrm{Vol}^{\mathrm{ageo}}_{\mathrm{ring}}(p)$ as a function of the uptake percentage $p\%$. Observe that the red dashed curves are (typically) above the blue curves. In subfigures (A) and (B) the underlying network contains many super-spreaders, while in subfigures (C) and (D) the network contains fewer: observe that the presence of fewer super-spreaders makes the red dashed curve, the volume difference in geometric networks between recommender and random uptake, more spiky.

forms geometric barriers of CTA-users around a central node. These barriers are hard for the infection to pass through.

On the empirical side, we mention that a way to introduce more geometry into the underlying contact network is via travel restrictions. Two of the authors previously studied the effect of such intervention on epidemic curves in [15], which found that already in simple scenarios where there are no CTAs, travel restrictions are indeed most effective in reducing the height of the (first) peak.

1.3.4 How costly is quarantine?

A common belief is that using a CTA might cause a high quarantine load on the population. However, row 3 of Fig 2 demonstrates that (as long as social-distancing measures make the underlying infection rate not too high) the average number of days per year that an individual has to quarantine does not rise very steeply with the CTA-uptake percentage $p\%$. Though it is beyond the scope of this study, note that quarantine time can be reduced by providing prompt testing to people who are notified to quarantine by the CTA.

User recommendations make CTAs more effective, in terms of reducing volume and hospital load. Fortunately, this is *not* at the expense of increased quarantining — in most of the graphs in Fig 2, the average quarantine time under recommender scenarios is actually less than what occurs with random uptake, see the third row of Fig 2. Note that the average quarantine time/year/person has a high variance, hence the non-smooth curves. Typically we have:

$$\mathrm{Quar}_{\mathrm{ring}}(p) \approx \mathrm{Quar}_{\mathrm{basic}}(p) \lesssim \mathrm{Quar}_{\mathrm{rand}}(p). \tag{7}$$

For one underlying network class (many super-spreaders and less geometry $\tau = 2.3, \alpha = 1.3$), all scenarios lead to about the same amount of quarantine per year, see Fig 2 third row first column. We note, however, that while the yearly quarantine load per person is lower using recommendation

scenarios, for *all but one* underlying network class these non-random CTA uptake scenarios *elongate* the epidemic:

$$\mathrm{Time}_{\mathrm{rand}}(p) < \mathrm{Time}_{\mathrm{basic}}(p) \approx \mathrm{Time}_{\mathrm{ring}}(p) \lesssim \mathrm{Time}_{\mathrm{deg}}(p) \tag{8}$$

We mention that the total duration of the epidemic has a quite high variance: this is due to a final stage when most of the population already had the disease but yet the infection manage to 'linger' in some hard-to-reach parts of the network.

Here we note a surprising feature that applies to underlying networks with relatively few super-spreaders ($\tau = 3.3$) in Fig 2. In this case, $\mathrm{Quar}(p)$ is maximised for some $p^\star\% < 100\%$, and decreases above $p^\star$, i.e., the curves are unimodal. For $p > p^\star$, we have

$$\mathrm{Quar}(p) < \mathrm{Quar}(p^\star) \tag{9}$$

meaning that using CTAs only increase the average quarantine time below uptake $p^\star$ but not above it. We mention that $p^\star$ *depends on the uptake scenario*, and comes earlier with recommendation scenarios than with random uptake. On the last row of Fig 10 the same phenomenon is observed for various levels of infectiousness of the disease: unimodality in (9) is robust for these underlying networks. It disappears, however, for networks where either there are more super-spreaders or the geometry is less apparent (Fig 10, first five rows).

The same phenomena of unimodality is also present for the total time of the epidemic: for some $p^\dagger$ depending on the uptake scenario

$$\mathrm{Time}(p) < \mathrm{Time}(p^\dagger), \tag{10}$$

where, interestingly, the uptake percentage $p^\dagger\% \geq p^\star\%$ and $p^\dagger$ roughly coincides with the uptake percentage p^{stop}that is necessary to effectively eradicate the epidemic. A possible explanation for this phenomenon is that for these underlying networks and level of infectiousness, the necessary uptake percentage to push down the epidemic from supercritical to subcritical (via the CTA) is around $p^\dagger\% < 100\%$: around $p^\dagger \approx p^{\mathrm{stop}}$, the epidemic is in a near-critical state where it lingers around for quite a log time in certain regions before it can spread further. On the other hand, for lower uptake scenarios, the CTA is not very effective and the epidemic can spread quicker, while for higher uptake rates, the CTA is very efficient on stopping the epidemic. This can be observed since $p^\dagger$ roughly coincides with the uptake rate p^{stop} above which the total volume becomes close to 0, see the first vs last row of Fig 2. For other networks, this effect is less pronounced and can be seen only for lower level of infectiousness, see the first column of Figs 8 and 11.

The conclusion is that, in this particular setting, using CTAs at high uptake rates is effective not only for reducing the volume and maximum hospital load of the epidemic but also in reducing the average quarantine length and the length of the epidemic.

1.3.5 Other parameter values

The experiments in Fig 2 assumed an epidemic model (see Section 2) with a low rate of infection $\beta = 0.05$. This means that each infectious node infects each of its susceptible neighbours with probability β at every time-step (e.g., every day). Low rates of infection can be achieved by social-distancing measures. The results that we have mentioned so far are robust over other infection rates. See Figs 8, 9, and 10 where the rate of infection is varied, for volume, maximum hospital load and quarantine load.

1.4 Experiment 2: Varying the strictness of quarantine measures

It has been observed that the severity of symptoms varies with the individual [16, 17, 18, 19]. In our model, we assume that there is an underlying probability distribution that determines whether an individual would be symptomatic, in the case that this individual is infected, and also that there is an underlying probability distribution determining the severity of these symptoms. Policy makers may then impose quarantine or home isolation on individuals that have symptoms above a certain severity-threshold; setting the threshold is a political/economic decision. In our model, the strictness of quarantine measures is represented by a value $q \in [0, 1]$. This quantity q is the fraction of individuals whose symptoms would be so severe that they would go into home-quarantine if infected. Strict quarantine measures correspond to large values of q – in the extreme, taking $q = 1$ means that all infected individuals would go into home-quarantine. Recall that we work under the assumption that either a single virus is present in the population (so all relevant symptoms are caused by this virus) or testing is immediate and available to those who exceed the symptom severity threshold. This assumption is important for the results in this section.

For any given epidemic, there is an underlying probability q_{symp}, which is the proportion of infected individuals that experiences any symptoms upon infection (mathematically, the probability of showing symptoms upon infection). Implementing a quarantine-strictness $q > q_{\text{symp}}$ would require a program of mass testing. The implementation of such a testing program is beyond the scope of this paper, however we make two remarks:

(1) For many epidemics, the value of q_{symp} is fairly high (for COVID-19, the literature is varied, [20] finds a $q_{\text{symp}} \approx 0.3$ over all population, varying in age, [21, 22] find it ≈ 0.6 and ≈ 0.7, respectively, while [23] finds it as high as ≈ 0.85).

(2) our results indicate that setting quarantine-strictness even below q_{symp} can still be effective for decreasing the epidemic volume and maximum hospital load.

The social benefits of quarantining measures are nuanced. While quarantining clearly helps to reduce the volume of the epidemic (Vol) and the maximum hospital load (HMax), it also comes at the high societal cost of reduced work capacity. As a policy maker, one might like to balance the reduction in Vol and HMax against Quar, the average number of days spent in quarantine per person per year, over the course of the epidemic. In Experiment 1 we took the fixed value $q = 0.6$. In Experiment 2, we study how the setting of the quarantine-strictness q affects all of these KPIs. In this experiment, we vary the quarantining strength q and the CTA-uptake percentage $p\%$ for the four CTA-uptake scenarios. Our results are summarised in Fig 5, that we elaborate below.

1.4.1 Reducing total volume & hospital load by stricter quarantining

Increasing the quarantine-strictness q (i.e., requiring individuals to quarantine even with mild symptoms) reduces the number of individuals who are able to spread the virus, and hence reduces the volume and maximum hospital load of the epidemic. This is a very natural observation, valid across all parameter settings and underlying contact networks. The top row of Fig 5 shows how the volume of the epidemic decreases, for various CTA-uptake percentages $p\%$, as the quarantine-strictness q increases. We observe that for $q_1 > q_2$, for any fixed p and uptake scenario $\square = \mathrm{rand}, \mathrm{basic}, \mathrm{ring}, \mathrm{deg}$,

$$\begin{aligned} \mathrm{Vol}_{\square,p}(q_1) &\leq \mathrm{Vol}_{\square,p}(q_2), \\ \mathrm{HMax}_{\square,p}(q_1) &\leq \mathrm{HMax}_{\square,p}(q_2). \end{aligned} \tag{11}$$

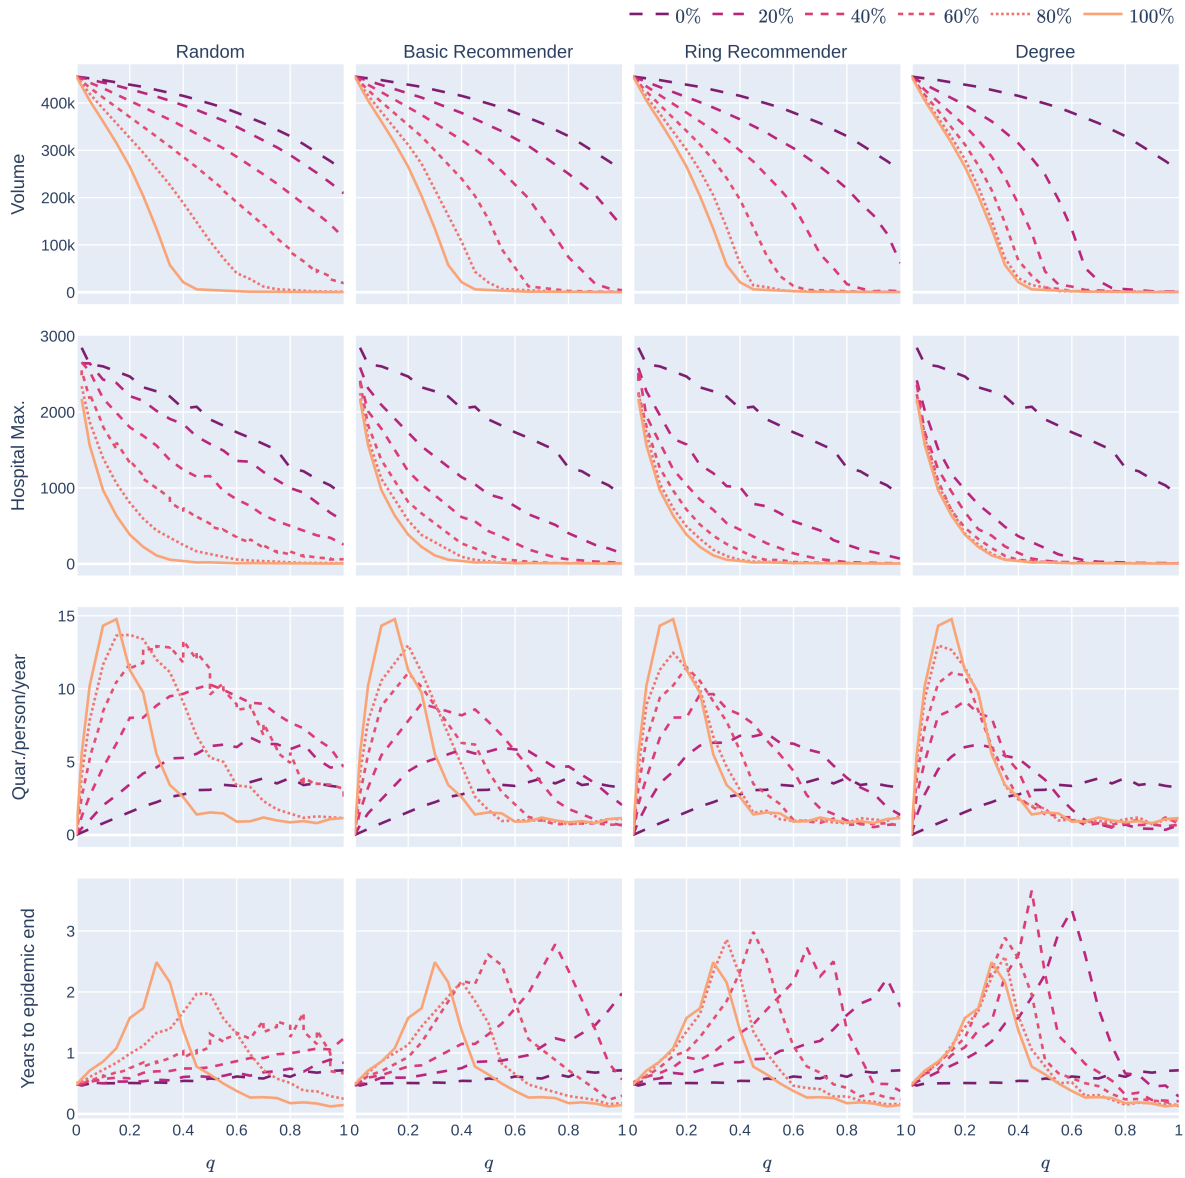

Figure 5: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios: the KPIs plotted against q, the quarantine strictness. We work here with a fixed underlying network, a GIRG with $\tau = 3.3$ (fewer super-spreaders) and $\alpha = 2.3$ (more geometric). Each column represents a given uptake scenario (random, basic-recommender and ring-recommender, and degree-targetted). Fig 12 below is qualitatively similar and considers an underlying network that is less geometric.

However, the rate of decrease varies across uptake scenarios. Figs 18, 19, 20 and 21 show that typically

$$\begin{aligned} \mathrm{Vol}_{\mathrm{deg},p}(q) < \mathrm{Vol}_{\mathrm{ring},p}(q) < \mathrm{Vol}_{\mathrm{basic},p}(q) < \mathrm{Vol}_{\mathrm{rand},p}(q), \\ \mathrm{HMax}_{\mathrm{deg},p}(q) < \mathrm{HMax}_{\mathrm{ring},p}(q) < \mathrm{HMax}_{\mathrm{basic},p}(q) < \mathrm{Hmax}_{\mathrm{rand},p}(q), \end{aligned} \tag{12}$$

by plotting the various uptake scenarios on the same diagram: the curve is steepest under degree-targetted uptake, ring recommendation is second best, while basic recommendation still performs better than random uptake. This is similar to what we have already observed in Section 1.3.

For network with more super-spreaders (Figs 13 and 14), these effects are even more exaggerated. It is interesting to observe that the curves of $\mathrm{Vol}_{\square}(p), \mathrm{HMax}_{\square}(p)$ for $\square = \mathrm{basic}, \mathrm{ring}, \mathrm{deg}$ drop so significantly compared to $\mathrm{Vol}_{\mathrm{rand}}(0), \mathrm{HMqx}_{\mathrm{rand}}(0)$, respectively, that it appears to be a discontinuity at 0. While we believe these are due to finite size effects, the observation that recommending helps significantly already at very low uptake rates remains valid.

1.4.2 Improving CTA efficacy by slightly increasing quarantine strength

We next study the shape of the curves $\mathrm{Vol}_{\square,p}(q)$ and $\mathrm{HMax}_{\square,p}(q)$ for fixed p as q varies and $\square = \mathrm{rand}, \mathrm{basic}, \mathrm{ring}, \mathrm{deg}$. Intuitively, we examine what happens when we slightly increase the quarantine strength q, especially in the mild-quarantine scenario where $q < 0.5$. It is important to note that without a CTA present ($p = 0\%$), HMax decreases linearly as a function of q, so for some s:

$$\mathrm{HMax}_{p=0\%}(q) \approx (1 - sq)\mathrm{HMax}_{p=0\%}(0).$$

Whenever a CTA is present, ($p\% > 0\%$) HMax becomes a *steeply decreasing convex* function of q for low q:

$$\mathrm{Vol}_p(q) - \mathrm{Vol}_p(q + \mathrm{d}q) \gg s\mathrm{d}s \cdot \mathrm{HMax}_{p=0\%}(0).$$

The take-home message is: a somewhat stricter quarantine rule can already lower the maximum hospital load and thus "*flatten the curve*" very well in the presence of a CTA. A similar observation is valid for $\mathrm{Vol}_p(q)$: the curve is concave for $p\% = 0\%$, while it becomes linear or convex when p increases.

When the value of q is small (so the quarantine measures are not strict), Fig 5 also shows that another way to flatten the curve is to increase the uptake percentage: the rate at which Vol and HMax decrease is much higher for large CTA-uptake percentages: when $p_1 > p_2$ and dq is a small positive number, then

$$\begin{aligned} \mathrm{Vol}_{p_1}(q) - \mathrm{Vol}_{p_1}(q + \mathrm{d}q) \gg \mathrm{Vol}_{p_2}(q) - \mathrm{Vol}_{p_2}(q + \mathrm{d}q), \\ \mathrm{HMax}_{p_1}(q) - \mathrm{HMax}_{p_1}(q + \mathrm{d}q) \gg \mathrm{HMax}_{p_2}(q) - \mathrm{HMax}_{p_2}(q + \mathrm{d}q). \end{aligned}$$

Thus a high(er) CTA uptake rate makes it possible to impose less strict quarantine measures and yet keep the volume and hospital load under control.

We mention that these findings are robust across underlying networks: while Fig 5 shows the KPIs on a fairly geometric network with not too many and not too large super-spreaders, Fig 12 shows the same KPIs for a different class of networks; where many long-range edges are present. The two scenarios are qualitatively similar. For network with more super-spreaders (Figs 13 and 14), these effects are even more exaggerated. It is interesting to observe that the curves of $\mathrm{Vol}_{\square,p}(q), \mathrm{HMax}_{\square,p}(q)$ for $\square = \mathrm{basic}, \mathrm{ring}, \mathrm{deg}$ drop so significantly compared to $\mathrm{Vol}_{\square,0}(0), \mathrm{HMAx}_{\square,0}(0)$, respectively, that it appears to be a discontinuity at 0 in q: seemingly, e.g. for $p = 20\%$

$$\mathrm{HMax}_{\square,20}(0) < \mathrm{HMax}_{\square,0}(0).$$

While we believe these are due to finite size effects, the observation that recommending helps significantly already at very low uptake rates remains valid.

1.4.3 Reducing quarantining by quarantining

Although it is valuable to increase quarantine-strictness in order to reduce Vol and HMax, one may wonder whether stricter quarantine measures require more social sacrifice. Our experiments demonstrate that the picture is not this simple: Row 3 of Fig 5 shows that, for every positive CTA-uptake percentage $p\%$ and every uptake scenario $\square = \mathrm{rand}, \mathrm{basic}, \mathrm{ring}, \mathrm{deg}$, there is a pessimal quarantine strictness $q^{\star}_{\square,p}$ that is strictly less than 1. In other words, the curves $\mathrm{Quar}_{\square,p}(q)$ are unimodal as a function of q, so for $q \neq q^{\star}_{p}$:

$$\mathrm{Quar}_{\square,p}(q) \leq \mathrm{Quar}_{\square,p}(q^{\star}_{\square,p}).$$

Thus increasing the quarantine-strictness q to a value above $q^{\star}_{\square,p}$ actually decreases the average number of days per year that people have to quarantine, as compared to the number that they would have to quarantine with value $q^{*}_{\square,p}$. This observation is most useful if the CTA-uptake percentage $p\%$ is high enough that $q_{\square,p} < q_{\mathrm{symp}}$, so that quarantine-strictness quantities $q > q^{*}_{\square,p}$ are easy to implement in reality. Note that $q^{\star}_{\square,p}$ also depends on the infectiousness of the disease. Stricter social distancing, corresponding to lower reproduction number, leads to lower value of $q^{\star}_{\square,p}$. Thus, when social-distancing measures are in place, it will be more possible to reduce the average time that people quarantine by setting $q > q^{*}_{\square,p}$.

A similar observation is valid for the total time of the epidemic. The curve $\mathrm{Time}(q)$ is unimodal in q: there is a $q^{\dagger}$ depending on the uptake percentage $p\%$ and the uptake scenario such that the total duration of the epidemic is maximised at $q^{\dagger}$,

$$\mathrm{Time}_{\square,p}(q) \leq \mathrm{Time}_{\square,p}(q^{\dagger}_{\square,p}), \tag{13}$$

and the point of switching comes earliest for degree-targetted uptake, followed by ring and basic recommendation, and finally random uptake, i.e.,

$$q^{\dagger}_{\mathrm{deg},p} \leq q^{\dagger}_{\mathrm{ring},p} \leq q^{\dagger}_{\mathrm{basic},p} \leq q^{\dagger}_{\mathrm{rand},p}.$$

This phenomenon is similar to that observed in Section 1.3.4, by observing that $q^{\dagger} \approx q^{\mathrm{stop}}$ coincides with the necessary quarantine strictness to make the epidemic turn from super to subcritical. Such a value of q exists (at least in theory, for not too high levels of infectiousness) since if every individual quarantines who is infected, the virus has little chance to spread. At the critical proportion $q^{\dagger}$ the epidemic is in a near-critical phase and 'lingers' for a long time in the population.

These phenomena can also be observed on other graphs with more long-range edges, see the last two rows of Fig 12.

For network with more super-spreaders (Figs 13 and 14), the situation is somewhat different. The average quarantine time curve as a function of q seems to *stabilise* rather than peak. In this case, excluding very low and high values of q, we see

$$\mathrm{Quar}_{\square,p}(q) \approx t_{\square,p}$$

for some constant $t_{\square,p}$. Although the time of the epidemic still roughly satisfies Equation (13) in these networks, the situation is more straightforward. For small values of q the time of the epidemic increases with q. Then there is a wide range of values of q (independent of the uptake percentage $p\%$) where the time of the epidemic stabilises around its maximum value (with plenty of low-order variance). Then for q outside of this range, the time of the epidemic decreases with q.

1.4.4 Geometry helps

Fig 15 shows the effect of geometry on the volume and maximum hospital load as a function of the quarantine strictness function q, for two levels of infectiousness (high and low). Observe that networks with more geometry ($\alpha = 2.3$) have lower volume and lower maximum hospital load than networks with many long-range edges and hence less geometry ($\alpha = 1.3$), in this respect, the conclusion is similar to that in Experiment 1, see Section 1.3.3.

Fig 16 shows that in case there are many super-spreaders in the network ($\tau = 2.3$), these form a well-connected core so that the value of α becomes insignificant: in this case geometry is not described by this parameter. See more on this in Section 3.1 below.

1.4.5 Other parameter values

The experiments in Fig 5 assumed an epidemic model (see Fig 2) with a low rate of infection $\beta = 0.05$. As we noted earlier, low rates of infection can be achieved by social-distancing measures. Fig 15 repeats the experiment for other values of β.

1.5 Remark: Social distancing and basic reproduction number.

As one would expect, social-distancing measures, which reduce the probability that any particular interaction leads to an infection, increase the effectiveness of CTAs. This is observed in our experiments — in all scenarios, the volume is monotonic in β. See Fig 17. Similar monotonicity for the maximum hospital load can be observed by comparing the columns in Fig 9.

2 The epidemic model

The epidemic model that we use is a refinement of the agent-based version of the discrete-time SEIR model. The spreading process changes at discrete time steps, $t = \{0, 1, 2, \dots\}$, each time step corresponds to, say, a day. We fix the network G in advance. We think of nodes in the network as individuals, and denote the set of nodes by $\mathcal{V}$. The neighbours of a node u are nodes with a direct connection (also called link, or edge) to u. A connection may correspond to a friendship or an acquaintance, or simply a contact event.

As we explained in the context of Experiment 2, the model has a parameter $q \in [0, 1]$ corresponding to quarantine-strictness. The quantity q is the fraction of individuals whose symptoms would be so severe that they would go into home-quarantine if infected. As we already explained, determining the value of q is a social/economic decision. The model reflects this decision as follows: For each node, independently, and in advance of the epidemic simulation, it is determined randomly, with probability q, whether the individual corresponding to this node will, upon infection, develop symptoms exceeding the socially-determined symptom severity threshold, and then go into home-quarantine. The set of nodes which will do so is called $\mathcal{V}_{\text{sev}}$ and the set of nodes which will not is called $\mathcal{V}_{\text{mild}}$. During the epidemic, nodes in $\mathcal{V}_{\text{sev}}$ that become infected will go into home-quarantine. One may worry about uncooperative individuals who do not self-isolate, despite having sufficiently severe symptoms. We do not need to adjust to model to account for these, since the quarantine-strictness q can be viewed as the fraction of individuals that (1) would develop sufficiently-severe symptoms, and (2) would be willing to self-isolate. Note that the actual distribution of the symptom-strength in the population is irrelevant to the problems that we study; the only relevant parameter is the final proportion q. The assumption

that we do we make, however, is that each individual has an independent symptom-strength variable, making $\mathcal{V}_{\text{sev}}$ a random set of nodes.

Definition 1 (SEISeR model)**.** *Determine the node sets $\mathcal{V}_{\text{sev}}$ and $\mathcal{V}_{\text{mild}}$. Nodes in $\mathcal{V}_{\text{sev}}$ can be in five possible states:* susceptible (S), exposed (E), infectious but not yet showing severe symptoms (I), infectious and showing severe symptoms (Se), or removed (R)*, while nodes in $\mathcal{V}_{\text{mild}}$ can only be in state $(S), (E), (I)$ or (R). At $t = 0$, a subset of nodes $\mathcal{E}_0 \subset \mathcal{V}$ is put in state (E), all other nodes are in state (S) (susceptible). The discrete time dynamics between the states are described as follows (see also Fig 6):*

- **Infecting:** *Each infectious node (in state (I)),* infects *each of its susceptible neighbours within the network with probability β at every time-step. Infections to different neighbours happen independently. These neighbour nodes enter state (E), exposed to the virus.*
- **Exposed → Infectious (not (yet) severely symptomatic)** *When exposed to the virus, each node becomes infectious (transitions to state (I)) with probability γ at every time-step, independently of other nodes.*
- **Infectious → Severely Symptomatic** *When a node in $\mathcal{V}_{\text{sev}}$ is in state (I), (infectious but does not show severe symptoms yet), it transitions to state (Se) with probability $\gamma_{\text{sev}} = 1/2.5$ at each time-step.*
- **Infectious → Removed** *When a node in $\mathcal{V}_{\text{mild}}$ is in state (I), infectious but does not show severe symptoms, it heals and transitions to state (R) with probability $\eta_{mild} = 1/7$ at each time-step.*
- **Severe Symptomatic → Removed** *When a node is showing severe symptoms, it* heals *with probability $\eta_{sev} = 1/4.5$ at every time-step, independently of other nodes. Upon healing, the node enters state R (removed).*

We emphasise that in our model, severely symptomatic nodes do not infect anymore. We assume that they self-quarantine until being removed. In reality, some infectious individuals may choose not to self-quarantine, but we have already incorporated this into the model Recall that $\mathcal{V}_{\text{sev}}$ contains only nodes that would develop sufficiently severe symptoms, and would go into home-quarantine.

To keep the parameter space tractable, we set $\gamma = 1/3$, $\eta_{\text{mild}} = 1/7$, $\gamma_{\text{sev}} = 1/2.5$, and $\eta_{\text{sev}} = 1/4.5$. Thus, in the model, after being exposed to the virus, it takes on average 3 days to become infectious and a further 7 days to heal completely for each individual. Nodes in $\mathcal{V}_{\text{sev}}$ infect for an average of 2.5 days before showing symptoms and moving to isolation (that lasts until their perfect healing) while nodes in $\mathcal{V}_{\text{mild}}$ are infectious for an average of 7 days before removal. These values are based on empirical findings e.g. in [24, 25, 1, 26], where we emphasise that our results are robust in these parameters. Experiment 2 varies the value of q, which is the expectation of $|\mathcal{V}_{\text{sev}}|/|\mathcal{V}|$.

Next, we describe our model for the CTA. First we give the details of the four CTA-uptake scenarios from Section 1.3 then we define the necessary modifications to the dynamics of the SEISeR model above to incorporate the presence of a CTA.

Definition 2 (Random, basic-recommender, ring-recommender, and degree-targetted CTA uptake)**.** *Before the epidemic starts, we determine the sets $\mathcal{V}_u, \mathcal{V}_{nu}$ of users and non-users of the application. We denote by $p = 100|\mathcal{V}_u|/|\mathcal{V}|$ the empirical uptake percentage. We say that the user uptake is*

- **Random***: When the nodes that use that application are chosen uniformly at random.*
- **Degree-targetted***: When the $p \cdot |\mathcal{V}|/100$ nodes with the highest degree (number of connections) are targetted to use the CTA.*

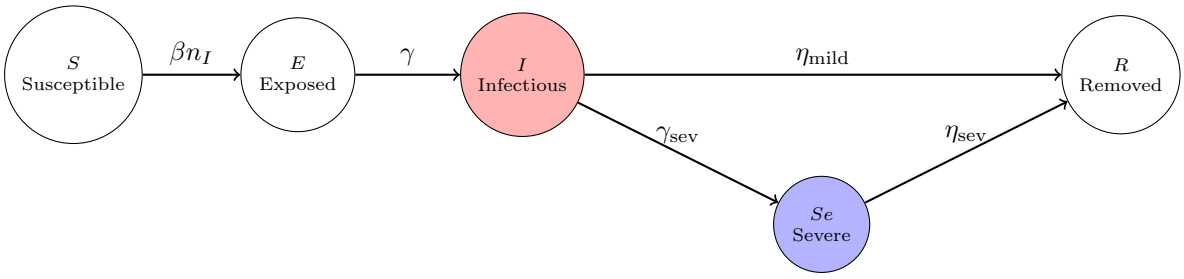

Figure 6: An illustration of the dynamical change of states. On the first transition arrow, n_I denotes the number of neighbours in state (I) of a susceptible node. Only nodes in $\mathcal{V}_{\text{sev}}$ transition to state (Se). By setting $1/\eta_{\text{mild}} = 1/\gamma_{\text{sev}} + 1/\eta_{\text{sev}}$ the average duration of having the virus in an individual's system is $1/\eta_{\text{mild}}$ for each node. Only nodes in the red circle can infect, while nodes in the blue circle stay in home isolation.

- **Basic-Recommender:** *When an initial set of nodes is chosen uniformly at random. Each node in this initial set uses the CTA, and also recommends it to one of its neighbour chosen uniformly at random, who will also use the CTA. The resulting number of users is* $p \cdot |\mathcal{V}|/100$
- **Ring-Recommender:** *When an initial set of nodes is chosen uniformly at random. Each node in this initial set uses the CTA, and also recommends it to all of its neighbour, half of whom (chosen uniformly at random) will also use the CTA. Again, the resulting number of users if* $p \cdot |\mathcal{V}|/100$.

We observe that degree-targetted uptake is an unrealistic scenario to consider, since it requires a complete knowledge of the contact network. However, studying the performance of this hypothetical scenario gives a baseline comparison for determining the potential performance of a given uptake percentage. We also emphasise that in the basic-recommendation and ring-recommendation uptake scenarios, the performance of the final uptake percentage $p\% = 100 \cdot |\mathcal{V}_u|/|\mathcal{V}|\%$ is compared to the performance of other uptake scenarios with the same $p\%$. In other words, the basis for comparison is not the proportion of the initially chosen set. Instead, it is the final uptake percentage. See Fig 1 for the relation between the initial and final uptake percentage for the four networks.

Next, we describe the necessary modifications to the SEISeR model in the presence of a CTA. Informally, a CTA user, upon either being tested positive or simply showing sufficient symptoms, will notify its contacts to stay in quarantine for a duration of T days ($T = 14$ in most countries), regardless of whether these contacts themselves show symptoms.

Definition 3 (SEISeR with a contact-tracing application (CTA))**.** *In addition to the original states (S), (E), (I), (Se) and (R) of the SEISeR model, there are five new states (NS), (NE), (NI), (NSe) and (NR), corresponding to "notified" versions of the original states — we refer to these as "N-states" and we refer to the original states as "O-states". Only CTA-users (nodes in $\mathcal{V}_{\text{u}}$) will ever enter the N-states. They enter the N-states when they receive a notification, via the CTA. They stay in the N-states until $T = 14$ days have elapsed since the last notification received. While in N-states, they self-quarantine, so cannot spread infection.*

The transitions between the N-states are exactly the same as the transitions between the original versions of these states, except that there is no transition from (NS) to (NE), because the notified nodes in state (NS) will be quarantining, so will not become exposed. See the black arrows denoting transitions in Fig 7.

All nodes start in O-states, exactly as in the SEISeR model. After each time-step t, nodes move between the O-states and the N-states as follows:

- **Notification:** *If a node in $\mathcal{V}_u$ starts to show sufficient symptoms at time t (i.e., either it transitions from (I) to (Se) at time t or it transitions from (NI) to (NSe) at time t) then it sends a notification to all of its network-neighbours that are in $\mathcal{V}_u$.*
- **Moving to N-states:** *If a node in an O-state receives a notification, it moves to its corresponding N-state (following the corresponding red arrow in Fig 7).*
- **Moving to O-states:** *If a node is in an N-state and it has not received a notification for $T = 14$ time-steps, then it moves to its corresponding O-state (following the corresponding blue arrow in Fig 7).*

Justification of modelling choices and robustness: We make a few comments to justify our modelling choices. Nodes in the N-states and in state (Se) are assumed to be self-quarantining, so they cannot spread infection. These are the blue nodes in Fig 7). Only nodes in state (I) (coloured red) can spread infection.

Sending Notifications: The reason that nodes entering state (NSe) send notifications (even though they are already self-quarantining) is that they may have spread infection before entering the N-states (and starting to quarantine): their neighbours may be in exposed or infectious states when receiving the notification. In fact, while this model does not include testing, such a scenario also happens in real life when an individual, already part of a contact tracing chain, receives a positive test. In this case, the rest of the contacts of this individual need to be notified as well.

Removed nodes in quarantine: The underlying SEISeR model assumes that individuals can only be infected once (so "removed" nodes do not become susceptible again). Despite this, in light of possible re-infection, many countries (e.g. the Netherlands, where two of the authors are located) do ask individuals who are notified to self-isolate, even if they have already have Covid-19. While transitioning from (R) to (NR) has no effect on the course of the epidemic, it does have an effect on the total population's effective workforce (which we also study). Indeed, the self-isolating blue nodes represent a social cost, since all of these nodes are self-isolating, and cannot go to work. A node who is say, exposed when it receives a notification via a CTA may go through all the phases (NE), (NI), (NSe), (NR) before the T days are over. In this case it cannot become simply removed and go back to work immediately, but has to wait out the T days and then transition to state (R). For the same reason we make a distinction between nodes in (Se) and (NSe): nodes in (Se) may leave home isolation immediately after healing, while nodes in (NSe) have to wait for the necessary total T days even after healing to be able to leave isolation.

Multiple notifications extend the quarantine time: We make a further comment on why we choose to extend the quarantine by T days after an additional notification is sent to a node already in an N-state. The reason for this is again based on real life experience: since not all contacts and hence notifications lead to infection, a second notification should not be ignored. Whether individuals are actually willing to comply with such rules belongs to behavioural science and is out of the scope of the current study. We assume an idealised scenario where each node in $\mathcal{V}_u$ complies with the rules. Varying the size of V_u captures the effect of less compliance as well.

Robustness of choices: To conclude the justification, we emphasise that we conduct a *qualitative and comparative* study. The model is sufficiently robust that the conclusions of our study also apply under small changes to the model. For example, if an uptake scenario performs better than another with regard to this particular model, it will also perform better in a model with slight changes, e.g.

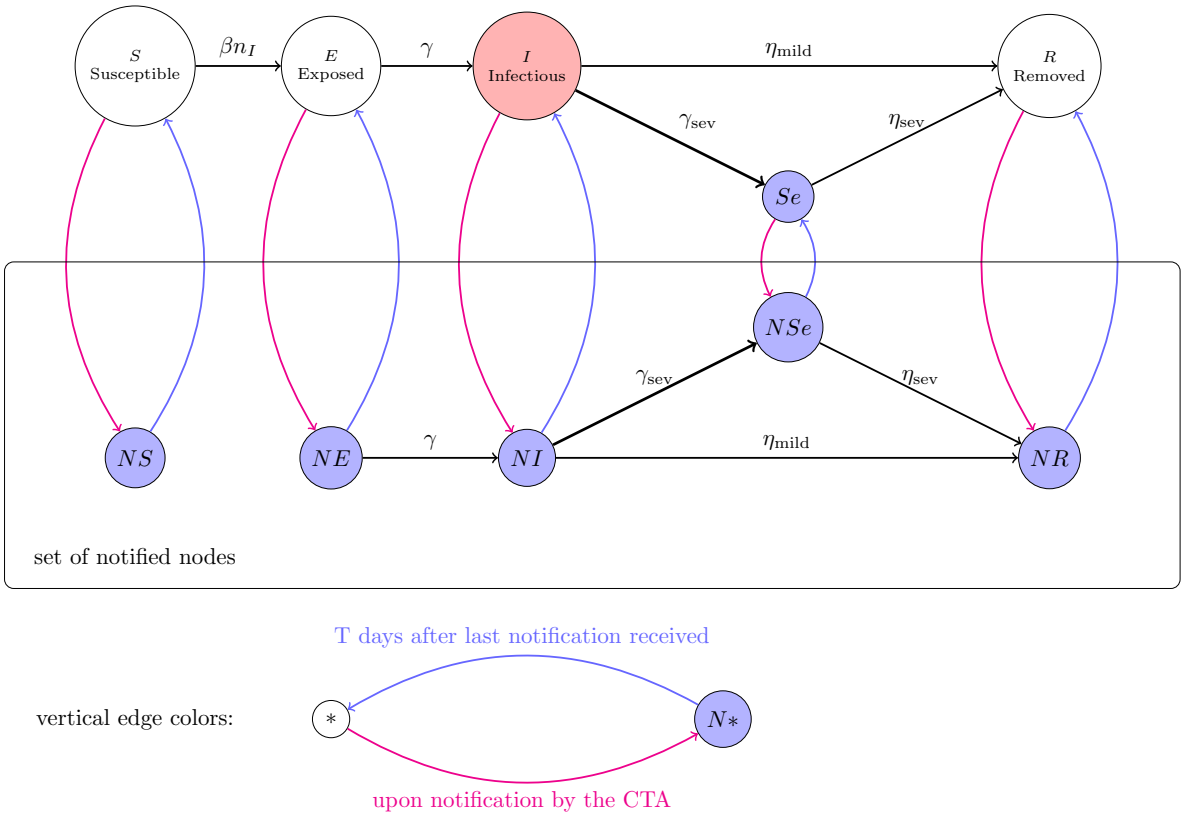

Figure 7: An illustration of the dynamical state changes in the presence of a contract-tracing application. The black arrows represent the transitions of the SEISeR model and the corresponding transitions in the notified versions of the states. After each time step, nodes move between the O-states (S), (E), (I), (Se) and (R) and the N-states (NS), (NE), (NI), (NSe) and (NR). Moves to N-states (depicted by red arrows) follow from CTA notifications. CTA users (nodes in $\mathcal{V}_u$) receive such notifications when neighbouring CTA users show sufficient symptoms, thus transitioning from (I) to (Se) or from (NI) to (NSe). Nodes in blue states are required to self-quarantine so cannot spread infection. Moves to O-states (depicted by blue arrows) occur when a node has received no new notifications for $T = 14$ days.

a model that does not extend the time of home isolation when a second notification is received, or a model that does not require already-removed nodes to go into quarantine.

Next we describe the underlying contact network of the agent-based epidemic.

3 Agent-based population model: spatial scale free network models

To be able to carry out our study, we represent the underlying contact network of individuals (nodes, or agents) as a network. Based on the universality results in [15], our primary choice for the underlying contact network is a mixture of pure geometric and purely random network models, called Geometric Inhomogeneous Random Graphs, which possess geometric features and can match the statistical properties of real human contact networks, such as degree distributions and clustering, as well as long-range connections. For baseline comparison, we also study the case when the underlying network is a "mean-field network", ignoring the spatial component. For this scenario we use the *configuration model,* which can mimic the local statistical properties of real human contact networks, such as degree distributions.

3.1 Geometric Inhomogeneous Random Graphs

We model populations embedded in geometric space using Geometric Inhomogeneous Random Graphs (GIRGs). Recently, several spatial random graph models were developed to mimic properties of real networks features: hyperbolic random graphs [27, 28, 29], scale-free percolation [30], and GIRGs [31, 4, 32]. The qualitative behaviour of these models is the same, in fact, they can be unified into a general model containing all three models as special cases. For sake of simplicity we decided to work with GIRGs.

The advantage of these spatial models are that they can mimic several aspects of real contact networks: individuals are embedded in space, just like in real life, allowing for local community structures to be present in the contact network. This results in strong clustering [32, 33]. Beyond mimicking the degree distribution of real networks (e.g. high variability) GIRGs also incorporate connections (edges) bridging spatial distance on all scales, short as well as long-range edges. The model is very flexible, and for properly chosen parameter settings it is *scale-free* in two respects: both in spatial distance that edges cover, and in node-degree variability [32].

Contact- or activity networks of humans have been found to show similar behaviour, including heavy-tailed degree distributions, strong clustering, and community structures [34, 35, 36, 37, 38], as well as heavy-tailed distance-distribution for edges, see references in [39]. For epidemic spread specifically, lighter tailed degree distributions such as negative binomial have been observed e.g. in [1, 25]. GIRGs also allow for the incorporation of such distributions, by setting light-tailed fitness distributions.

The following definition from [15] is general, and in its last sentence we specify it for a GIRG. The underlying space can be the earth's surface, or $\mathbb{R}^2$, the two-dimensional Euclidean space. We denote by $x \wedge y$ the minimum of two numbers x, y.

Definition 4 (Geometric Inhomogeneous Random Graph (GIRG))**.** *Fix $N \geq 1$ the number of nodes. Assign to each node $u \in \{1, 2, \ldots, n\}$ a* fitness $w_u > 0$*, and a* spatial location $\Phi(u)$*. Fix $\alpha > 0$. For any pair of nodes u, v with fixed $w_u, w_v, \Phi(u), \Phi(v)$, connect them by an edge with probability*

$$\mathbf{Prob}(u \textit{ is connected to } v) \asymp \left(\frac{w_u w_v}{\mathrm{dist}(\Phi(u), \Phi(v))^2}\right)^{\alpha} \wedge 1, \tag{14}$$

where dist *is a distance function, i.e., it measures the distance between the locations* $\Phi(u)$ *and* $\Phi(v)$*. When* $\Phi(u)$ *is chosen uniformly in a box of volume* N*, we obtain a GIRG.*

Mathematically, the distance function dist in Definition 4 is a metric on the underlying space, (e.g., a metric such as $\|\cdot\|_2$). In the definition, the formula is given for a two-dimensional model. In higher dimensions, the exponent 2α in the denominator would change to $\dim \cdot \alpha$, with dim being the dimension of the model.

GIRGs have a natural interpretation: the fitnesses express the ability of nodes to have many connections, Φ embeds them in space, and α is the *long-range* parameter: the smaller α is, the more the model favours longer connections. The parameter space of GIRG is rich enough to model many desired features observed in real networks:

(a) extreme variability of the number of neighbours (degrees),
(b) connections present on all length-scales,
(c) small and ultra-small distances,
(d) strong clustering,
(e) local communities.

In real-life networks, an extreme variability of node degree is often observed, see [34, 40, 37]. Extreme node degree variability results in the presence of a few individuals with extreme influence on spreading processes, the *hubs or super-spreaders* [37]. Mathematically, this extreme degree variability can be expressed using the empirical distribution of node degrees, that follows a *power-law*:

$$\mathbf{Prob}(\deg(u) \geq x) \asymp \frac{1}{x^{\tau-1}}. \tag{15}$$

for some $\tau > 2$. Setting a power-law fitness distribution for w_u in a GIRG yields that the degrees satisfy (15). As we shall see, the parameter α in (14) determines the presence of long-range connections: as α gets smaller, it is more likely that there are long-range connections. Thus, the two parameters of the model are τ and α.

The parameters τ and α determine the average graph distance between nodes in a GIRG [41, 31, 30, 42]. Letting $d_G(u,v)$ denote the graph distance between u and v (the number of connections on the shortest path between nodes u and v), the average distance is given by

$$\overline{\mathrm{Dist}}(N) = \frac{1}{\binom{N}{2}} \sum_{u,v} d_G(u,v) = \begin{cases} \Theta(\log\log N) & \text{when } \tau \in (2,3), \alpha > 1 \\ \Theta\big((\log N)^{\zeta}\big) & \text{when } \tau > 3, \alpha \in (1,2) \\ \Theta(\sqrt{N}) & \text{when } \tau > 3, \alpha > 2. \end{cases} \tag{16}$$

As [15] has found, these average distances give a good indication as to the speed at which epidemics can spread on a network, as well as an indication of the shape of the epidemic curve [15].

Equation (16) demonstrates that increasing α decreases the number of long-range edges when $\tau > 3$. To see this, note that when $\tau > 3$, the node fitnesses w_u (which satisfy the same power law as (15)) are typically less than $\sqrt{N}$ so the ratio

$$R_{u,v} := \frac{w_u w_v}{\mathrm{dist}(\Phi(u), \Phi(v))^2}$$

is typically less than 1. In this case, since $R_{u,v}$ is raise to the power α in the connectivity probability (Equation (14)), increasing α decreases the probabilities of long-range edges. However, when $\tau \in (2,3)$, node pairs with ratio $R_{u,v} > 1$ happen frequently enough, and for these pairs the connection

probability is unaffected by the change of α. So, networks with $\tau \in (2,3)$ has many long-range edges irrespective of α.

Comparing (16) to the average distance in a-geometric network models such as the configuration model (see (17) below) and to the polynomial average distance in lattice models, we observe that geometry and long-range connections play a role in GIRGs with $\tau > 3$, and the GIRG model interpolates between the small-world configuration model and the lattice.

Clustering in a network corresponds to the phenomenon known as "a friend of a friend is also likely to be my friend", and it manifests in the presence of *triangles* in the network. Communities and clustering are naturally present in GIRGs, because the model favours connections between nodes that are close to each other in space. Thus, a GIRG incorporates all five desired features.

For the spread of information or infections in GIRGs and similar networks, it is known that large-degree nodes and many triangles have opposing effects. On the one hand, nodes of large degree (also called hubs, super-spreaders, or influencers) contribute to fast dissemination, and foster explosive propagation of information or infections [43, 44, 45, 46]. On the other hand, clustering and community structures provide natural barriers that slow down the process, while long-range edges accelerate the spread [47, 48, 49, 50, 51].

3.2 The configuration model and other a-geometric network models

A-geometric random network models are often used as null-models to compare network data to purely random networks. Here we describe a commonly used model, the configuration model This is a well-studied random-graph model, dating back to Bender and Canfield [52], Bollobás [53] and Wormald [54], and introduced in network science by Molloy and Reed [55, 56]. The main advantage of this simple model is that it can mimic the degree distribution of real-life networks.

Definition 5 (Configuration model)**.** *Fix $N \geq 1$ the number of nodes. Prescribe to each node $u \in \{1, 2, \ldots, n\}$ its node-degree $\deg(u) \geq 0$, so that the total degree $h_N := \sum_{u \leq N} \deg(u)$ is even. To form a graph, to each node u we assign $\deg(u)$ half-edges. The half-edges are then paired uniformly at random to form edges. Self-loops and multiple edges are then erased from the resulting multi-graph.*

It is known that the erasure of self-loops and multiple edges does not affect the empirical degree distribution of the model in any significant way [57].

In comparison to GIRGs, the configuration model can incorporate the desired properties (a), (c) but not (b), (d) and (e). For (a), highly variable degree distributions can be incorporated since the degree sequence is prescribed.

Small distances are an important feature of a-geometric network models in general, aka the *small-world property*. Heuristically, this means that two arbitrary nodes in the network can be connected via very short paths, using only a few connections. Mathematically, when setting power law degree distribution, the average distance in the model scales as [58, 59]

$$\overline{\mathrm{Dist}}(N) = \begin{cases} \Theta(\log N) & \text{when } \tau > 3 \\ \Theta(\log \log N) & \text{when } \tau \in (2,3), \end{cases} \tag{17}$$

The second row, showing doubly-logarithmic distances, is often called ultra-small world. When $\tau > 3$, the shape of the epidemic curve is shown to converge in the large network limit, see [60, 61, 62].

While the configuration model and other a-geometric models easily accommodate power-law node degrees, and have the small-world property, they neither contain communities nor clustering [40, 33].

Other popular a-geometric models include the Chung-Lu and the similar Norros-Reittu model [63, 64, 65], and preferential attachment models (also known as the Barabási-Albert model) [66]. In the Chung-Lu and the Norros-Reittu model, only the expected degrees of nodes are prescribed, rather than their exact degree, while in preferential attachment models, the power law exponent can be tuned. Using similar parameter settings, these models behave qualitatively similarly to the configuration model, although numerical values might differ, see [57, 67, 40, 68, 69] for references. As a result, we choose the configuration model for baseline comparison.

In [15] we have compared the behaviour of a similar epidemic on the configuration model and on GIRGs. There, we found the epidemic curves are similar for models with the same degree distribution whenever this distribution has a power law with some $\tau \in (2, 3)$. Furthermore, GIRGs with any parameter $\tau > 3$ and long-range parameter $\alpha \in (1, 2)$ are comparable with a configuration model with similar degree distribution. This phenomenon is best explained by comparing the average distance, i.e., (16) to (17): when $\tau \in (2, 3)$, both models have doubly-logarithmic distances and thus allow for very quick spread, while when $\tau > 3$ and $\alpha \in (1, 2)$, poly-logarithmic distances in GIRGs are comparable to logarithmic distances in the a-geometric counterpart, hence the similarity of the curves.

While we expect the same universality to occur for the more complicated SEISeR model, for baseline comparison, we also carry out some of our simulations on the configuration model with power law degree distributions, as in (15).

4 Supplementary figures

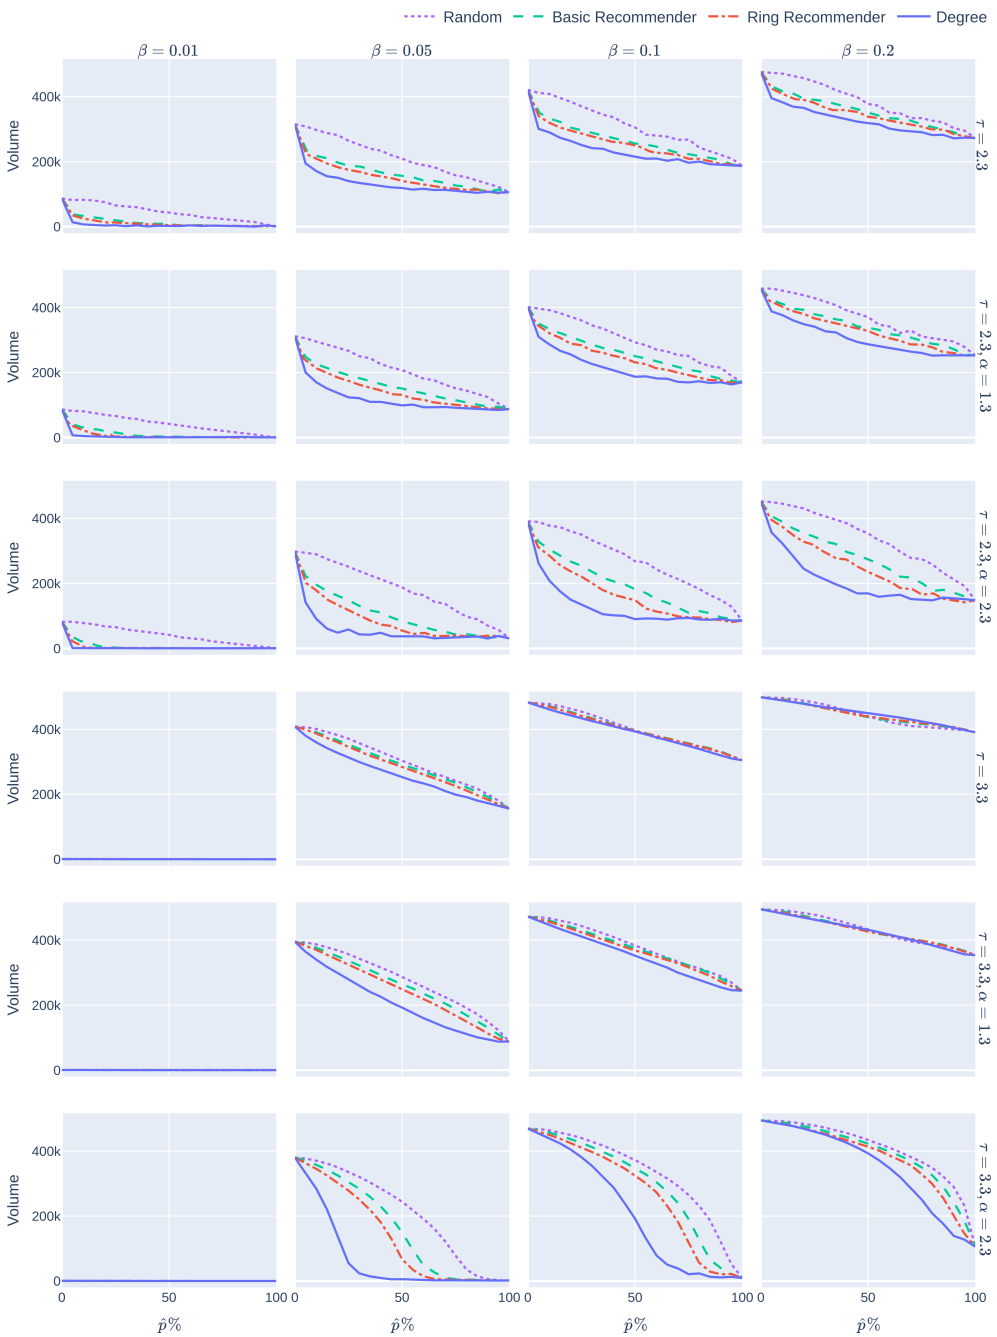

Figure 8: $\mathrm{Vol}(p)$: The total volume of the epidemic as a function of the uptake percentage $p\%$ in four different uptake scenarios with $q = 0.6$, on four different networks. Dotted: $\mathrm{Vol}_{\mathrm{rand}}(p)$, dashed-dotted $\mathrm{Vol}_{\mathrm{basic}}(p)$, dashed $\mathrm{Vol}_{\mathrm{ring}}(p)$, continuous $\mathrm{Vol}_{\mathrm{deg}}(p)$. The x-axes vary the uptake percentage p from 0 to 100%, the y-axes shows the total number of nodes reached by the epidemic. Different columns indicate varying infection probabilities β. Each row represents a different network. First two columns: GIRGs with many super-spreaders ($\tau = 2.3$), third column: GIRG with not many super-spreaders but many long-range connections ($\tau = 3.3$, $\alpha = 1.3$), fourth column: GIRG with not many super-spreaders and also mostly local connections($\tau = 3.3$, $\alpha = 2.3$). The volume data in Fig 2 corresponds to the $\beta = 0.05$ column.

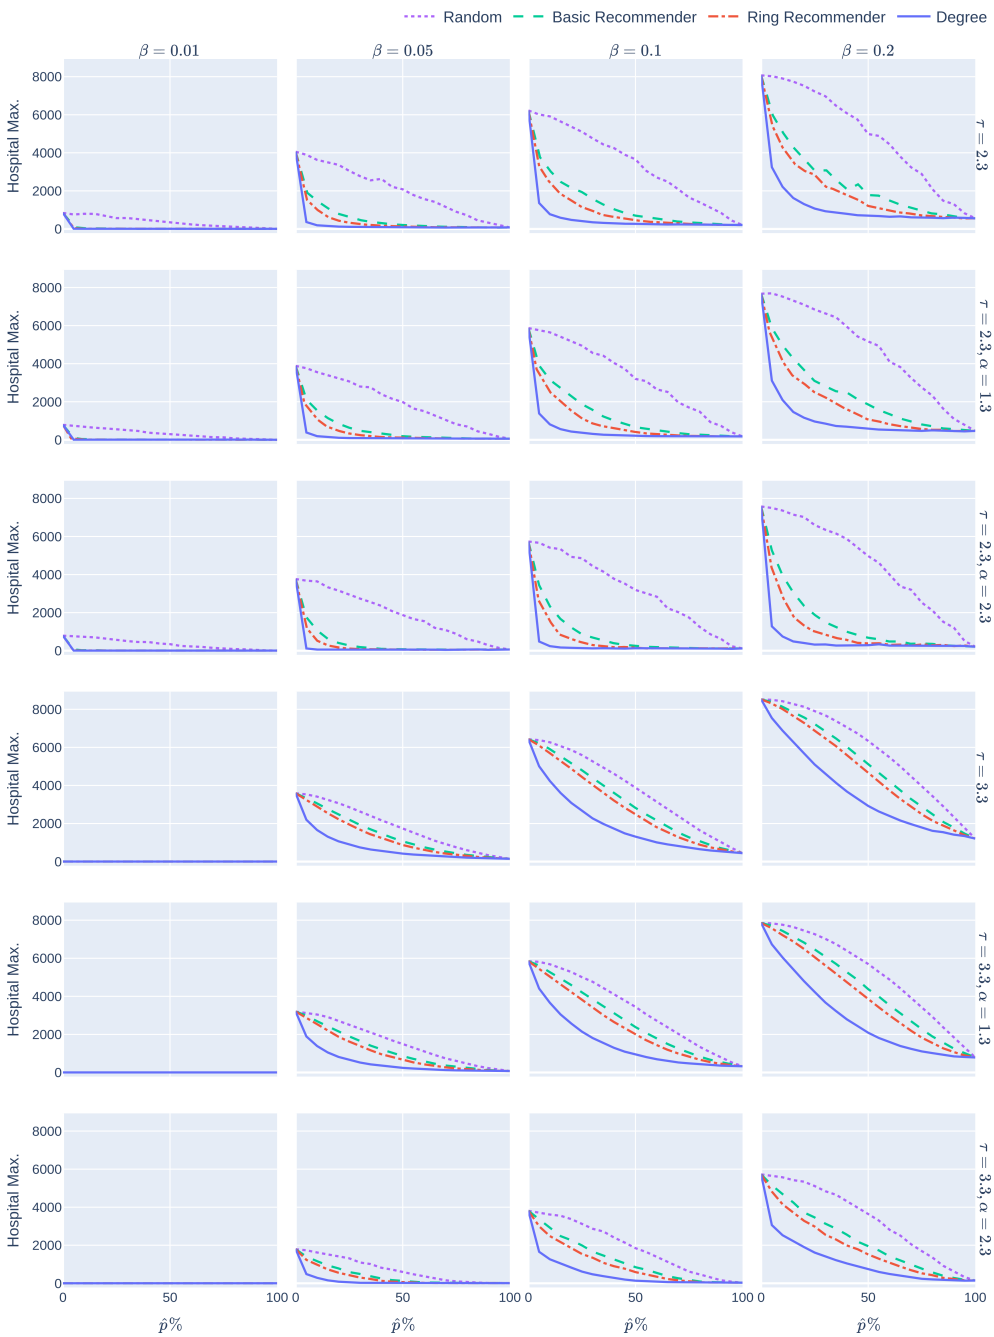

Figure 9: $\mathrm{HMax}(p)$: The maximum hospital load as a function of the uptake percentage $p\%$ in four different uptake scenarios with $q = 0.6$, on four different networks. Dotted: $\mathrm{HMax}_{\text{rand}}(p)$, dashed-dotted $\mathrm{HMax}_{\text{basic}}(p)$, dashed $\mathrm{HMax}_{\text{ring}}(p)$, continuous $\mathrm{Hmax}_{\text{deg}}(p)$. The x-axes vary the uptake percentage p from 0 to 100%, the y-axes shows the maximum hospital load. Different columns indicate varying infection probabilities β. Each row represents a different network. Note that ring recommendation outperforms degree-targetted uptake in row 3, column 5. First two columns: GIRGs with many super-spreaders ($\tau = 2.3$), third column: GIRG with not many super-spreaders but many long-range connections ($\tau = 3.3$, $\alpha = 1.3$), fourth column: GIRG with not many super-spreaders and also mostly local connections($\tau = 3.3$, $\alpha = 2.3$). The HMax data in Fig 2 corresponds to the $\beta = 0.05$ column.

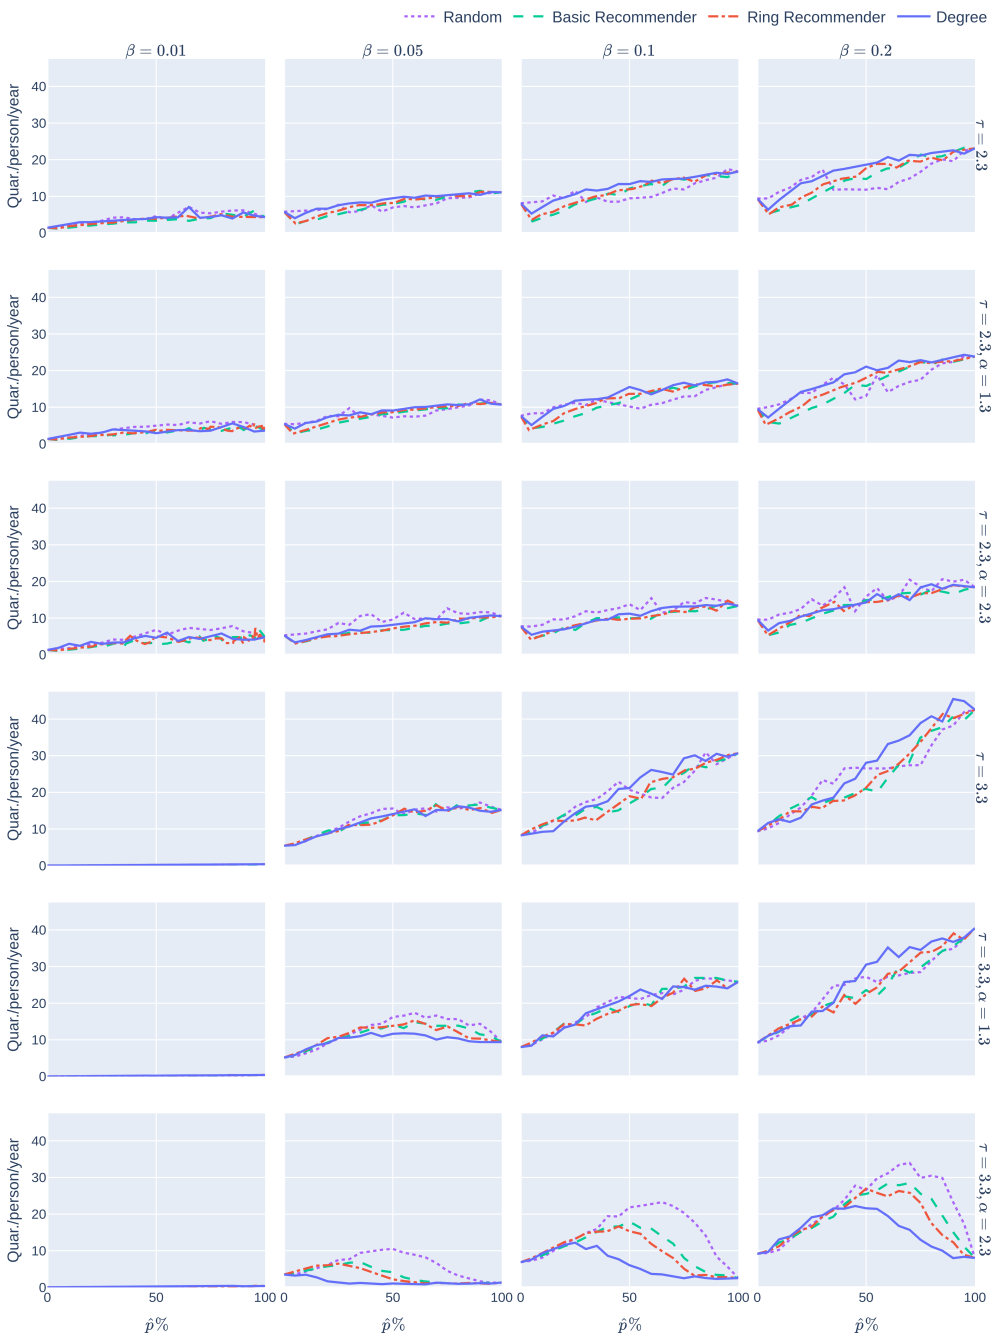

Figure 10: $\mathrm{Quar}(p)$: The average number of days spent in quarantine per person per year, over the course of the epidemic as a function of the uptake percentage $p\%$ in four different uptake scenarios with $q = 0.6$, on four different networks. Dotted: $\mathrm{Quar}_{\mathrm{rand}}(p)$, dashed-dotted $\mathrm{Quar}_{\mathrm{basic}}(p)$, dashed $\mathrm{Quar}_{\mathrm{ring}}(p)$, continuous $\mathrm{Quar}_{\mathrm{deg}}(p)$. The x-axes vary the uptake percentage p from 0 to 100%, the y-axes shows $\mathrm{Quar}(p)$. Different columns indicate varying infection probabilities β. Each row represents a different network. First two columns: GIRGs with many super-spreaders ($\tau = 2.3$), third column: GIRG with not many super-spreaders but many long-range connections ($\tau = 3.3$, $\alpha = 1.3$), fourth column: GIRG with not many super-spreaders and also mostly local connections($\tau = 3.3$, $\alpha = 2.3$). The Quar data in Fig 2 corresponds to the $\beta = 0.05$ column.

Figure 11: $\mathrm{Time}(p)$: The duration of the epidemic, in years, as a function of the uptake percentage $p\%$ in four different uptake scenarios with $q = 0.6$, on four different networks. Dotted: $\mathrm{Time}_{\text{rand}}(p)$, dashed-dotted $\mathrm{Time}_{\text{basic}}(p)$, dashed $\mathrm{Time}_{\text{ring}}(p)$, continuous $\mathrm{Time}_{\text{deg}}(p)$. The x-axes vary the uptake percentage p from 0 to 100%, the y-axes shows $\mathrm{Time}(p)$. Note that this has a lot of variance Different columns indicate varying infection probabilities β. Each row represents a different network. First two columns: GIRGs with many super-spreaders ($\tau = 2.3$), third column: GIRG with not many super-spreaders but many long-range connections ($\tau = 3.3$, $\alpha = 1.3$), fourth column: GIRG with not many super-spreaders and also mostly local connections($\tau = 3.3$, $\alpha = 2.3$). The Time data in Fig 2 corresponds to the $\beta = 0.05$ column.

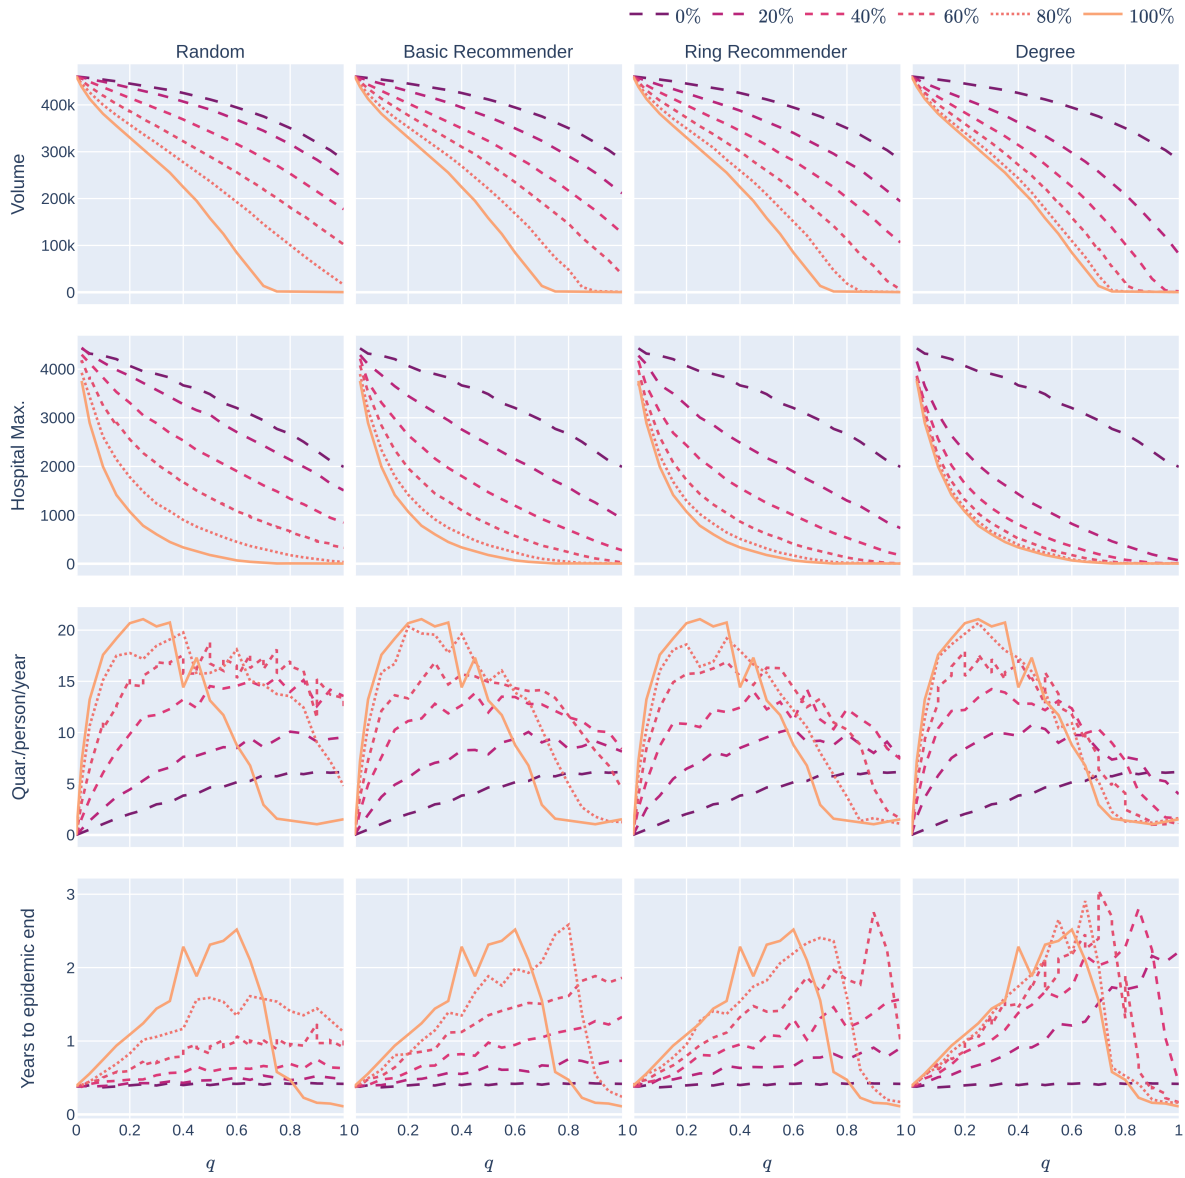

Figure 12: Experiment 2 for a less geometric network: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a spatial scale free network with $\tau = 3.3, \alpha = 1.3$.

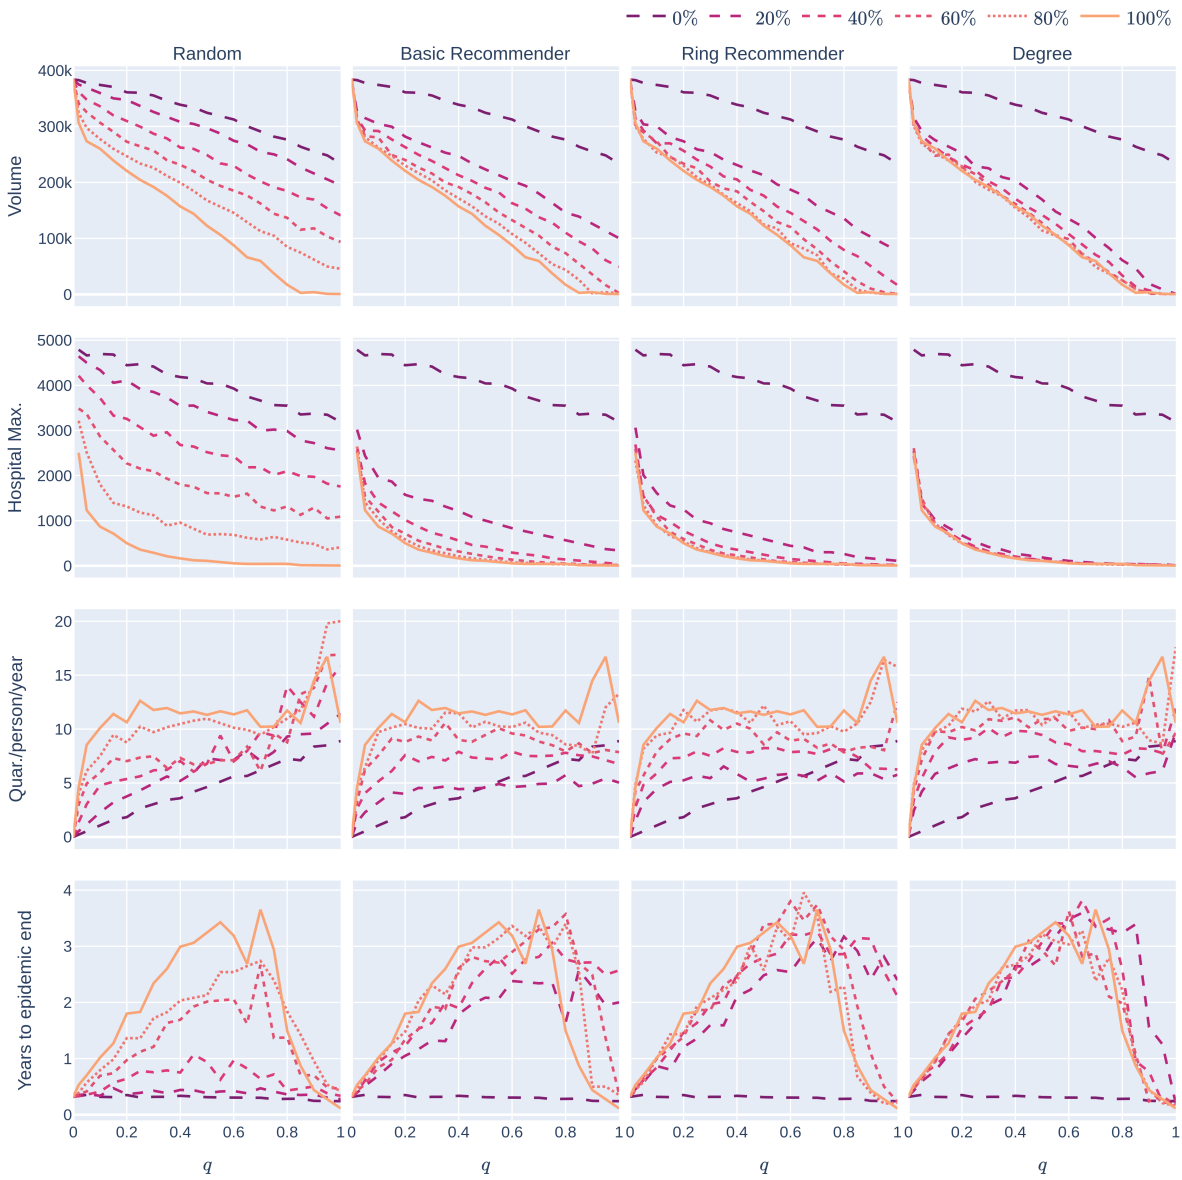

Figure 13: Experiment 2 on a network with many super-spreaders: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a GIRG with $\tau = 2.3, \alpha = 1.3$. In this case, the parameter α does not have a significant effect since long-range edges are present because of the super-spreaders: the figures are very similar to Fig 14 below.

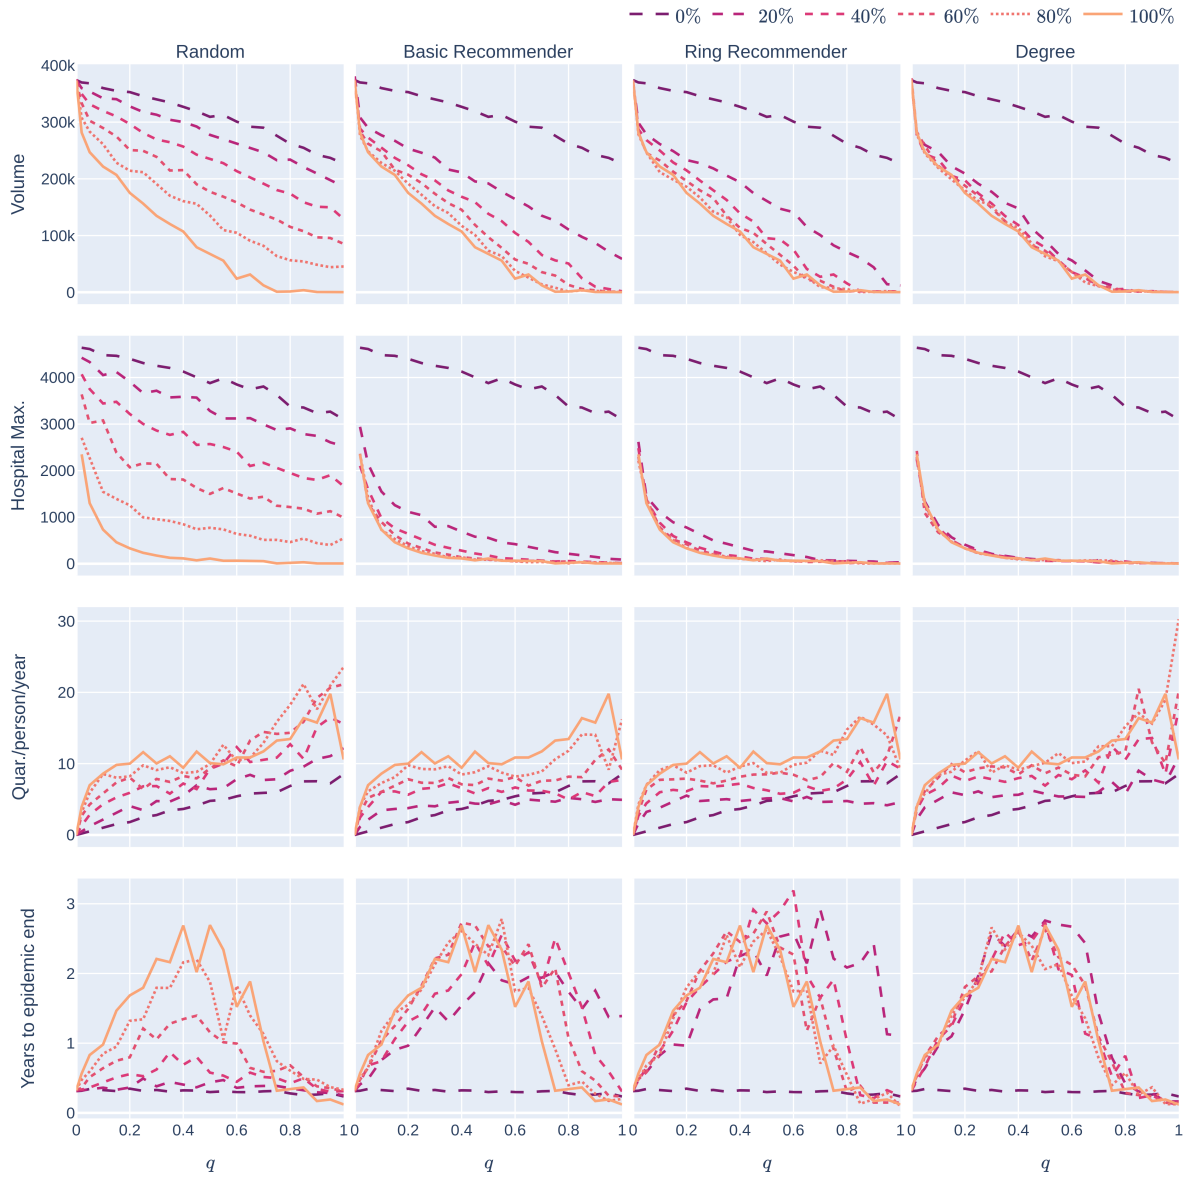

Figure 14: Experiment 2 on a network with many super-spreaders: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a GIRG with $\tau = 2.3, \alpha = 2.3$. Observe the similarity with Fig 13.

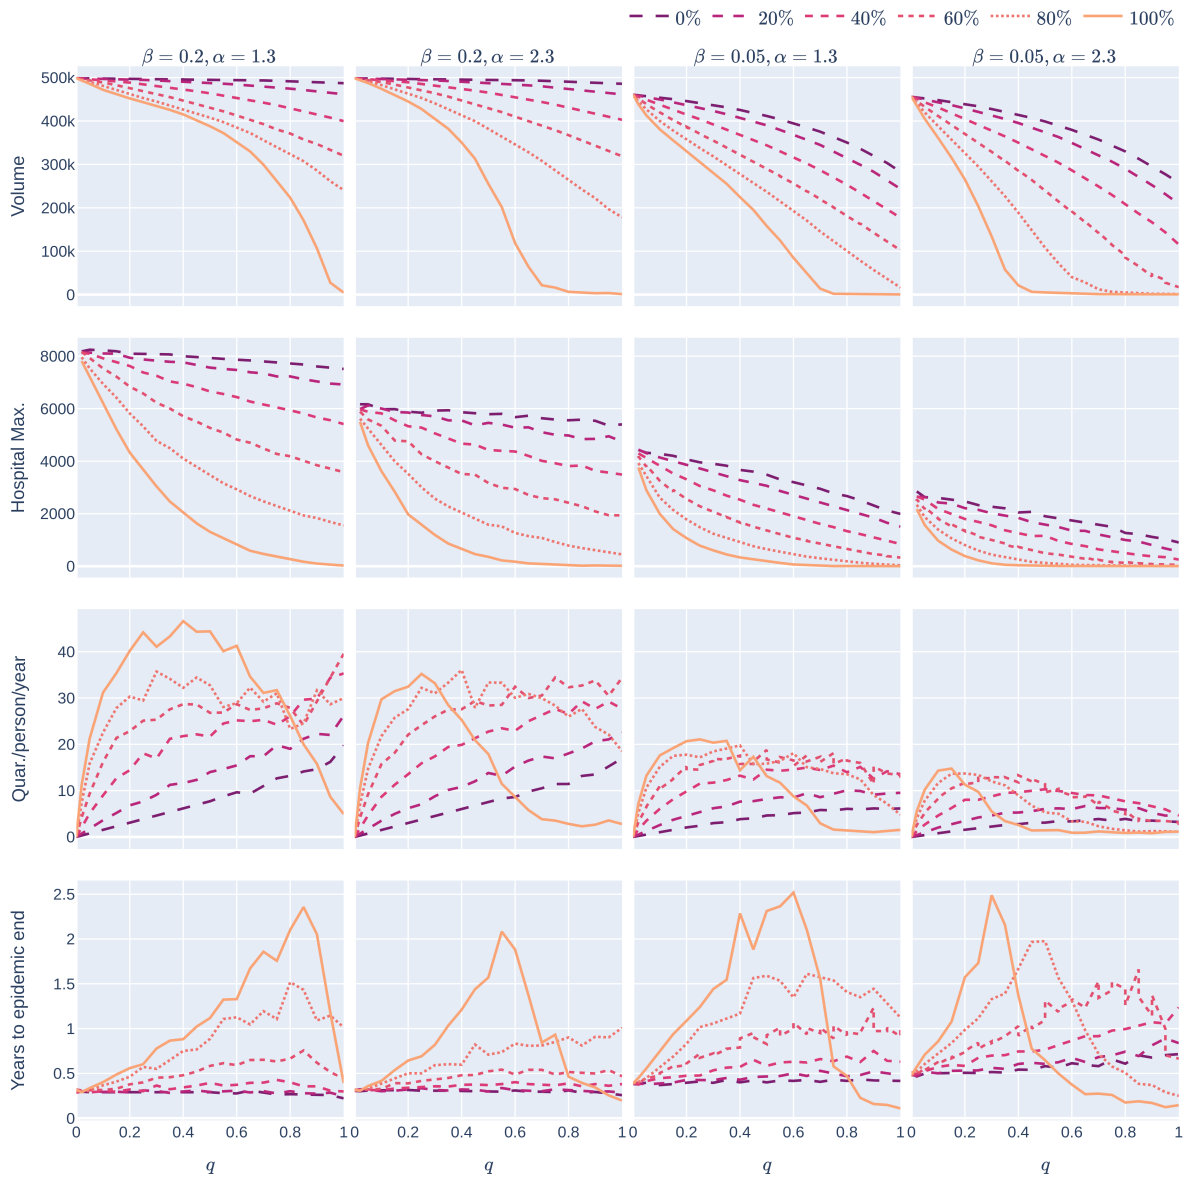

Figure 15: The effect of geometry and level of infectiousness on Experiment 2: The different curves correspond to different CTA uptake-rates ($p\%$), all in the random-uptake scenario in a network with few super-spreaders ($\tau = 3.3$). The horizontal axis corresponds to q, the quarantine strictness. The four columns represent different scenarios – the first column has $\beta = 0.2$ (a high infection rate, perhaps caused by less social distancing) and $\alpha = 1.3$ (a less geometric network). The second column also has $\beta = 0.2$, but with more geometry ($\alpha = 2.3$). The final two columns represent scenarios with a lower underlying infection rate ($\beta = 0.05$), with varying levels of geometry. The four rows correspond to the four KPIs: The first row represents the total volume of the epidemic. The second column represents the maximum hospital load. The third row represents the average number of days per year that individuals spend in quarantine, and the final row represents the length of the epidemic.

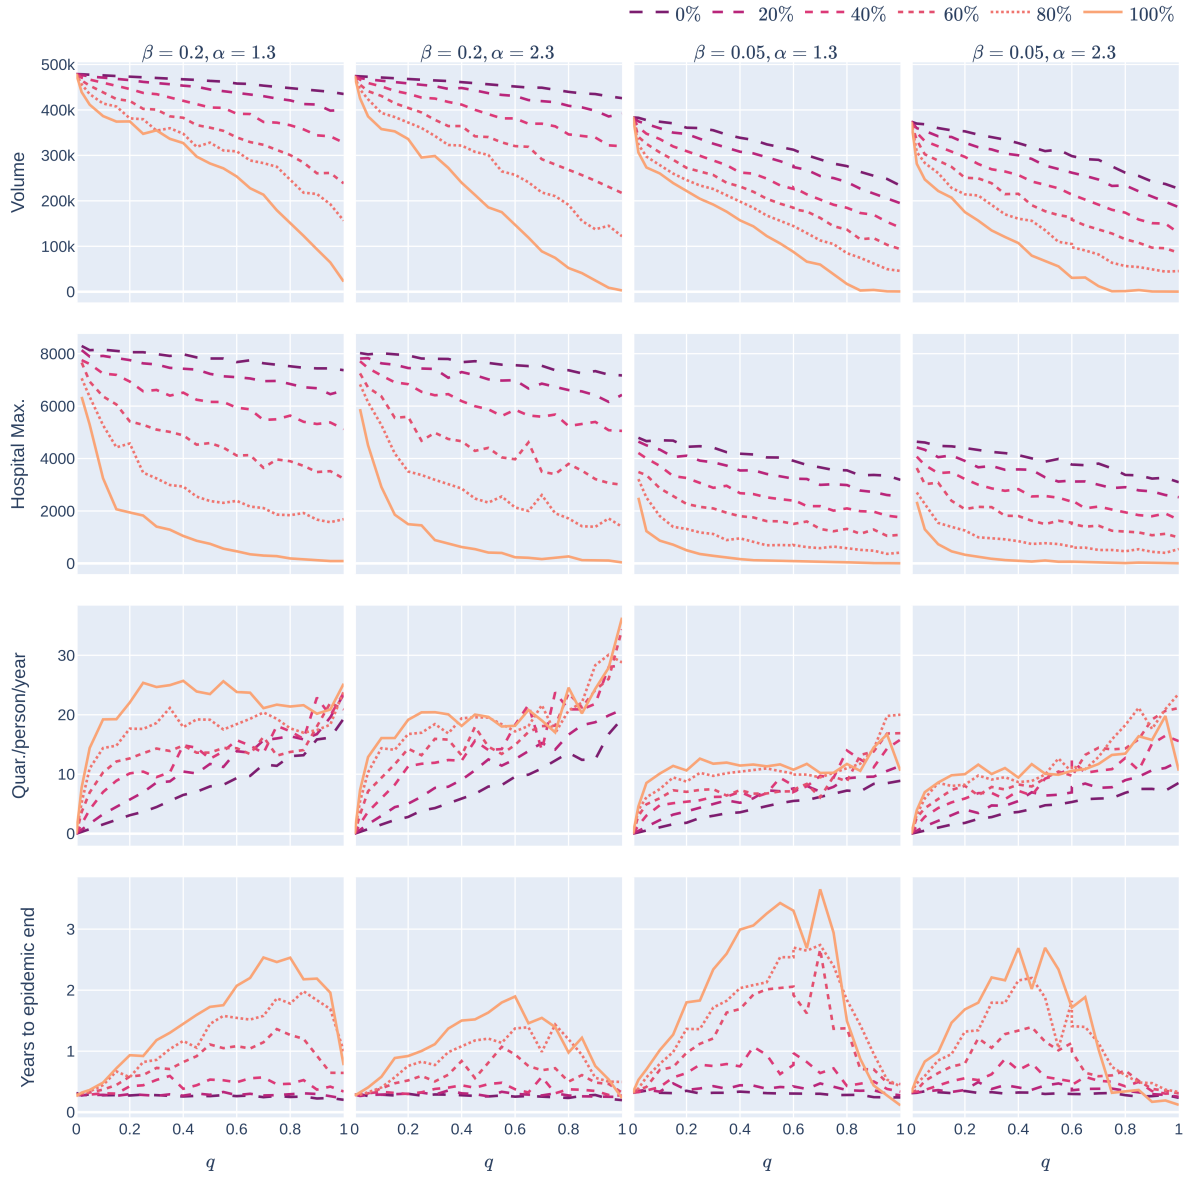

Figure 16: The effect of geometry and level of infectiousness on Experiment 2: The different curves correspond to different CTA uptake-rates ($p\%$), all in the random-uptake scenario in a network with many super-spreaders ($\tau = 2.3$). The horizontal axis corresponds to q, the quarantine strictness. The four columns represent different scenarios – the first column has $\beta = 0.2$ (a high infection rate, perhaps caused by less social distancing) and $\alpha = 1.3$. The second column also has $\beta = 0.2$, but with a higher α, implying fewer (but not significantly fewer) long-range edges. Observe that the effect of increasing α is not very significant, corresponding to the explanation in Section 1.4.4. The final two columns represent scenarios with a lower underlying infection rate ($\beta = 0.05$), with varying levels of geometry. The four rows correspond to the four KPIs: The first row represents the total volume of the epidemic. The second column represents the maximum hospital load. The third row represents the average number of days per year that individuals spend in quarantine, and the final row represents the length of the epidemic.

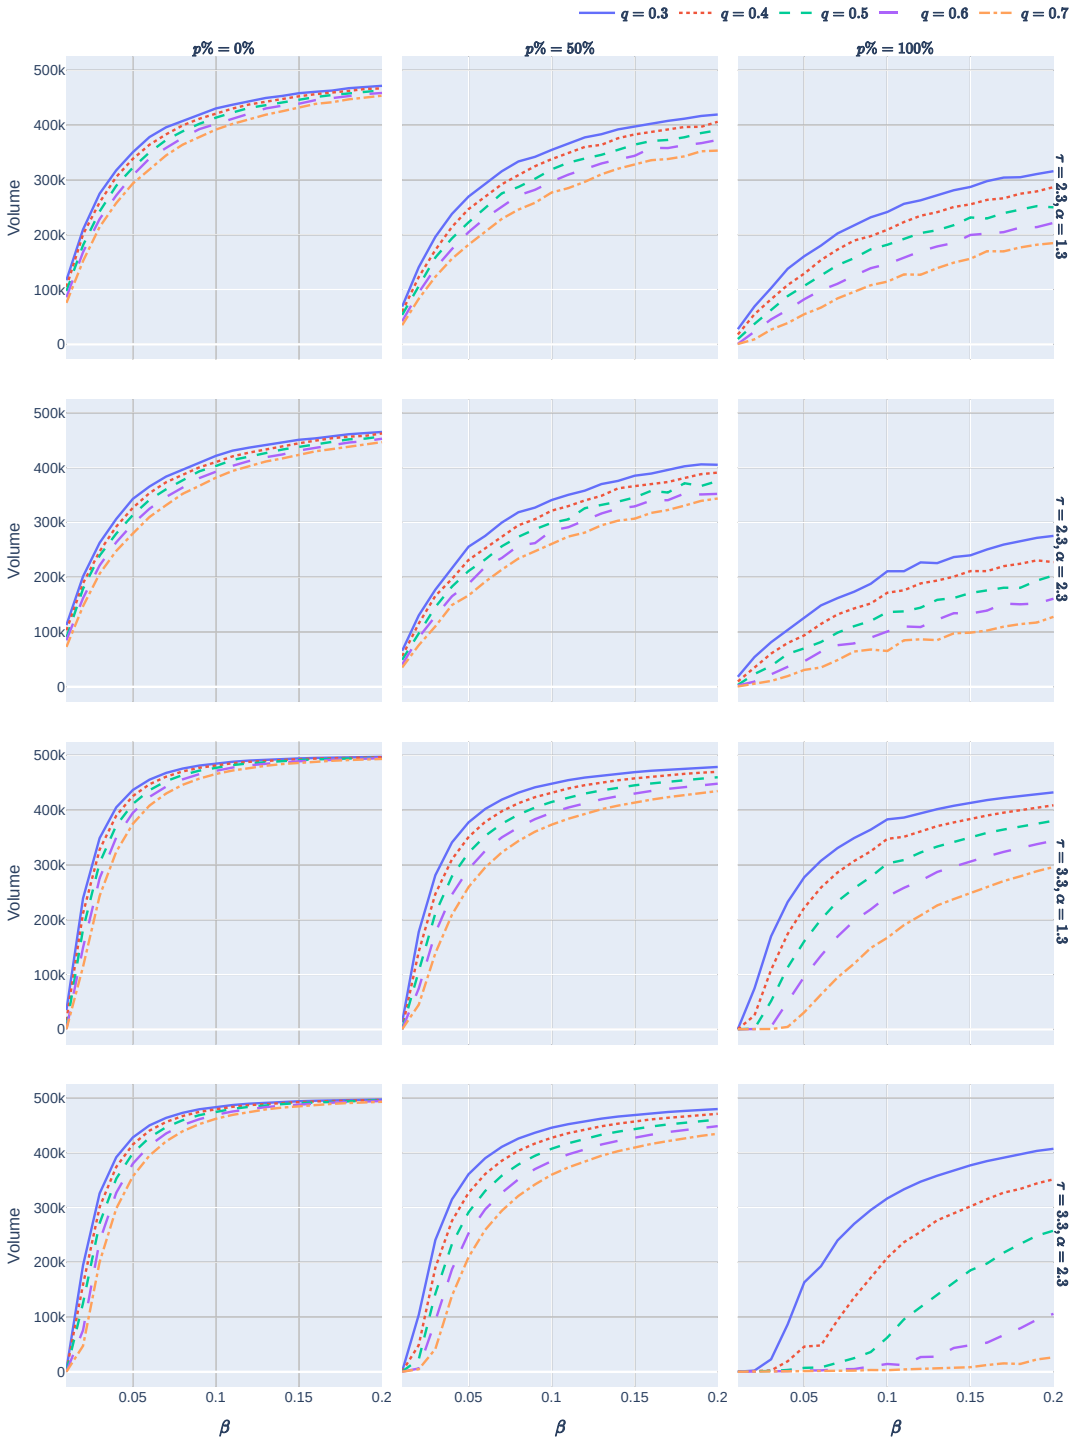

Figure 17: Social distancing: Social distancing measures decrease β, which is shown here on the horizontal axis. The vertical axis is the total volume of the epidemic. There are curves for various CTA-uptake percentages $p\%$ and quarantine strictness values q under the random uptake scenario.

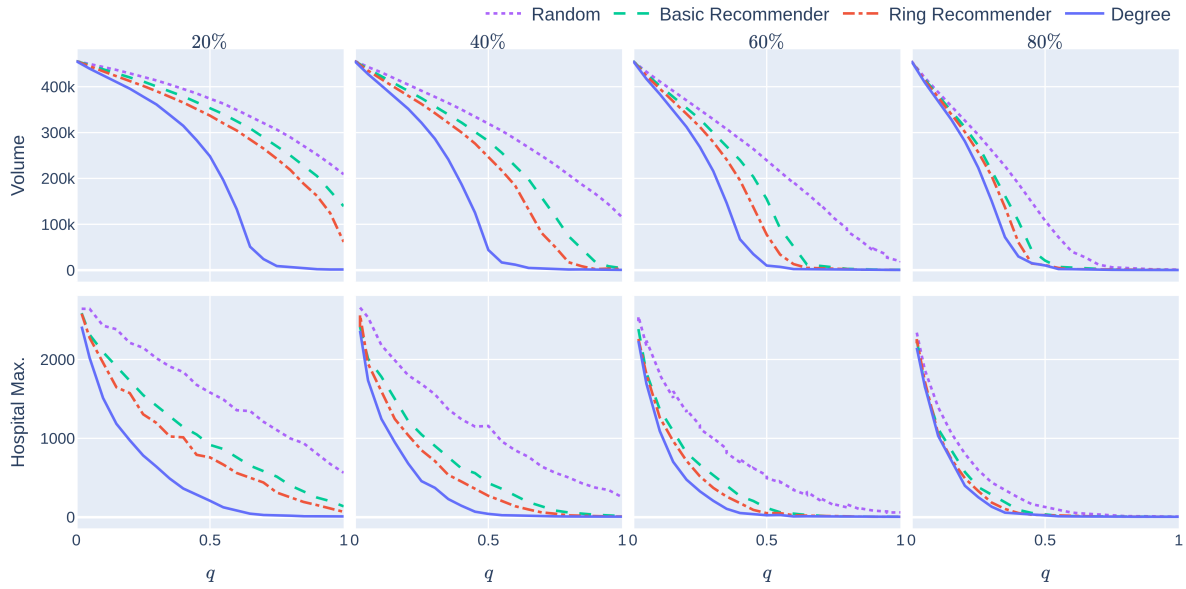

Figure 18: Experiment 2 on a geometric network with not too many super-spreaders: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a GIRG with $\tau = 3.3$ (fewer super-spreaders) and $\alpha = 2.3$ (more geometric).

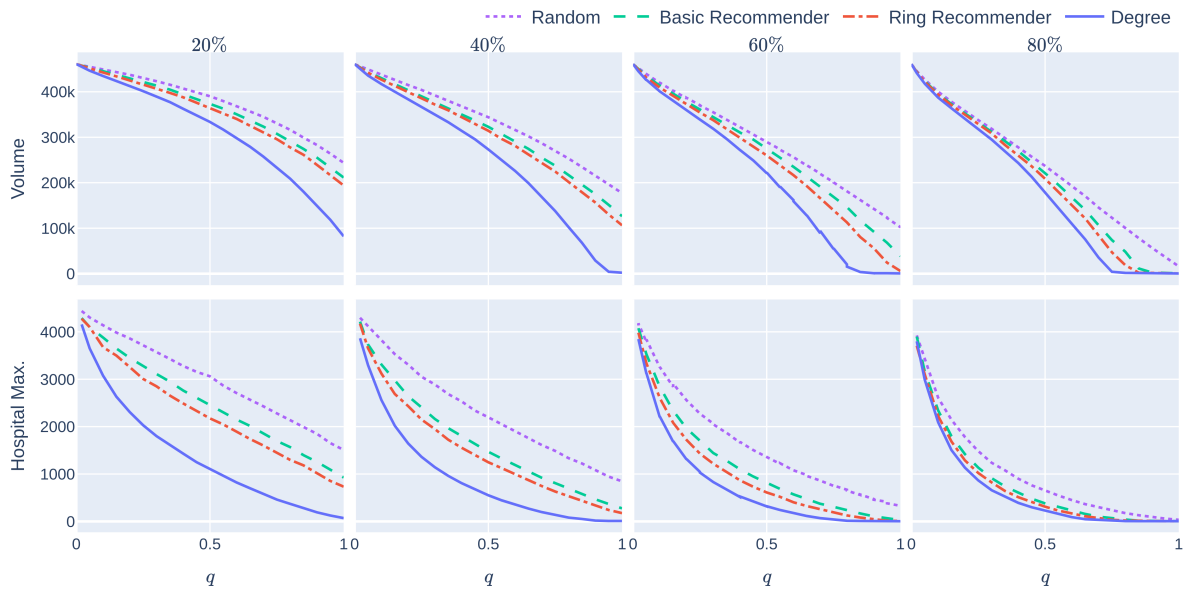

Figure 19: Experiment 2 on a less geometric network with not too many super-spreaders: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a GIRG with $\tau = 3.3$ (fewer super-spreaders) and $\alpha = 1.3$ (less geometric).

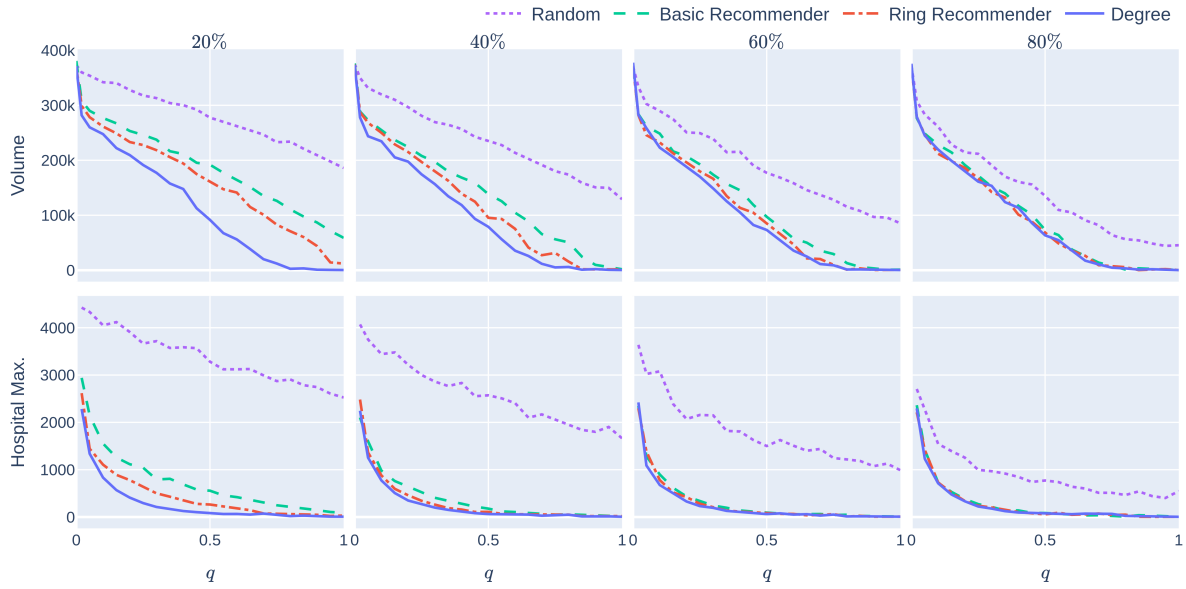

Figure 20: Experiment 2 on a geometric network with many super-spreaders: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a GIRG with $\tau = 2.3$ (more super-spreaders) and $\alpha = 2.3$ (more geometric).

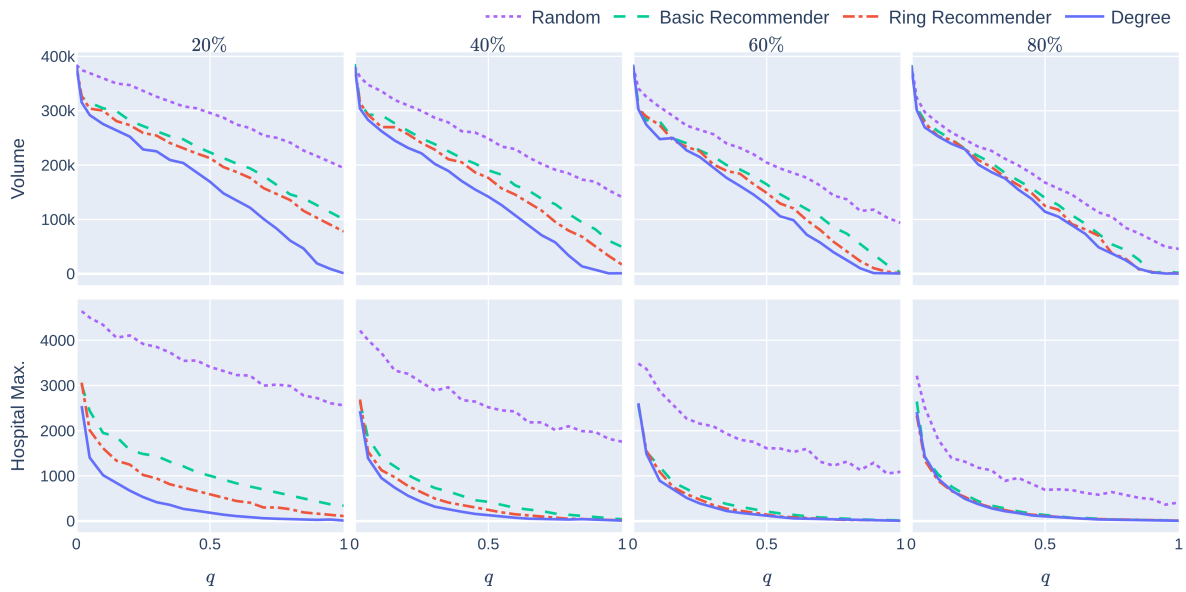

Figure 21: Experiment 2 on a less geometric network with many super-spreaders: Influence of the quarantine strictness q for various uptake percentages $p\%$ and the four uptake scenarios at infection rate $\beta = 0.05$. The underlying network is a GIRG with $\tau = 2.3$ (more super-spreaders) and $\alpha = 1.3$ (less geometric).

CRediT Author Contributions

LG contributed to Conceptualization, Investigation, Writing — Original Draft Preparation, Writing — Review & Editing; JJ contributed to Conceptualization, Data Curation, Investigation, Software, Visualisation; Writing — Review & Editing; JK contributed to Conceptualization, Investigation, Writing — Original Draft Preparation, Writing — Review & Editing; JL contributed to Investigation, Software, Validation.

Acknowledgements

The work of JJ and JK is partly supported by the Netherlands Organisation for Scientific Research (NWO) through grant NWO 613.009.122.

References

[1] Kretzschmar ME, Rozhnova G, Bootsma MC, van Boven M, van de Wijgert JH, Bonten MJ. Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. The Lancet Public Health. 2020;5(8):e452–e459.

[2] Abueg M, Hinch R, Wu N, Liu L, Probert WJM, Wu A, et al. Modeling the combined effect of digital exposure notification and non-pharmaceutical interventions on the COVID-19 epidemic in Washington state. medRxiv. 2020;doi:10.1101/2020.08.29.20184135.

[3] Lambert A. A mathematical assessment of the efficiency of quarantining and contact tracing in curbing the COVID-19 epidemic. medRxiv. 2020;doi:10.1101/2020.05.04.20091009.

[4] Bringmann K, Keusch R, Lengler J. Sampling Geometric Inhomogeneous Random Graphs in Linear Time. In: 25th Annual European Symposium on Algorithms (ESA 2017). vol. 87 of Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany; 2017. p. 20:1–20:15.

[5] Davis S, Abbasi B, Shah S, Telfer S, Begon M. Spatial analyses of wildlife contact networks. Journal of The Royal Society Interface. 2015;12(102):20141004. doi:10.1098/rsif.2014.1004.

[6] González MC, Hidalgo CA, Barabási AL. Understanding individual human mobility patterns. Nature. 2008;453:779 – 782.

[7] Viboud C, Bjørnstad ON, Smith DL, Simonsen L, Miller MA, Grenfell BT. Synchrony, waves, and spatial hierarchies in the spread of influenza. science. 2006;312(5772):447–451.

[8] Strassburg MA. The global eradication of smallpox. American journal of infection control. 1982;10(2):53–59.

[9] Henao-Restrepo AM, Longini IM, Egger M, Dean NE, Edmunds WJ, Camacho A, et al. Efficacy and effectiveness of an rVSV-vectored vaccine expressing Ebola surface glycoprotein: interim results from the Guinea ring vaccination cluster-randomised trial. The Lancet. 2015;386(9996):857–866.

[10] Holme P, Litvak N. Cost-efficient vaccination protocols for network epidemiology. PLOS Computational Biology. 2017;doi:https://doi.org/10.1371/journal.pcbi.1005696.

[11] Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, et al. Social contacts and mixing patterns relevant to the spread of infectious diseases. PLoS Med. 2008;5(3):e74.

[12] Tange O. GNU Parallel - The Command-Line Power Tool. The USENIX Magazine. 2011;36(1):42–47. doi:http://dx.doi.org/10.5281/zenodo.16303.

[13] Bläsius T, Friedrich T, Katzmann M, Meyer U, Penschuck M, Weyand C. Efficiently Generating Geometric Inhomogeneous and Hyperbolic Random Graphs. In: Proceedings of the 27th Annual European Symposium on Algorithms (ESA 2019). vol. 144 of Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany; 2019. p. 21:1–21:14.

[14] Feld SL. Why your friends have more friends than you do. American Journal of Sociology. 1991;96(6):1464–1477.

[15] Jorritsma J, Hulshof T, Komjáthy J. Not all interventions are equal for the height of the second peak. Chaos, Solitons & Fractals. 2020;139:109965.

[16] Clark A, Jit M, Warren-Gash C, Guthrie B, Wang HH, Mercer SW, et al. Global, regional, and national estimates of the population at increased risk of severe COVID-19 due to underlying health conditions in 2020: a modelling study. The Lancet Global Health. 2020;8(8):e1003–e1017.

[17] Menni C, Valdes AM, Freidin MB, Sudre CH, Nguyen LH, Drew DA, et al. Real-time tracking of self-reported symptoms to predict potential COVID-19. Nature medicine. 2020; p. 1–4.

[18] Pan L, Mu M, Yang P, Sun Y, Wang R, Yan J, et al. Clinical characteristics of COVID-19 patients with digestive symptoms in Hubei, China: a descriptive, cross-sectional, multicenter study. The American journal of gastroenterology. 2020;115.

[19] Wang HY, Li XL, Yan ZR, Sun XP, Han J, Zhang BW. Potential neurological symptoms of COVID-19. Therapeutic Advances in Neurological Disorders. 2020;13:1756286420917830.

[20] Poletti P, Tirani M, Cereda D, Trentini F, Guzzetta G, Sabatino G, et al. Probability of symptoms and critical disease after SARS-CoV-2 infection. arXiv preprint arXiv:200608471. 2020;.

[21] Mizumoto K, Kagaya K, Zarebski A, Chowell G. Estimating the asymptomatic ratio of 2019 novel coronavirus onboard the princess cruises ship, 2020. MedRxiv. 2020;.

[22] Nishiura H, Kobayashi T, Miyama T, Suzuki A, Jung Sm, Hayashi K, et al. Estimation of the asymptomatic ratio of novel coronavirus infections (COVID-19). International journal of infectious diseases. 2020;94:154.

[23] He J, Guo Y, Mao R, Zhang J. Proportion of asymptomatic coronavirus disease 2019: A systematic review and meta-analysis. Journal of medical virology. 2020;.

[24] Han H. Estimate the incubation period of coronavirus 2019 (COVID-19). medRxiv. 2020;.

[25] Kretzschmar M, Rozhnova G, van Boven M. Isolation and contact tracing can tip the scale to containment of COVID-19 in populations with social distancing. Available at SSRN 3562458. 2020;.

[26] Lauer SA, Grantz KH, Bi Q, Jones FK, Zheng Q, Meredith HR, et al. The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application. Annals of internal medicine. 2020;172(9):577–582.

[27] Bode M, Fountoulakis N, Müller T. On the Largest Component of a Hyperbolic Model of Complex Networks. The Electronic Journal of Combinatorics. 2015;22:1–43.

[28] Papadopoulos F, Krioukov D, Boguñá M, Vahdat A. Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces. In: Proceedings of the International Conference on Computer Communications (INFOCOM 2010). IEEE; 2010. p. 1–9.

[29] Gugelmann L, Panagiotou K, Peter U. Random hyperbolic graphs: degree sequence and clustering. In: 39th International Colloquium on Automata, Languages, and Programming (ICALP); 2012. p. 573–585.

[30] Deijfen M, van der Hofstad R, Hooghiemstra G; Institut Henri Poincaré. Scale-free percolation. Annales de l'Institut Henri Poincaré, Probabilités et Statistiques. 2013;49(3):817–838.

[31] Bringmann K, Keusch R, Lengler J. Average distance in a general class of scale-free networks with underlying geometry; 2016.

[32] Bringmann K, Keusch R, Lengler J. Geometric inhomogeneous random graphs. Theoretical Computer Science. 2019;760:35–54.

[33] Stegehuis C, van der Hofstad R, van Leeuwaarden JSH. Scale-free network clustering in hyperbolic and other random graphs. Journal of Physics A: Mathematical and Theoretical. 2019;52(29):295101. doi:10.1088/1751-8121/ab2269.

[34] Albert R, Barabási AL. Statistical mechanics of complex networks. Reviews of Modern Physics. 2002;74(1):47.

[35] Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU. Complex networks: Structure and dynamics. Physics Reports. 2006;424(4-5):175–308.

[36] Muchnik L, Pei S, Parra LC, Reis SDS, Andrade JSJ, Havlin S, et al. Origins of power-law degree distribution in the heterogeneity of human activity in social networks. Scientific Reports. 2013;3(1783).

[37] Newman M, Barabasi AL, Watts DJ. The structure and dynamics of networks. vol. 19. Princeton University Press; 2011.

[38] Newman M. Networks. Oxford university press; 2018.

[39] Illenberger J, Nagel K, Flötteröd G. The Role of Spatial Interaction in Social Networks. Networks and Spatial Economics. 2013; p. 255–282. doi:10.1007/s11067-012-9180-4.

[40] Dorogovtsev SN, Mendes JFF. Evolution of networks. Advances in Physics. 2002;51(4):1079–1187. doi:10.1080/00018730110112519.

[41] Biskup M. On the scaling of the chemical distance in long-range percolation models. The Annals of Probability. 2004;32(4):2938–2977. doi:10.1214/009117904000000577.

[42] Deprez P, Hazra RS, Wüthrich MV. Inhomogeneous Long-Range Percolation for Real-Life Network Modeling. Risks. 2015;3(1):1–23.

[43] Gruhl D, Guha R, Liben-Nowell D, Tomkins A. Information diffusion through blogspace. In: Proceedings of the 13th international conference on World Wide Web. ACM; 2004. p. 491–501.

[44] Pastor-Satorras R, Castellano C, Van Mieghem P, Vespignani A. Epidemic processes in complex networks. Reviews of modern physics. 2015;87(3):925.

[45] Pastor-Satorras R, Vespignani A. Evolution and structure of the Internet: A statistical physics approach. Cambridge University Press; 2007.

[46] Dorogovtsev SN, Goltsev AV, Mendes JF. Critical phenomena in complex networks. Reviews of Modern Physics. 2008;80(4):1275.

[47] Karsai M, Kivelä M, Pan RK, Kaski K, Kertész J, Barabási AL, et al. Small but slow world: How network topology and burstiness slow down spreading. Physical Review E. 2011;83(2):025102.

[48] Merler S, Ajelli M, Fumanelli L, Gomes MF, y Piontti AP, Rossi L, et al. Spatiotemporal spread of the 2014 outbreak of Ebola virus disease in Liberia and the effectiveness of non-pharmaceutical interventions: a computational modelling analysis. The Lancet Infectious Diseases. 2015;15(2):204–211.

[49] Bajardi P, Barrat A, Natale F, Savini L, Colizza V. Dynamical patterns of cattle trade movements. Public Library of Science PLoS One. 2011;6(5):e19869.

[50] Janssen J, Mehrabian A. Rumors Spread Slowly in a Small-World Spatial Network. SIAM Journal on Discrete Mathematics. 2017;31(4):2414–2428.

[51] Isham V, Kaczmarska J, Nekovee M. Spread of information and infection on finite random networks. Physical Review E. 2011;83(4):046128.

[52] Bender EA, Canfield ER. The asymptotic number of labeled graphs with given degree sequences. J Combinatorial Theory Ser A. 1978;24(3):296–307. doi:10.1016/0097-3165(78)90059-6.

[53] Bollobás B. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. European J Combin. 1980;1(4):311–316. doi:10.1016/S0195-6698(80)80030-8.

[54] Wormald NC. The asymptotic connectivity of labelled regular graphs. J Combin Theory Ser B. 1981;31(2):156–167. doi:10.1016/S0095-8956(81)80021-4.

[55] Molloy M, Reed B. A critical point for random graphs with a given degree sequence. In: The Structure and Dynamics of Networks. Princeton University Press; 2011. p. 240–258.

[56] Molloy M, Reed B. The Size of the Giant Component of a Random Graph with a Given Degree Sequence. Combinatorics, Probability and Computing. 1998;7(3):295–305.

[57] Hofstad Rvd. Random graphs and complex networks. Vol. 1. Cambridge Series in Statistical and Probabilistic Mathematics, [43]. Cambridge University Press, Cambridge; 2017. Available from: https://doi.org/10.1017/9781316779422.

[58] van Der Hofstad R, Hooghiemstra G. Universality for distances in power-law random graphs. Journal of Mathematical Physics. 2008;49(12).

[59] Hofstad Rvd, Hooghiemstra G, Znamenski D. Distances in random graphs with finite mean and infinite variance degrees. Electronic Journal Of Probability. 2007;12:703–766.

[60] Barbour A, Reinert G. Approximating the epidemic curve. Electronic Journal Of Probability. 2013;18:1–30.

[61] Bhamidi S, van Der Hofstad R, Komjáthy J. The front of the epidemic spread and first passage percolation. Journal Of Applied Probability. 2014;51(A):101–121.

[62] Janson S, Luczak M, Windridge P. Law of large numbers for the SIR epidemic on a random graph with given degrees. Random Structures & Algorithms. 2014;45(4):726–763. doi:10.1002/rsa.20575.

[63] Chung F, Lu L. Connected Components in Random Graphs with Given Expected Degree Sequences. Annals of Combinatorics. 2002;6(2):125–145.

[64] Chung F, Lu L. The average distances in random graphs with given expected degrees. Proceedings of the National Academy of Sciences of the United States of America. 2002;99(25).

[65] Norros I, Reittu H. On a conditionally Poissonian graph process. Advances in Applied Probability. 2006;38(1):59–75.

[66] Barabási A, Albert R. Emergence of scaling in random networks. Science (New York, NY). 1999;286(5439).

[67] Berger N, Borgs C, Chayes JT, Saberi A. On the spread of viruses on the internet. In: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms. Society for Industrial and Applied Mathematics; 2005. p. 301–310.

[68] Dereich S, Mörters P. Random networks with sublinear preferential attachment: degree evolutions. Electron J Probab. 2009;14(43):1222–1267.

[69] Dommers S, Hofstad R, Hooghiemstra G. Diameters in Preferential Attachment Models. Journal of Statistical Physics. 2010;139(1):72–107.