Time evolution of predictability of epidemics on networks

Petter Holme[1,2,3,*] and Taro Takaguchi[4,5]

[1]*Department of Energy Science, Sungkyunkwan University, Suwon 440-746, Korea*
[2]*Department of Physics, Umeå University, 90187 Umeå, Sweden*
[3]*Department of Sociology, Stockholm University, 10961 Stockholm, Sweden*
[4]*National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan*
[5]*JST, ERATO, Kawarabayashi Large Graph Project, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan*

Epidemic outbreaks of new pathogens, or known pathogens in new populations, cause a great deal of fear because they are hard to predict. For theoretical models of disease spreading, on the other hand, quantities characterizing the outbreak converge to deterministic functions of time. Our goal in this paper is to shed some light on this apparent discrepancy. We measure the diversity of (and, thus, the predictability of) outbreak sizes and extinction times as functions of time given different scenarios of the amount of information available. Under the assumption of perfect information—i.e., knowing the state of each individual with respect to the disease—the predictability decreases exponentially, or faster, with time. The decay is slowest for intermediate values of the per-contact transmission probability. With a weaker assumption on the information available, assuming that we know only the fraction of currently infectious, recovered, or susceptible individuals, the predictability also decreases exponentially most of the time. There are, however, some peculiar regions in this scenario where the predictability decreases. In other words, to predict its final size with a given accuracy, we would need increasingly more information about the outbreak.

PACS number(s): 89.65.−s, 64.60.aq, 87.23.Cc

I. INTRODUCTION

Outbreaks of serious infectious diseases can be very frightening, even though many times the diseases die out before reaching the people who are worried about it. Part of the reason for this is that it is hard to forecast disease outbreaks (see, e.g., Refs. [1,2]), and this uncertainty adds to the fear. At the same time, our common models for disease spreading behave, once they take hold in the population [3], like deterministic, perfectly predictable quantities [4]. Such models have two components. First, they model the person-to-person contagion and the history of one individual with respect to the disease. This is done by dividing individuals into states (or "compartments") with respect to the disease, and assigning rules for transitions between the states. A canonical model of diseases that make the infectious person immune upon recovery (or kill their host) is the susceptible-infectious-recovered (SIR) model. This model has three classes: susceptible individuals can acquire the disease, infectious individuals can spread it further, and recovered individuals (who, technically speaking, can also be dead) cannot get the disease nor can they spread it. The transition rules are that a susceptible individual, upon meeting an infectious individual, can (with some probability) become infectious, and after some time, or with some probability, that individual recovers or dies. These transitions are assumed to be instantaneous (which is quite a coarse simplification with respect to real diseases). The second component of epidemic models describes the population level processes over which the pathogens propagate. For decades, theoretical epidemiologists have ignored this issue and taken the "well-mixed assumption," i.e., that any two individuals have the same chance of meeting during an interval of time. Recently, researchers have recognized this problem and

represented the contact structure as a network [5–7]. The regularities, or *structure*, of the network have a great impact on the evolution of disease outbreaks. Among the network structures, the one with the strongest influence is perhaps the degree distribution—the probability distribution of the number of neighbors. This structure explained the existence of superspreaders [8] and challenged the existence of an epidemic threshold below which an outbreak would always die out quickly [9]. Now, if one studies the SIR process on the configuration model [10]—random networks with an arbitrary degree distribution—the outbreak turns out to be completely predictable in the long-time and large-N limits [4]. To be specific, population-averaged quantities such as the fraction of infected individuals converge to deterministic functions of time. The major uncertainty is in the very beginning, whether the outbreak would die out immediately or not. There is thus an apparent contradiction—our canonical compartmental models seem unable to capture the uncertainty of real-world outbreaks.

In this work, we investigate how the predictability of SIR processes evolves with the outbreak itself. We study questions such as the following: Assuming an outbreak started some time t ago, then what can we say about the final number of infected people Ω and the extinction time τ as a function of t? How does our ability to predict Ω and τ depend on the network structure and the type of information we have about the outbreak? In particular, we consider the two scenarios in which the state of each individual is known (see Fig. 1 for a schematic illustration and Sec. III) and only the fractions of currently infectious, recovered, or susceptible individuals are known (see Sec. IV).

II. PRELIMINARIES

In this section, we will clarify the methods and precise model definitions used in the rest of the paper. We also mention some computational considerations.
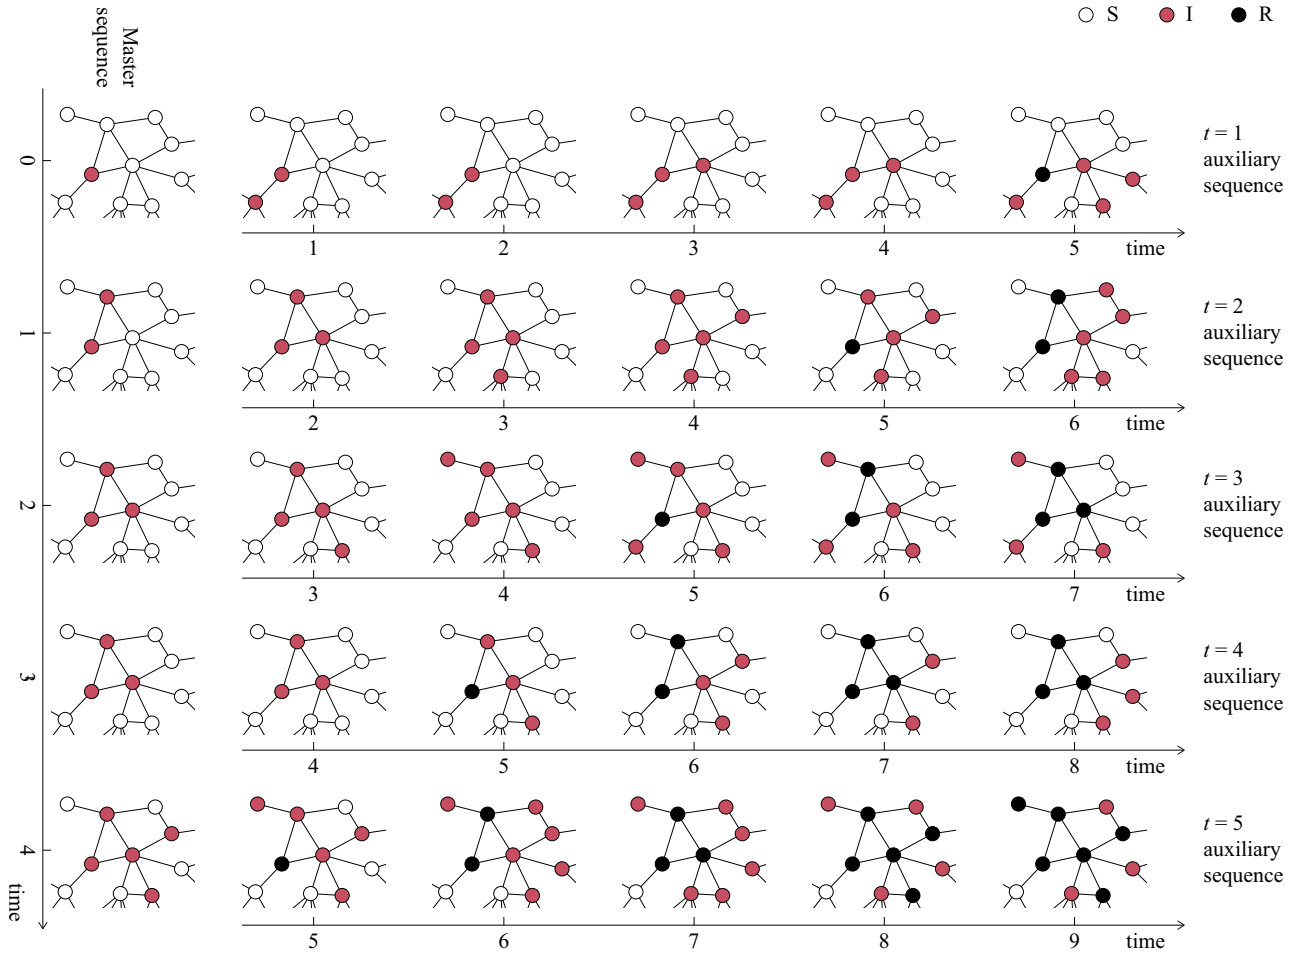
---

*holme@skku.edu

FIG. 1. (Color online) Illustration of the methodology to study the decay of predictability in the case of maximal information of the system. First we fun a master sequence simulation of the SIR process on a network. At every time step, we break this simulation and use the configuration as a seed for 1000 auxiliary simulations. The standard deviation of the final outbreak size and time to extinction thus captures how accurately the outbreak can be predicted given the state of the system at the breaking point.

A. SIR simulation

We assume there is a disease spreading over a static underlying network represented as a graph, $G = (V, E)$. Set V is a set of vertices representing individuals of the population; E is a set of edges representing (unordered) pairs of individuals in close enough contact for the disease to spread. The number of vertices is denoted by N and the number of edges is by M. The vertices are, at a given time, in one of three states—S (meaning that they do not have the disease, but they can get it), I (meaning that they have the disease and they can spread it), and R (meaning that they do not have the disease and they cannot get it). We assume a disease outbreak starts at $t = 0$. At the beginning, all vertices belong to the state S except for a randomly chosen vertex that is in state I. If a pair $(i, j) \in E$ consists of one I and one S vertex, then the S vertex has a chance β of becoming I (otherwise it stays S). We call this an *infection event*. Every vertex with state I has a chance ν per time unit of becoming R (in a *recovery event*). The state of the system at a certain time could thus be fully described by two vectors: $\mathbf{s} = (s_1, \ldots, s_N)$ giving the state of each vertex ($s_i \in \{S,I,R\}$), and a vector $\mathbf{t_I} = (t_1^I, \ldots, t_N^I)$ giving the infection time of the vertices.

Most important quantities describing the outbreak will only depend on the ratio between β and ν (in the well-mixed SIR model, this ratio is called R_0, but to not confuse things, we do not use this name). The actual values of β and ν are only needed to calculate the real time to reach the peak prevalence, extinction, etc. That does not interest us in the present paper, so we measure time in units of $1/\nu$. In a simulation, this can conveniently be done by (at an iteration of the algorithm) performing a random infection event with a probability

$$P = \frac{\beta \Sigma}{\beta \Sigma + I}, \qquad (1)$$

where Σ is the number of edges between an I and an S vertex, and I is the prevalence (note that roman letters symbolize a state and italicized letters represent the number of vertices in the corresponding state) [11]. The factor $\Delta t = (\beta \Sigma + I)^{-1}$ represents the time increment since the last iteration. Thus, to keep track of the time (in units of $1/\nu$), one adds Δt to a variable representing time. If an infection event is not performed (which happens with a probability $1 - P$), we perform a recovery event. In an infection event, the SI edge (to become II) is chosen randomly among all SI edges. Similarly,

TABLE I. Summarizing the network structure of the six classes of networks we study. All networks have 2500 vertices. Except for the large-world network (which is unique), we use the same 1000 independent realizations as in the rest of the work. In the second row, we list the standard deviation (once again with the exception for the large-world network).

Network model	M	s	d	C
Large world	5000	1	312.8	0.5
Small world	5000	1	36.6	0.484
	0	0	3.69	0.00236
Random regular	5000	1	6.47	0.000964
	0	0	0.00401	0.000440
Scale free, $\gamma = 2$	6696	0.963	3.24	0.0666
	365	0.00783	0.125	0.0132
Scale free, $\gamma = 2.5$	3039	0.785	4.54	0.0153
	160	0.0271	0.464	0.00497
Scale free, $\gamma = 3$	2046	0.46	6.88	0.00277
	69.3	0.0459	1.01	0.00161

in a recovery event, the I vertex (to become R) is selected with uniform randomness among all I vertices.

B. Network models

To study the effects of network topology on the outbreak predictability, we use a collection of six network models. We choose the models to reflect a variety of stylized network topologies. This methodology is inspired by Ref. [12]. Throughout this paper, we will use graphs of size $N = 2500$. We use 1000 realizations of each network model. A summary of the basic network structure in these networks can be found in Table I.

1. Large- and small-world networks

First, we use the Watts-Strogatz small-world network model [13]. In this model, all vertices have the same degree k (number of neighbors). They are initially arranged on a ring and connected to their $k/2$ nearest neighbors, at either side, of the ring. Counting modulo N, these neighbors can be enumerated as $i - k/2, i - k/2 + 1, \ldots, i - 1, i + 1, \ldots, i + k/2$. Then one goes through all vertices i around the ring, and for every edge (i,j) pointing to a neighbor ahead (i.e., $j > i$ modulo N), with probability p, one replaces (or "rewires") it by (i,j') [where j' is a different vertex from i and $(i,j') \notin E$]. By construction, if $p = 0$ and $k > 2$, this model gives networks with a high density of triangles. Furthermore, for these parameter values, the average distances grow like N. Watts and Strogatz show that if p is slightly larger than zero, the distances scale much slower with N. The number of triangles is, on the other hand, not so sensitive to p. The reason for this behavior is that the rewired edges (usually called "long-range edges," with reference to the ring topology) connect distant parts of the graph. For the triangles, only the directly involved vertices are affected by a rewiring, but for the distances the extended neighborhoods of these vertices are affected. In this work, we use the Watts-Strogatz model with $p = 0$ (calling it *large-world networks*) and $p = 0.01$ (calling it *small-world networks*). We use $M = 5000$, so the small-world networks

have an expected number of 50 rewired edges. The large-world networks, of course, are all isomorphic and only one copy of them is needed as a substrate for our simulations.

2. Random regular graphs

Random regular graphs [14] are designed to be as uniform as possible with respect to the positions of vertices. Like the Watts-Strogatz model, all vertices have the same degree $k = 4$ (giving $M = 5000$). Other than that, they are as random as possible. In graph models with fewer restrictions (such as Erdös-Rényi random graphs [15], where there is a fixed probability for any vertex pair to belong to E), the degrees vary, which differentiates the vertices. On the other hand, for our purposes, Erdös-Rényi graphs would probably give very similar results.

3. Scale-free networks

So far, all the model networks we discussed have uniform degree distributions. This is not very realistic. Rather, the degrees of empirical networks are in general often broadly and skewedly distributed. This is also true for some of the particular networks that diseases spread over, such as sexual networks [16] or networks of contacts between patients in hospitals [17]. To model such networks, we use the configuration model [15] with an emerging power-law degree distribution. We draw N integers $\{k_i\}_{i=1}^{N}$ from a probability distribution

$$P(k) = \begin{cases} k^{-\gamma} & \text{if} \quad 1 \leqslant k \leqslant N-1, \\ 0 & \text{otherwise}, \end{cases} \tag{2}$$

until $\sum_{i=1}^{N} k_i$ is even. Then we randomly attach edges between vertices until vertex i has k_i neighbors. We do not forbid multiple edges or self-edges. That is, the resulting object is a multigraph. The upper limit $N - 1$ is somewhat arbitrarily chosen to be the same as the maximum degree in a simple graph. This is needed to keep the fluctuations down (that comes from the extreme variation if k is generated by small γ values). Finally, we construct a simple graph by removing multiple edges and self-edges. We try three different γ values from the typical range of empirical networks [18]: 2, 2.5, and 3.

4. Theoretical threshold values

As mentioned earlier, an interesting feature of epidemic models is that they can have a threshold behavior where the average fraction of infected individuals is finite as $N \to \infty$ provided that β is larger than a threshold value β_c. By analogy to models in statistical physics, this could be described as a continuous phase transition [7]. This is not the focus of our study, but we will review a few theoretical results for our models to give some context.

The large-world and scale-free networks have thresholds for the extreme values of β. In the large-world network case, the epidemics spread along a one-dimensional chain (Fig. 2). For any $\beta < 1$ value, there is a finite chance that the outbreak will stop, so there is an expected distance it will propagate from the seed. The outbreaks are limited by this finite distance, so if $N \to \infty$, the fraction of affected vertices is zero. For the scale-free networks of the exponents that we study, the threshold is $\beta_c = 0$ [7,9]. The random regular networks have
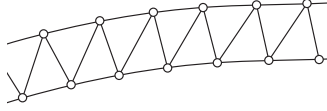
FIG. 2. The local structure of the Watts-Strogatz model networks with $k = 4$.

a threshold $\beta_c = 1/2$ [7]. The small-world networks in the limit $p = 1$ have $\beta_c = 2/7$ [19], but we are not aware of any derivations of the threshold for $0 < p < 1$.

III. PREDICTABILITY GIVEN THE STATE OF THE SYSTEM

The predictability of any kind of phenomenon depends on the information available and the ability to use it. In this paper, we assume that the disease spreading is determined by the SIR process. In addition, we assume that we know the precise state of the system, i.e., the underlying network, the state of all vertices, and when they changed state. In that case, one cannot, by definition, do better in predicting the future than evaluating the SIR process with the state of the system as the input. In this section, we focus on this limit of maximum information about an outbreak.

A. Methods

We want to understand how the predictability of Ω and τ depends on t given that we know $\mathbf{s}$ and $\mathbf{t}^{\mathrm{I}}$. To this end, we run an SIR simulation producing a *master sequence* $(\mathbf{s}_0, \mathbf{t}_0^{\mathrm{I}})$, and at every time step we start 1000 simulation runs with $(\mathbf{s}_0, \mathbf{t}_0^{\mathrm{I}})$ as the initial condition. The standard deviations of Ω and τ of these *auxiliary sequences* measure the unpredictability evolution of the master sequence. Finally, we average these standard deviations over at least 2000 master sequences and call them σ_Ω (for the outbreak size) and σ_τ (for the extinction

time). For each such sequence, we take a random network realization from a pool of 1000. When an outbreak is dead (when $\Sigma = 0$), it contributes with zero to the average and standard deviation. This procedure is illustrated in Fig. 1.

B. Results

Now we will turn to our numerical results relating to the predictability.

1. Predicting the final outbreak size

We start by investigating the σ_Ω for our six classes of networks for an exponentially increasing sequence of 11 β values. We choose the β values so that we cover outbreaks of all sizes, in all network models. We plot the results in Fig. 3. Our first observation is that none of the curves (i.e., for no β value and for no topology) decays slower than exponential. Most of them decay roughly exponentially, while some decay more rapidly. The faster-than-exponential decay is clearest for the highest β curves of the small- and large-world networks [see Figs. 3(a) and 3(b)]. The random-regular graphs in Fig. 3(c) do not have the same fast drop-off. Since all three models in Figs. 3(a)–3(c) have uniform degree distributions, the explanation must be something else. Locally, the large- and small-world networks look the same—bands of stacked triangles (see Fig. 2). We believe the cutoff of the exponential decay relates to the typical length of these bands. Thus, for example, the $\beta = 8$ curve in Fig. 3(b) bends down around $t = 18$. Our interpretation is that the early, slower decay is a period in which none of the outbreaks has reached the entire population. For the other models, since the outbreak can reach the entire population much faster, these two regions get blurred and the result is just one exponential decay. Comparing curves of different β values, we see that the ones with the slowest decay are the ones with the longest extinction times. Presumably, as an extinct outbreak is perfectly predictable, the extinction
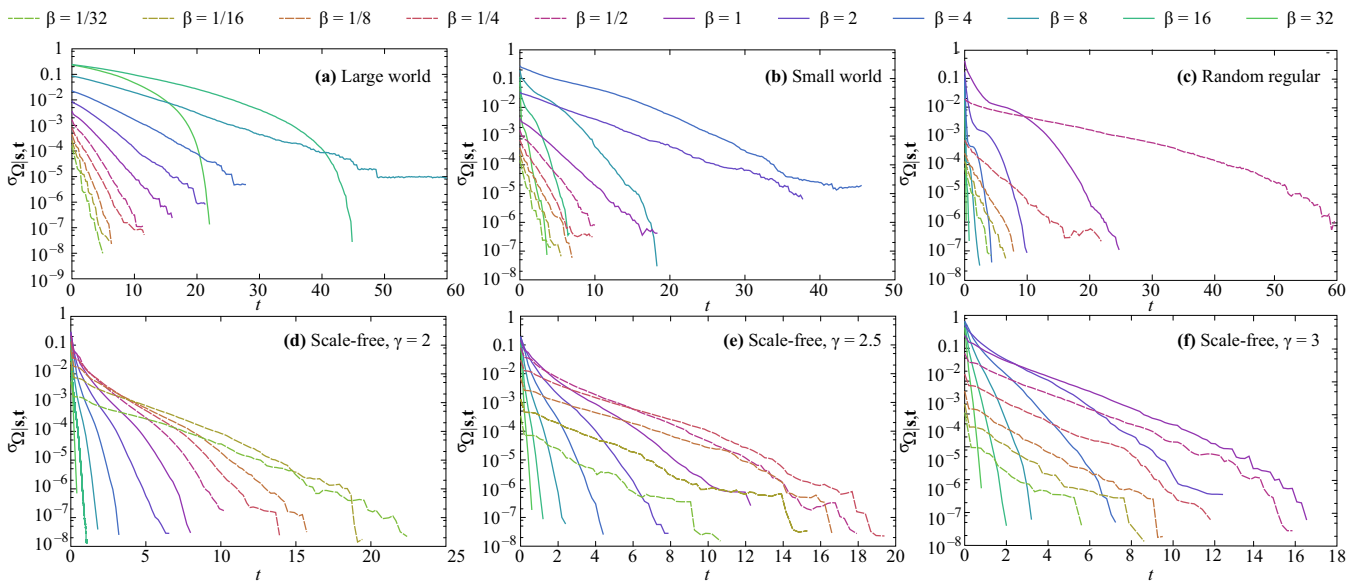

FIG. 3. (Color online) Standard deviation of the final outbreak size given the state at t, as a function of t. As σ_Ω measures the outbreak diversity, its increase reflects the decay of predictability. The curves end when the all the simulated outbreaks have died. In principle, σ_Ω is defined for any $t > 0$, but after the outbreaks are over it is zero and thus not visible in this figure.
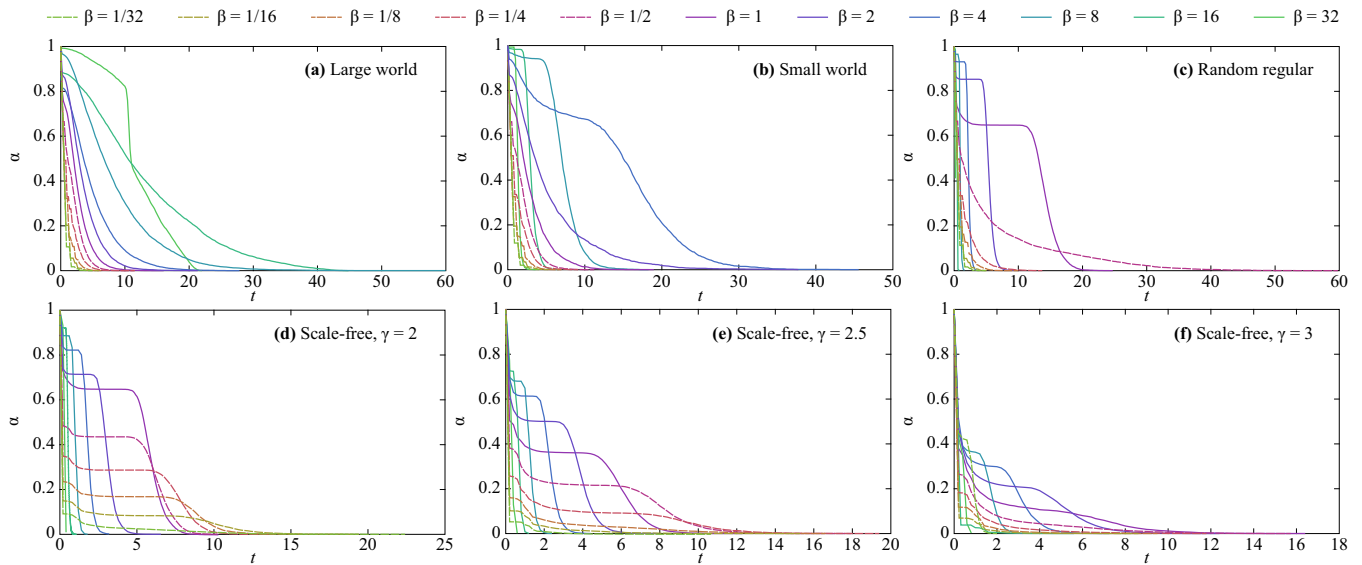
FIG. 4. (Color online) The fraction of surviving outbreaks as a function of t. The panels and parameter values are the same as in Fig. 3.

time is an important factor in determining the time evolution of predictability. As a reference, to see the progression of the outbreak, we plot the fraction of surviving outbreaks as a function of time in Fig. 4 and the average prevalence in Fig. 5. The increasing effect of perfectly predictable dead outbreaks would explain why intermediate β values have the slowest decay of predictability—extinction times will initially grow with β (as the chance for an early extinction decreases) and then decrease (due to the increasingly fast spread and subsequent burn-out in the population [20]). This is, however, not the case for the scale-free networks in the same parameter range. In Figs. 3(d)–3(f), we plot the results for scale-free networks. For $\gamma = 2.5$ and 3, when β is small, the decay rate of σ_Ω is almost constant. The σ_Ω value still increases with β, but this increase happens in the early die-off (mentioned in the Introduction). In other words, smaller β affects σ_Ω by

increasing the chance that the disease will die in the very early stage. Once the outbreak takes hold in the population, the predictability decays independent of β. This can also be seen in Fig. 4 as plateaus after the initial die-off. For larger β, over the peak in extinction time, the lifetime decreases with β. One interesting thing to note from Fig. 5 is how long the outbreaks live after the peak. The prevalence for many of the curves in Fig. 5 has reached very low values before the corresponding curve in the fraction of surviving outbreaks (Fig. 4).

To investigate more closely the relationship between the decay of σ_Ω and β, we assume the scaling form $\sigma_\Omega \sim \exp(-t/T)$ and measure T [21]. Note that this assumption is not justified by any rigorous theory—it should be regarded as a somewhat sketchy summary of Fig. 3. In Fig. 6, we show the results. Some of the curves (as mentioned) do not have an exponential tail and are excluded from this analysis. This
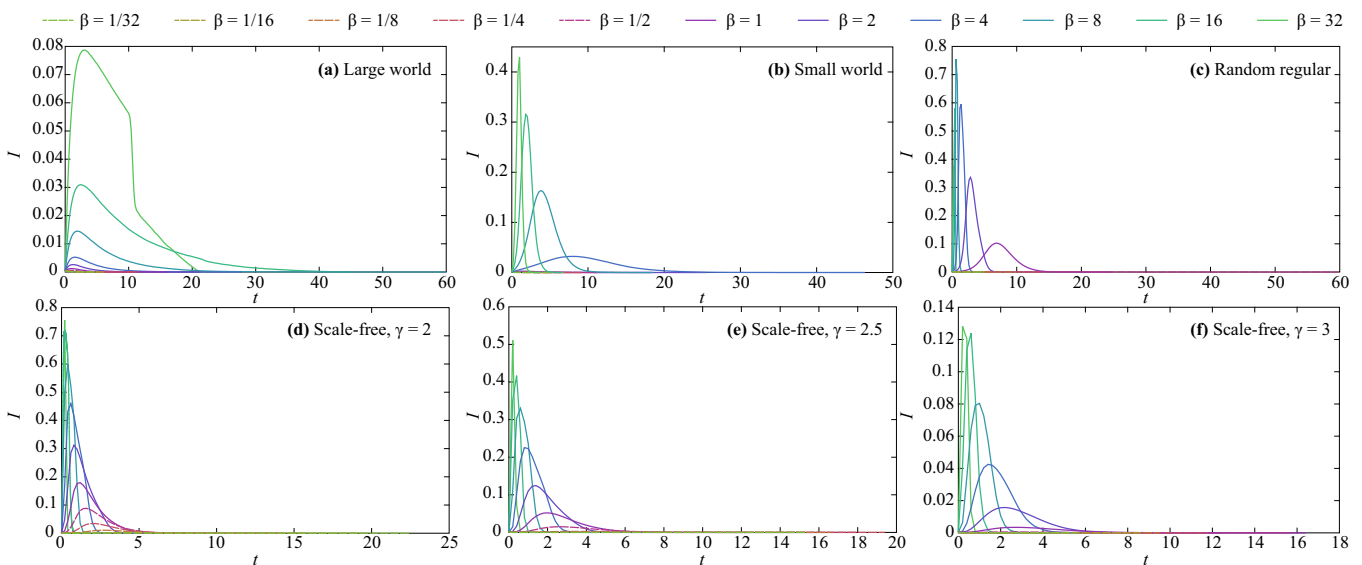

FIG. 5. (Color online) The prevalence (fraction of infectious vertices) as a function of t. The panels and parameter values are the same as in Fig. 3.
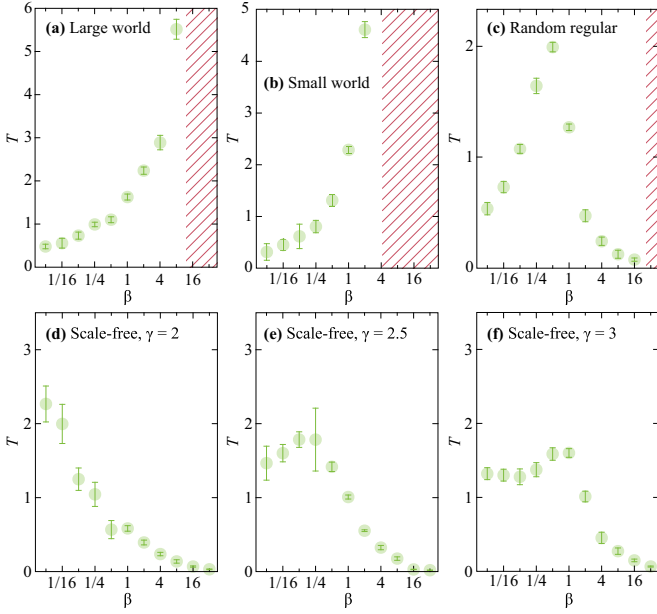
FIG. 6. (Color online) The decay constants from nonlinear least-squares fits to exponential forms of the curves shown in Fig. 3. The shaded regions indicate where σ_Ω does not fit an exponential function. The error bars represent standard errors.

figure illustrates how, except for maybe the $\gamma = 2$ scale-free network, all curves of Fig. 3 have the slowest decay of σ_Ω for intermediate β values. It seems reasonable that the location of this peak converges to the epidemic threshold as $N \to \infty$. On the other hand, none of the Fig. 3 curves has a concave shape indicative of a slower-than exponential decay (which one could expect if σ_Ω was a critical parameter, diverging at the threshold). An alternative hypothesis is that this peak coincides with the maximal extinction time, which is thought

to be larger than, but distinct from, the epidemic threshold [20]. We will leave this as a question for future studies.

2. Predicting the extinction time

In addition to studying the (un)predictability with respect to outbreak size, we also study σ_τ—the corresponding quantity for the extinction time τ—in Fig. 7. The general picture from Fig. 3 holds—the decay is roughly exponential. However, there is more structure in these curves. The decay fits the exponential form worse than σ_Ω does. Several of the σ_τ curves for the random-regular and scale-free graphs plateau at intermediate β values. This means there are times when predicting the outbreak size gets more precise with time, but predicting how long the epidemics will last does not.

IV. PREDICTABILITY GIVEN THE SIZES OF THE COMPARTMENTS

In Sec. III, we studied the scenario in which we have maximum information about an outbreak. This is of course an idealized situation, and in this section we turn to the more realistic scenario in which we know the number of people infected and recovered, but not who is in what state.

To assess the predictability based on the summarized information about the process, we redefine the measure of predictability σ_Ω by the standard deviation of Ω conditioned on the number of I vertices (I), R vertices (R), and I or R vertices. We denote the measures redefined as $\sigma_{\Omega|I}$, $\sigma_{\Omega|R}$, and $\sigma_{\Omega|I+R}$, respectively. In other words, $\sigma_{\Omega|I}$, for example, quantifies the uncertainty of the final outbreak size Ω conditioned on the fact that we know the I value at time t. If $\sigma_{\Omega|I}$ of t takes a small value, the final outcome of the epidemic process is predictable at time t in the sense that the fluctuation in Ω conditioned on I at t is small. Note that, because the total number of vertices is fixed to N, $\sigma_{\Omega|I+R}$ is equivalent to that conditioned on the number of S vertices. We calculate these measures as follows (we take $\sigma_{\Omega|I}$ as an example). Instead of setting a
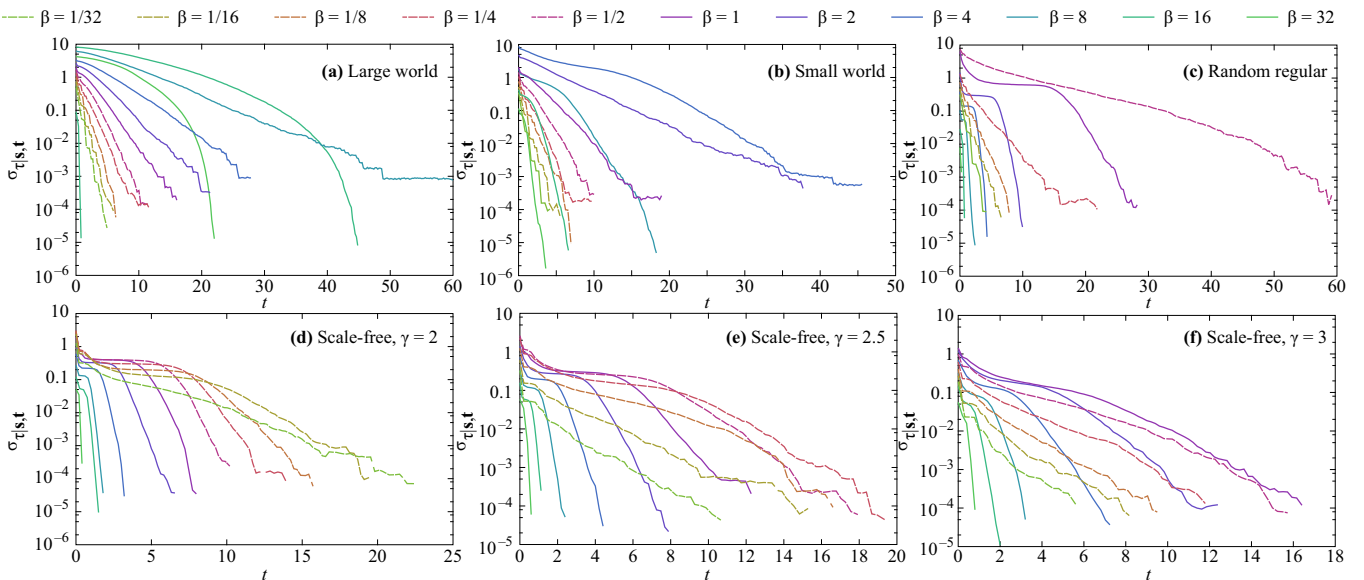

FIG. 7. (Color online) This figure corresponds to Fig. 3, but shows the standard deviation of the extinction time rather than the outbreak size.
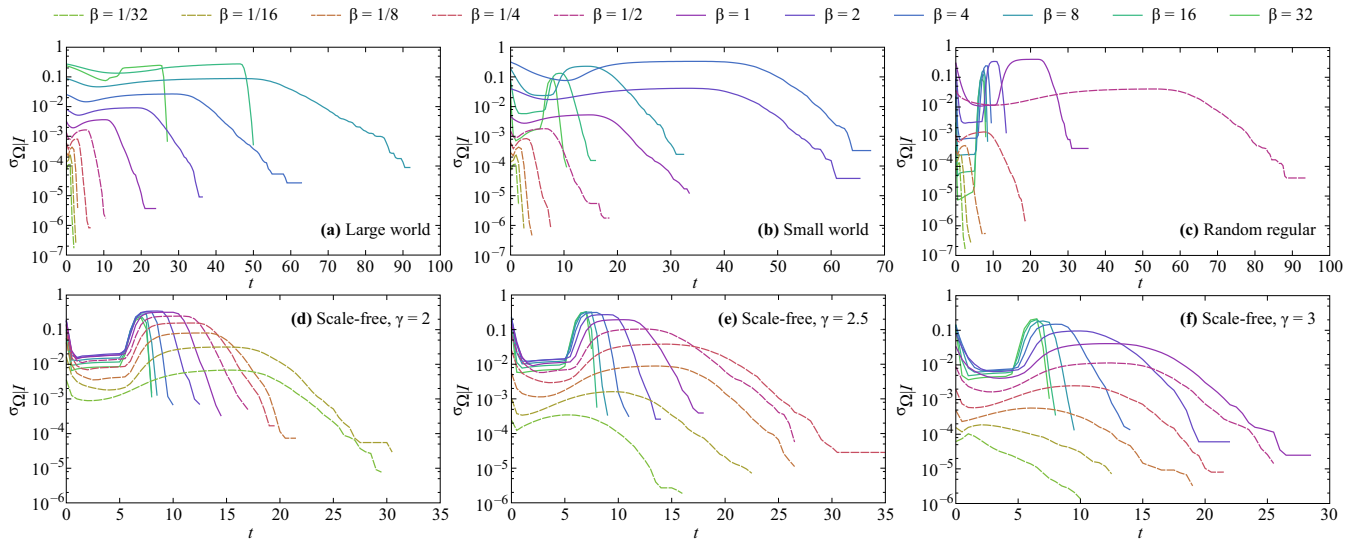
FIG. 8. (Color online) Standard deviation of the final outbreak sizes conditioned on I, that is, the number of infected vertices at t. Otherwise (the parameter values, number of averages, etc.) the plot corresponds to Fig. 3.

master sequence as we did in Sec. III, we first simulate 10 000 outbreaks with the initial seed chosen randomly on the given network realization. We pick out the outbreaks with the same I value at t among the outbreaks and calculate the standard deviation of Ω. We do this calculation for all I values observed at t. Finally, we obtain $\sigma_{\Omega|I}$ of t by calculating the expectation value of the standard deviation of Ω over the distribution of the I value observed at t. For deriving $\sigma_{\Omega|R}$ and $\sigma_{\Omega|I+R}$, we perform the same calculation based on R and $I + R$.

The resultant predictability measures $\sigma_{\Omega|I}$, $\sigma_{\Omega|R}$, and $\sigma_{\Omega|I+R}$ are plotted as a function of time t in Figs. 8, 9, and 10, respectively. Note that we use the same set of network realizations that we used in Sec. III. It also should be noted that the ranges of extinction time shown in Figs. 8, 9, and 10 are larger than those shown in Figs. 3 and 7, simply because we consider a larger number of simulation outbreaks than we

did in Sec. III. To be precise, here we consider 10^7 simulation runs, i.e., 10^4 runs on each of 10^3 network realizations.

For all of the network models, we observe a bump in $\sigma_{\Omega|I}$ (see Fig. 8). When we increase the β value, we find the bump is located at the range of larger t up to a certain β value and is located closer to $t = 0$ when we further increase β. In particular, for scale-free networks with $\gamma = 2$ [Fig. 8(d)], we observe that the bump is located at the largest t when $\beta = 1/32$. The presence of a bump in the $\sigma_{\Omega|I}$ curves can be explained as follows. At the beginning of the process, $\sigma_{\Omega|I}$ decreases with time t because the standard deviation of Ω among the outbreaks with the same I values gets smaller as the outbreaks are separated into ones rapidly dying out and those eventually spreading to a large scale. In the middle term of the process, $\sigma_{\Omega|I}$ increases with time t because a number of outbreaks, regardless of Ω, fall into the range of the small I
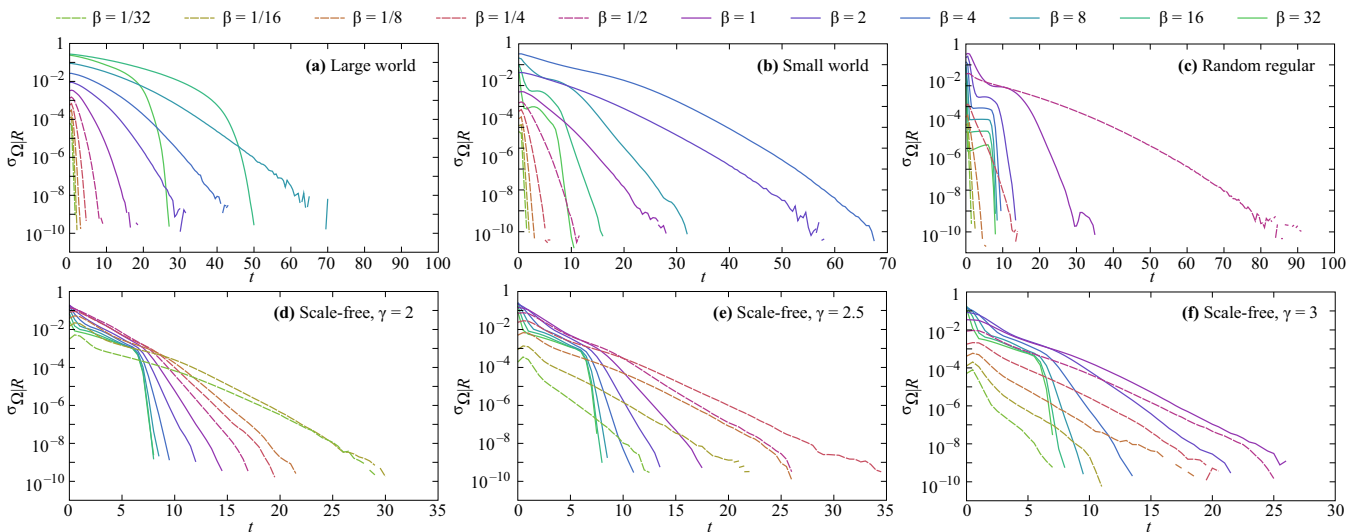

FIG. 9. (Color online) Plot corresponding to Fig. 8 for the standard deviation of the final outbreak sizes conditioned on R, that is, the number of recovered vertices at t.
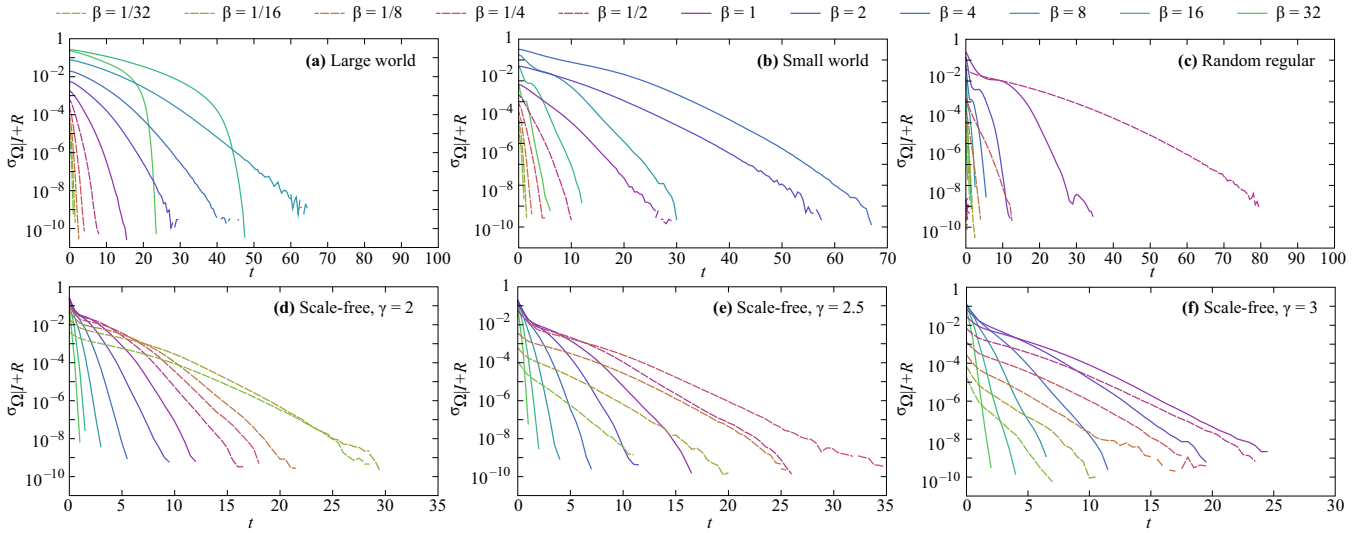
FIG. 10. (Color online) Plot corresponding to Fig. 8 for the standard deviation of the final outbreak sizes conditioned on $I + R$, that is, the number of infected and recovered vertices at t.

values near extinction. In the end of the process, $\sigma_{\Omega|I}$ decreases with time t because all the outbreaks die out and become perfectly predictable (as $I = 0$). These observations suggest that $\sigma_{\Omega|I}$ is a good predictor early in the outbreak, but that it becomes less reliable with time.

By contrast to the nonmonotonic behavior of $\sigma_{\Omega|I}$, the other measures $\sigma_{\Omega|R}$ and $\sigma_{\Omega|I+R}$ are nonincreasing functions with time except for the range of very small t (see Figs. 9 and 10). This difference between $\sigma_{\Omega|I}$ and $\sigma_{\Omega|R}$ ($\sigma_{\Omega|I+R}$) can be explained with the intrinsic behavior of I and R as follows. It is well known that, during the epidemic process, the I value increases and then decreases to zero over time [7] (also see Fig. 3). This nonmonotonicity of I results in the bump of the $\sigma_{\Omega|I}$ curves, as we discussed above. On the other hand, the R ($I + R = S$) value increases (decreases) monotonically over time [7]. If we observe a large R value at t, we expect a larger Ω value in the end of the process. In this way, this monotonicity of R and $I + R$ results in the monotonic decrease of the $\sigma_{\Omega|R}$ and $\sigma_{\Omega|I+R}$ curves. The results shown in Figs. 9 and 10 imply that the predictability of the SIR processes increases with time if we know the R value (and the I value in addition). In particular, the $\sigma_{\Omega|I+R}$ curves exhibit the dependence on different β values, which is very similar to the σ_Ω curves shown in Fig. 3. This observation is supported by the decay exponents of the $\sigma_{\Omega|I+R}$ curves fitted with an exponential form (Fig. 11), which indicate a dependence on β similar to that of σ_Ω (see Fig. 6). In other words, with such a summarized information about the vertices' states, we can obtain the same accuracy in our predictions (in terms of $\sigma_{\Omega|I+R}$) as when we have the perfect information about the system for the network models we consider.

V. SUMMARY AND DISCUSSION

In this paper, we numerically investigated the predictability of outbreaks of the SIR epidemic process on static network models. We used the standard deviation of the final outbreak sizes conditioned on the state of vertices at a given time as a key quantity for predictability (technically speaking unpredictability, since zero standard deviation means that the

outcome is perfectly predictable). We considered the two scenarios of the information available: First, we have perfect information of the system [where we know the exact state of the vertices $(\mathbf{s}, \mathbf{t}^I)$]. Second, we have partial information—to be precise, only the number of infected I or recovered R. In the first scenario, the standard deviation $\sigma_{\Omega|\mathbf{s},\mathbf{t}}$ decays like an exponential function of time, or slower, for all the types of network models we considered. This result is consistent with the previous theoretical result that the characteristic quantities of the SIR process asymptotically converge to deterministic functions [4]. The time constant of the exponential decay, however, is highly model-dependent. We saw a monotonic
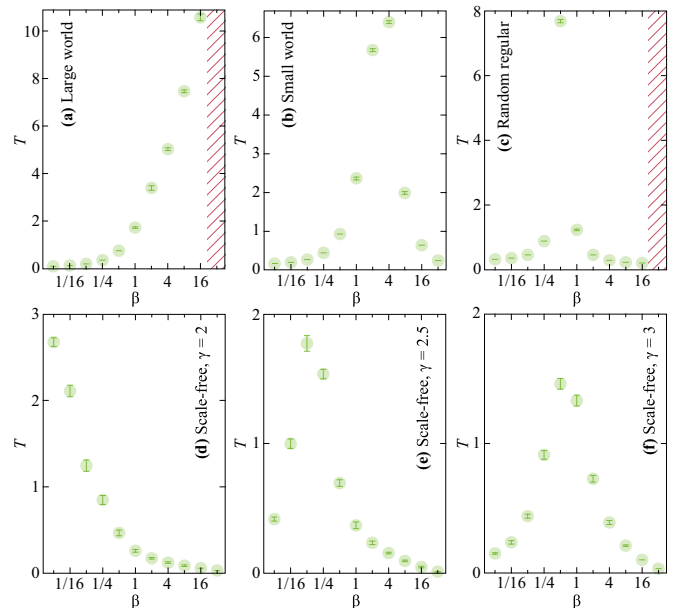

FIG. 11. (Color online) The decay constants from nonlinear least-squares fits to exponential forms of the curves of $\sigma_{\Omega|I+R}$ shown in Fig. 10. The shaded regions indicate where σ_Ω does not fit an exponential function. The error bars represent standard errors.

increase of the decay exponent for the large- and small-world networks, whereas a monotonic decrease was observed for the scale-free networks with $\gamma = 2$. For the random regular graph and the scale-free networks with $\gamma = 2.5$ and 3, there is a single peak of the decay exponent at an intermediate value of β. These results shed new light on the notion of predictability of the SIR process, i.e., the uncertainty in epidemic outbreaks is a nontrivial question even if we possess the information of contact network structure and vertices' states.

For the second scenario, the standard deviations $\sigma_{\Omega|R}$ and $\sigma_{\Omega|I+R}$ decay with exponential tails, while $\sigma_{\Omega|I}$ exhibits a non-monotonic change in time. Specifically, for large transmission rates it has a peak for intermediate times. In other words, some time after the outbreak starts, we can expect that knowing, e.g., the prevalence becomes less valuable. It is also in this region where there is a large difference between the two scenarios, i.e., knowing the exact configuration of infectious, susceptible, and recovered individuals improves the predictability a great deal. Another way of seeing this is that for most parts of the parameter space, the additional information in our first scenario is rarely of much use, although it is much more information ($2N$ numbers as opposed to one number).

Connecting back to our starting point—how could real outbreaks be hard to predict when the SIR model itself converges to deterministic quantities—we see, as mentioned, that for some parameter values, assuming only aggregate information, the convergence is slow. A possible answer is that the discrepancy comes from the fact that in practice we only have this kind of aggregate information. On the other hand, for many parameter values, also in the case of population-level information, the decay of predictability is very fast (also compared to the duration of the outbreak). For this reason, we also believe that the SIR model (i.e., the assumptions behind it), to some extent, underestimates the outbreak diversity.

This work paves the way for more detailed investigations of predictability of epidemic processes. For example, our work assumes that we know the starting point of the epidemic outbreak, which is fairly unrealistic. It would be interesting to investigate a situation in which that assumption is relaxed. This would connect our line of research to the question of identifying the source of an outbreak [22–24] or reconstructing likely transmission trees [25,26].

[1] D. J. Sencer and J. D. Miller, Emerg. Infect. Dis. **12**, 29 (2006).

[2] D. Lazer, R. Kennedy, G. King, and A. Vespignani, Science **343**, 1203 (2014).

[3] J. Iwai and S. Sasa, arXiv:1303.6606.

[4] S. Janson, M. Luczak, and P. Windridge, Random Struct. Algor. **45**, 726 (2014).

[5] M. J. Keeling and K. Eames, J. R. Soc. Interface **2**, 295 (2005).

[6] *Network Epidemiology: A Handbook for Survey Design and Data Collection*, edited by M. Morris (Oxford University Press, Oxford, 2004).

[7] R. Pastor-Satorras, C. Castellano, P. van Mieghem, and A. Vespignani, arXiv:1408.2701.

[8] J. O. Lloyd-Smith, S. J. Schreiber, P. E. Kopp, and W. M. Getz, Nature (London) **438**, 355 (2005).

[9] R. Pastor-Satorras and A. Vespignani, Phys. Rev. Lett. **86**, 3200 (2001).

[10] M. Molloy and B. A. Reed, Random Struct. Algor. **6**, 161 (1995).

[11] P. Holme, J. Logist. Eng. Univ. **30**, 1 (2014).

[12] S. Kininmonth, M. Drechsler, K. Johst, and H. P. Possingham, Mar. Ecol. Prog. Ser. **417**, 139 (2010).

[13] D. J. Watts and S. H. Strogatz, Nature (London) **393**, 440 (1998).

[14] B. D. MacKay, J. Algor. **11**, 52 (1990).

[15] M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).

[16] F. Liljeros, C. R. Edling, and L. A. N. Amaral, Microbes Infect. **5**, 189 (2003).

[17] F. Liljeros, J. Giesecke, and P. Holme, Math. Pop. Stud. **14**, 269 (2007).

[18] G. Caldarelli, *Scale-Free Networks: Complex Webs in Nature and Technology* (Oxford University Press, Oxford, 2007).

[19] Y. Moreno, R. Pastor-Satorras, and A. Vespignani, Eur. Phys. J. B **26**, 521 (2002).

[20] P. Holme, PLOS ONE **8**, e84429 (2013).

[21] We use the Gauss-Newton nonlinear least-squares method from the statistics software R, http://www.r-project.org/.

[22] A. Y. Lokhov, M. Mézard, H. Ohta, and L. Zdeborová, Phys. Rev. E **90**, 012801 (2014).

[23] C. H. Comin and L. da Fontoura Costa, Phys. Rev. E **84**, 056105 (2011).

[24] N. Antulov-Fantulin, A. Lancic, T. Smuc, H. Stefancic, and M. Sikic, arXiv:1406.2909.

[25] A. S. Walker *et al.*, PLOS Med. **9**, e1001172 (2012).

[26] T. Jombart, A. Cori, X. Didelot, S. Cauchemez, C. Fraser, and N. Ferguson, PLoS Comput. Biol. **10**, e1003457 (2014).