Society for
Mathematical
Biology

ORIGINAL ARTICLE

Estimating Initial Epidemic Growth Rates

**Junling Ma · Jonathan Dushoff ·
Benjamin M. Bolker · David J.D. Earn**

Abstract The initial exponential growth rate of an epidemic is an important measure of disease spread, and is commonly used to infer the basic reproduction number $\mathcal{R}_0$. While modern techniques (e.g., MCMC and particle filtering) for parameter estimation of mechanistic models have gained popularity, maximum likelihood fitting of phenomenological models remains important due to its simplicity, to the difficulty of using modern methods in the context of limited data, and to the fact that there is not always enough information available to choose an appropriate mechanistic model. However, it is often not clear which phenomenological model is appropriate for a given dataset. We compare the performance of four commonly used phenomenological models (exponential, Richards, logistic, and delayed logistic) in estimating initial epidemic growth rates by maximum likelihood, by fitting them to simulated epidemics with known parameters. For incidence data, both the logistic model and the Richards model yield accurate point estimates for fitting windows up to the epidemic peak. When observation errors are small, the Richards model yields confidence intervals with better coverage. For mortality data, the Richards model and the delayed logistic model yield the best growth rate estimates. We also investigate the width and coverage of the confidence intervals corresponding to these fits.

J. Ma (✉)
Mathematics and Statistics, University of Victoria, Victoria, BC, Canada
e-mail: junlingm@uvic.ca

J. Dushoff
Department of Biology, McMaster University, Hamilton, ON, Canada

B.M. Bolker · D.J.D. Earn
Mathematics and Statistics, McMaster University, Hamilton, ON, Canada

1 Introduction

The early phase of an epidemic of an infectious disease is usually characterized by exponential growth in the number of cases or deaths. The rate of this exponential growth, r, determines baseline projections of epidemic growth, and thus changes in r can be used as a robust measure of the effectiveness of control measures (or increases in pathogen transmissibility). Additionally, r can be used to estimate the basic reproduction number $\mathcal{R}_0$ (the average number of secondary infections caused by a typical infectious individual in a fully susceptible population). $\mathcal{R}_0$ is determined by r and the distribution of the generation interval (the time between a case and the secondary cases resulting from it) (Lipsitch et al. 2003), a fact that has been used in many applications (e.g., Chowell et al. 2003; Mills et al. 2004; Wearing et al. 2005; Wallinga and Lipsitch 2007; Roberts and Heesterbeek 2007; Nishiura et al. 2009b, 2010).

Despite the fact that discussion of $\mathcal{R}_0$ dominates both the theoretical and applied epidemiology literature, knowing r has many advantages compared with relying solely on the information contained in $\mathcal{R}_0$. Firstly, r is a measure of the speed of epidemic growth, conveying information about the time scale of disease spread. In contrast, $\mathcal{R}_0$ is a pure number with no associated time scale; epidemics with the same $\mathcal{R}_0$ can occur over vastly different time periods, ranging from days to years. Knowing the epidemic time frame can be critical for selection of disease control strategies. Secondly, r itself is independent of potentially uncertain knowledge about the generation interval distribution, and thus may be useful in comparing the severity of disease epidemics.

The growth rate r is most commonly estimated by fitting an exponential curve to the initial growth phase. The methods used can be as simple as least-squares fitting of an exponential curve to incidence (or cumulative incidence), or of a straight line to the logarithm of incidence (or of cumulative incidence) (e.g., Chowell et al. 2003; Mills et al. 2004; Vynnycky et al. 2007). Formal statistical inference from least-squares fitting is based on assuming independent, normally distributed errors with constant variance, assumptions that can easily be improved upon for this application. Consequently, over the years, more sophisticated methods have been applied, including Poisson regression (e.g., de Silva et al. 2009) and methods based on branching processes (Roberts and Heesterbeek 2007; Nishiura et al. 2009a, 2009b).

All of these methods, however, rely on fitting to an approximately exponential growth phase, which is usually short (due to depletion of susceptibles) and selection of a temporal fitting window requires an independent procedure. The fitting window is typically chosen by a reasonable heuristic (e.g., Chowell et al. 2007) use a goodness of fit method), but such methods are based on ad hoc decision rules; moreover, they do not alleviate the problem that there are usually very few data points during the exponential growth phase.

To use a longer sequence of data points, growth rate estimation methods must directly or indirectly take account of the depletion of susceptibles. This can sometimes be achieved with a mechanistic approach by fitting a transmission model to the epidemic curve, and estimating r simultaneously with other disease parameters (e.g., Pourabbas 2001; Chowell et al. 2006a, 2006b). This approach has the advantage of

facilitating estimation of the effects of external factors such as seasonality and control measures (Bootsma and Ferguson 2007; He et al. 2011), but it unavoidably relies on the appropriateness of the assumed transmission model. For example, most commonly used transmission models (unrealistically) assume a homogeneously mixed population, yet different underlying contact networks can result in very different parameter estimates that yield similar epidemic growth rates (e.g., Ma et al. 2013). In fact, different mechanistic models may fit equally well but give different answers (Wearing et al. 2005). Appropriate transmission models often require a relatively large number of parameters, making fits less stable, particularly if data are limited.

Phenomenological models (e.g., Hsieh et al. 2010) dispense with mechanistic assumptions. Instead, they make general assumptions about the shape of the incidence curve. In general, such simple phenomenological models use fewer parameters, make more straightforward assumptions, and are likely to give more robust estimates when applied in contexts with very limited data (a common situation for historical epidemics or near the start of an epidemic of an emerging or reemerging disease).

An important complication is that the best data for historical epidemics are often based on mortality records rather than disease incidence (e.g., Mills et al. 2004; Goldstein et al. 2009; He et al. 2011). From a deterministic perspective, this makes little difference: a disease spreading exponentially will have, on average, mortality rates that increase at the same exponential rate as incidence. When estimating parameters from noisy data, however, both the "sampling process" (only a fraction of infected individuals die), and the potentially broad distribution of times between infection and death, can substantially complicate the estimation process. Methods of accounting for mortality include sensitivity analyses, where artificial data are generated using a realistic death process and used to validate simple models (e.g., Mills et al. 2004); mechanistic models in which the death process is simulated as part of the fitting procedure (e.g., He et al. 2011); and deconvolution methods, which attempt to estimate the incidence time series implied by mortality data, case fatality proportions, and generation time distributions, as a separate first step in analysis (e.g., Goldstein et al. 2009).

Phenomenological models that are fit to a given dataset are currently chosen in an ad hoc manner. There is a lack of information on the performance of commonly used models. In this paper, we provide a guide for choosing the appropriate phenomenological model. We compare the performance of the exponential, the logistic, the Richards, and the delayed logistic models applied to simulated stochastic epidemics (for which the correct answer is known). The logistic model is used as a starting point for emulating the characteristic shape of *cumulative* incidence curves, which first grow exponentially, and then begin to level off. The logistic model can be generalized to the Richards model (Richards 1959; Hsieh et al. 2010), which has a parameter determining how fast the switch from an accelerating to decelerating cumulative epidemic curve occurs. For mortality data, we also consider the "delayed logistic model," which generalizes the logistic model to include an explicit delay from infection to death (Bootsma and Ferguson 2007). These phenomenological models are closely related to epidemic dynamics. In fact, the exponential model and the logistic model are the first- and second-order approximations to the growth phase of an epidemic curve produced by the stan-

dard Kermack–McKendrick SIR model (see, e.g., Kermack and McKendrick 1927; Kendall 1956).

Because observations drawn from the same cumulative curve are correlated, we base all of our fits on "interval incidences" (daily or weekly incidence). We fit maximum likelihood parameters for each of the four phenomenological models we consider, assuming that observation errors follow a Poisson distribution (see Sect. 2). We also investigate sensitivity to the choice of the fitting window, and the effect of the reporting ratio and time-to-death distribution.

2 Methods

2.1 Models Describing the Epidemic

Some of the phenomenological models that we consider in this paper use simple closed-form expressions to describe the *cumulative* growth of the epidemic. Most of our fitting, however, will be based on "interval" incidences, obtained by differencing cumulative expressions where necessary, e.g., $x(t) = c(t + \Delta t) - c(t)$, where $c(t)$ is cumulative incidence, and $x(t)$ is interval incidence (typically daily or weekly).

2.1.1 Exponential

The simplest model of exponential growth is

$$x(t) = x_0 e^{rt}. \tag{1}$$

This model has two parameters that need to be estimated: the initial value x_0 and the growth rate r.

2.1.2 Logistic

In an epidemic, cumulative incidence initially grows exponentially, but eventually slows and approaches a limit. This behavior is qualitatively similar to that of a logistic curve. Thus, a logistic model may allow us to use longer sequences of data from the beginning of an epidemic, by accounting for the epidemic slowing as it proceeds. In this model, the expected cumulative number of cases $c(t)$ is assumed to satisfy the following equation:

$$c'(t) = rc(t)\left[1 - \frac{c(t)}{K}\right], \tag{2}$$

where K is the final size of the epidemic, which $c(t)$ approaches. This equation has an explicit solution:

$$c(t) = \frac{K}{1 + [(K/c_0) - 1]e^{-rt}}, \tag{3}$$

where c_0 is the total number of cases observed at time $t = 0$. We obtain $x(t)$ by differencing as explained above.

2.1.3 Richards

In the Richards model (Richards 1959; Banks 1993; Hsieh et al. 2010), the cumulative epidemic curve $c(t)$ satisfies

$$c'(t) = rc(t)\left[1 - \left(\frac{c(t)}{K}\right)^a\right]. \tag{4}$$

This model is also called the power law logistic model, and the logistic model is a special case with $a = 1$. When $a \ll 1$, $[c(t)]^a \approx 1$. Consequently, for sufficiently small a there is effectively no density dependence in the growth rate in this model: $c'(t) \approx r(1 - 1/K^a)c(t)$. The contribution of K to the exponential growth rate makes it difficult to estimate r precisely. We therefore reparameterize the model, and use $r_0 = r[1 - (c_0/K)^a]$ (the exponential growth rate when $t = 0$) instead of r as our estimate of the exponential growth rate.

Equation (4) is also solvable and, in terms of the identifiable parameter r_0, $c(t)$ has the explicit form

$$c(t) = \frac{K}{(1 + [(K/c_0)^a - 1]\exp\{-r_0 t/[1 - (c_0/K)^a]\})^{1/a}}. \tag{5}$$

This model has three parameters: c_0, K, and the initial growth rate r_0.

2.1.4 Delayed Logistic (Mortality)

Our final model makes a small step away from phenomenological models to incorporate some of the mechanisms that may determine the shape of epidemic curves in the particular case of mortality (as opposed to case reporting) data. If the logistic model can be used to describe the cumulative incidence, then cumulative deaths can be explicitly modeled as a delay to the incidence curve $x(t)$. If we assume that the delay is exponentially distributed with rate m, the cumulative deaths $d(t)$ can be modeled as

$$d(t) = \int_0^t c_{\text{Log}}(s)e^{-m(t-s)}\,ds, \tag{6}$$

where $c_{\text{Log}}(t)$ is the solution to the logistic model (3).

We consider only the delayed logistic, and not the delayed Richards model, for mortality data because fitting either a or m, together with c_0 and r, early in the epidemic raises issues with statistical identifiability: A delayed Richards model appears to be impractical for our purposes.

2.2 Maximum Likelihood Estimation

We estimate parameters using maximum likelihood, assuming that the underlying process is deterministic, and that incidence during each reporting period follows a Poisson distribution. Here we are interested in the likelihood of our model, given an observed epidemic curve. We assume that the ith observation Y_i is a Poisson-distributed random variable, with mean equal to the model prediction $y(t)$. It is then

straightforward to calculate the likelihood. We then numerically search for the model parameters that maximize the likelihood function L (Bolker 2008, p. 170).

In addition to the maximum likelihood point estimate (MLE) of the parameters (we are particularly interested in r, the estimated growth rate), the likelihood approach also provides a framework for estimating confidence intervals. If the data set is sufficiently large, the confidence intervals of a focal parameter can be found by constructing the *likelihood profile*, the curve of maximum likelihood achievable by optimizing over all of the non-focal parameters when the focal parameter is held fixed at specific values away from its MLE, and finding the points on the profile where the difference in log-likelihood from the maximum is equal to half of the upper critical tail value (e.g., the 95th percentile) of the χ^2 distribution with one degree of freedom (Bolker 2008, p. 192).

2.3 Interval vs. Cumulative Epidemic Curves

The procedure we have described can be used to fit models to cumulative incidence (or mortality) rather than the corresponding interval incidence. Mathematically, cumulative and interval curves carry the same information, but statistically speaking, we expect different results from the two approaches. Likelihood calculations (including the special case of least-squares fitting) assume that the errors in individual observations are statistically independent. This assumption is particularly inappropriate for cumulative curves, where each observation contains all of the cases from prior observations. Some researchers use techniques such as parametric bootstrapping to address this issue (Chowell et al. 2006a, 2007), while others are apparently unaware of the problem (Roberts and Heesterbeek 2007). To assess the potential importance of correlated errors in fits to cumulative data, we compare results from fits to both cumulative and interval epidemic curves.

2.4 Process Error vs. Observation Error

Fitting epidemic data with a deterministic underlying curve makes an additional strong assumption: that the variation around the observed epidemic curve is due entirely to observation error with no contribution from process error (Bolker 2008, p. 344). In the presence of process error—stochastic variation in the underlying epidemic curve, not just in the observed curve—the observations will again be correlated, because random events will carry over across multiple time steps (because they affect the underlying epidemic process). Some epidemic modelers have used sophisticated methods to disentangle process from observation error (Bjørnstad et al. 2002; Ionides et al. 2006). However, such methods are extremely data-hungry, and are usually used only for diseases that persist over multiple years or decades. Consequently, they are unlikely to be applicable to outbreak prediction or retrospective analysis of individual epidemics.

2.5 The Simulated Epidemic Curves

We fit the phenomenological models to simulated epidemics for which the true values of the parameters are known. To begin with, we study the performance of these

Susceptible   Exposed   Infectious   Removed
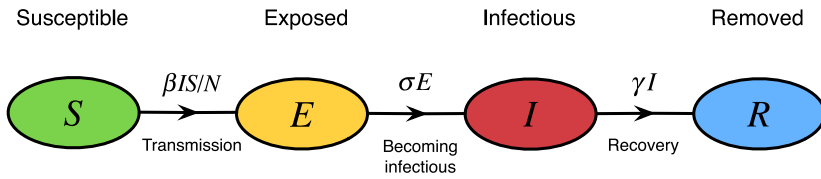
Fig. 1 Schematic representation of the SEIR model. The population is classified into 4 compartments: the susceptible (*S*), exposed (or latent, *E*), infectious (*I*), and recovered (*R*)

models in the absence of any noise in the data (i.e., we fit to a deterministic epidemic simulation). If a model does not perform well in this case, it should certainly not be applied to real data. We then fit the phenomenological models to epidemic curves generated by stochastic simulations. We first simulate without observation error, in order to test the performance of the phenomenological models on epidemic data for which noise is dominated by process errors. Binomial observation errors are then added to simulate either incidence data with a small reporting ratio or mortality data.

We use a standard compartmental Susceptible-Exposed-Infectious-Recovered (SEIR) model. All individuals are susceptible at the start of the epidemic except for the infected index case. Upon contact with an infectious individual, susceptible individuals become exposed (infected but not yet infectious). They leave the exposed state at constant rate σ and become infectious. They leave the infectious state at a constant rate γ, recovering to a state of permanent immunity. We assume that effective contacts (those that result in a new infection) occur at a rate $\frac{\beta}{N}SI$, where N is the population size, and S and I are respectively the numbers of susceptible and infectious individuals (Fig. 1).

In the deterministic limit (population size $N \to \infty$), the SEIR model can be expressed as a system of ordinary differential equations (Diekmann and Heesterbeek 2000),

$$\frac{dS}{dt} = -\frac{\beta}{N}SI, \tag{7a}$$

$$\frac{dE}{dt} = \frac{\beta}{N}SI - \sigma E, \tag{7b}$$

$$\frac{dI}{dt} = \sigma E - \gamma I, \tag{7c}$$

$$\frac{dR}{dt} = \gamma I. \tag{7d}$$

Linearizing this model about the disease-free equilibrium ($S = N$, $E = I = R = 0$), and computing the dominant eigenvalue, we find the initial exponential growth rate is

$$r = \frac{1}{2}\left[-(\sigma + \gamma) + \sqrt{(\sigma - \gamma)^2 + 4\beta\sigma}\right]. \tag{8}$$

Our stochastic simulations are carried out with population size $N = 10^6$ and discrete time steps of size $h = 0.1$ of a simulated day. At each time step, the number of

new infections is drawn from a Poisson-distributed random number generator with rate $\beta SIh/N$; similarly, the number of individuals becoming infectious or recovering are Poisson random variables with rates σEh and γIh, respectively. To simulate mortality curves, we assume that a fixed fraction ϕ of infected individuals will eventually die from the disease, and we assume the time from infection to death is Gamma distributed with shape parameter 3 and mean 30 days.

2.6 Protocol

We compare the performance of each of the phenomenological models described in Sect. 2.1 as epidemic growth rate estimators by fitting to simulated incidence and mortality curves generated with the stochastic transmission model described in Sect. 2.5. We also examine how each growth rate estimate depends on the time window selected from the epidemic curve.

3 Results

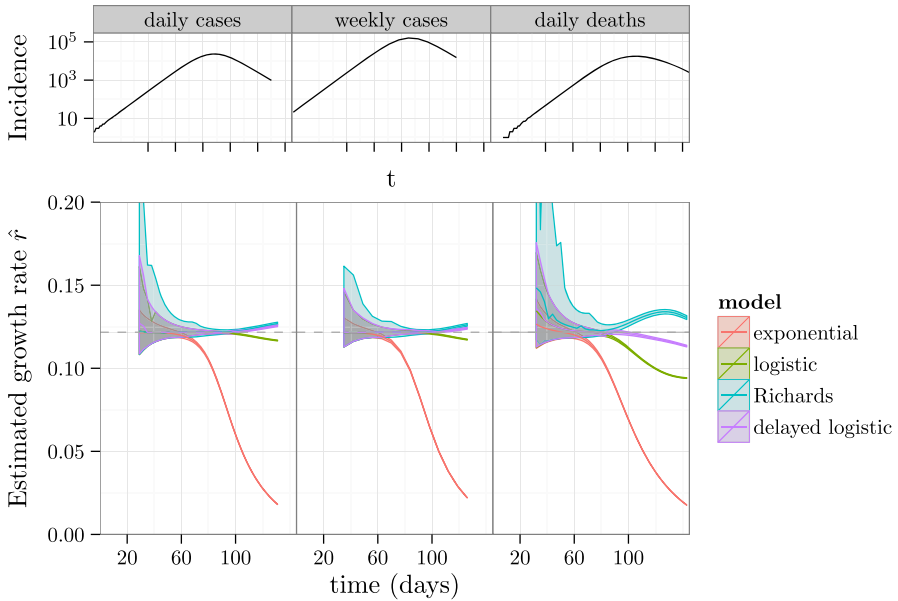### 3.1 Estimation of Initial Growth Rate r

Which model is best for estimating the initial growth rate of an epidemic depends on our fitting window, i.e., precisely which observed data points we include in our analysis. In a real-time application, data accumulate over time and we have the potential to keep increasing the length of our fitting window. In a retrospective study, we can in principle use the entire epidemic curve, but doing so will not necessarily yield the best estimate of the initial growth rate. In either case, we consider which model is most appropriate as a function of the fitting window.

Real-Time Growth Rate Estimation   With a real-time application in mind, we plot (for each model) our point estimate for the initial exponential growth as a function of the end point of the fitting window, fixing the starting point of the fitting window at the time when the disease in question was first detected in the population.

The left column in Fig. 2 shows the results of fitting each model to the incidence computed from a solution of the deterministic SEIR model. The plotted estimate at each time is based on a fitting window from time $t = 0$ until that time. Because the growth rate of the epidemic decreases as susceptibles are depleted, the point estimates from the exponential model deviate most quickly from the true growth rate, while all the other models stay close to the true value for all fitting windows up to (and even beyond) the peak of the epidemic. The middle column in Fig. 2 shows that weekly aggregation of the epidemic curve does not alter the results substantially. However, when fitting to simulated mortality data (the right column in Fig. 2), the estimated initial growth rate from the logistic model starts to deviate substantially from the true value after the epidemic reaches its peak, while the estimates from the Richards and delayed logistic models remain close to the true value over the whole period.

Figure 3 shows the results of fitting a single realization of the stochastic SEIR model, with the same disease parameters as with the deterministic model analyzed in Fig. 2. The estimates are similar, but stochasticity leads to wider confidence intervals.
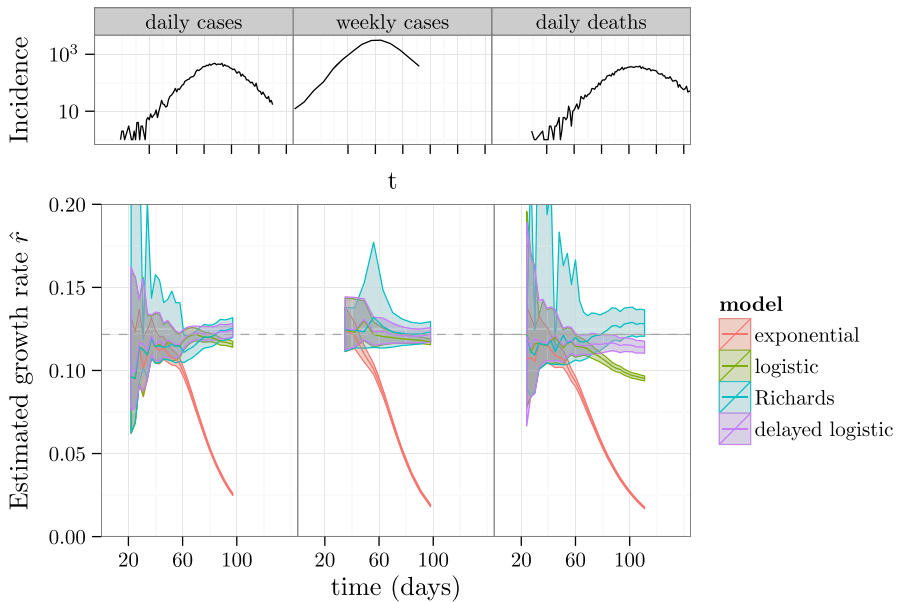
Fig. 2 Estimation of initial epidemic growth rates using each of the phenomenological models described in Sect. 2.1, for simulated daily incidence, weekly incidence, and daily mortality, obtained using the deterministic SEIR model (Eq. (7a)–(7d)) with true growth rate $r \approx 0.121$ (shown by *dashed gray horizontal line*); basic reproduction number $\mathcal{R}_0 = 2$; mean latent period $1/\sigma = 2$ days and mean infectious period $1/\gamma = 5$ days. The time from incidence to death is gamma distributed with shape parameter 3 and 30-day mean. The *top panel* shows the simulated epidemic curve on a logarithmic scale. The *lower panel* shows the point estimates and 95 % confidence intervals at the end time of each fitting window (which always start at time $t = 0$)

Retrospective Growth Rate Estimation   In a retrospective context, we have the luxury of being able to choose any fitting window we like, though we might be constrained by having missed the beginning of the epidemic or because reporting rates increased dramatically part way through the epidemic. This raises the question of how estimated initial growth rates depend on the start, as well as the end, of the fitting window. To explore this question, we fix the end point of the fitting window at the peak of the epidemic and consider starting points ranging from $t = 0$ to the time corresponding to 4 data points before the epidemic peak.

Figure 4 shows results for fits to simulated daily and weekly incidence and daily deaths. In all cases, the exponential model is sensitive to the starting time (the exponential fits always underestimate the true exponential growth rate because they fail to account for the slowing of exponential growth as the epidemic proceeds). For incidence data, all other models are insensitive to the starting time. For death data, the point estimates start to deviate from theoretical value as the starting time increases. The delayed logistic model is the least sensitive model. Fits to stochastic realizations (not shown) are similar.

Fig. 3 Estimation of initial growth rates from stochastic epidemic simulations. Each phenomenological model from Sect. 2.1 was fitted to an epidemic curve from a single realization of the stochastic SEIR model with the same parameters as for the deterministic model used for Fig. 2 (the layout of the figure is also the same as for Fig. 2)

3.2 Accuracy of Estimates of *r*: Coverage Probabilities

Process error (demographic noise) will lead to different initial growth rates for different stochastic realizations of any given epidemic model. Consequently, the initial growth rate that we estimate for a specific realization will not necessarily be close to the true growth rate of the model, even if it perfectly represents the initial growth of the particular simulation in question. Similarly, since real epidemics represent single realizations of a stochastic epidemic process, the initial growth rates we estimate from observed epidemics may not reflect the true growth rate associated with the underlying transmission process. This is not a problem if our goals do not extend beyond estimating the pattern of the specific epidemic in question, but if our ultimate goal—as is often the case—is to estimate the underlying properties of the disease, such as its basic reproduction number $\mathcal{R}_0$, then we are led to an important question. Does our confidence interval for the initial growth rate of a given realization (or observed epidemic) include the true growth rate of the underlying process?

   The probability that a confidence interval contains the true value of a parameter is called the *coverage probability*. If the confidence interval is computed by correctly accounting for all sources of variation in the data, then the coverage probability should equal the confidence level of the interval. In the present context, examining the coverage probability is not simply a consistency check, because our confidence intervals are computed by assuming that all stochastic variation arises from observation error, with no contribution from process error.
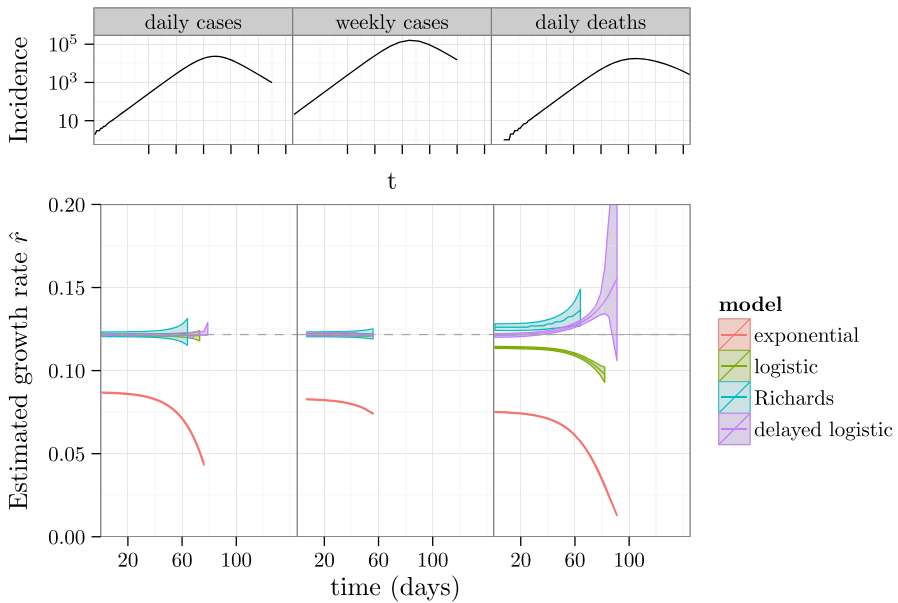
Fig. 4 Dependence of estimated initial epidemic growth rates on the start time of the fitting window. Each of the models of Sect. 2.1 was fitted to a simulated deterministic epidemic curve (generated by the deterministic SEIR model). The *bottom panels* show the estimated growth rate (and confidence intervals) versus the start time of the fitting window (the end of the fitting window is fixed at the peak of the epidemic). The disease parameters and figure layout are the same as in Fig. 2

We estimated coverage by fitting each model to 1,000 realizations of a stochastic SEIR model, with $N = 10^6$ and parameters chosen to be similar to estimates for the 1918 influenza pandemic (Mills et al. 2004; He et al. 2011): transmission rate $\beta = 0.4$ per day, mean latent period $1/\sigma = 2$ days, and mean infectious period $1/\gamma = 5$ days. The underlying deterministic model thus has a basic reproduction number $\mathcal{R}_0 = 2$ and initial growth rate $r = 0.12$.

The first three columns in the upper panel of Fig. 5 show the coverage probability of the 95 % confidence interval of the growth rate as a function of the length of the fitting window when fitted to daily incidence with perfect reporting, daily incidence with a 2 % reporting ratio, and daily mortality with 2 % case fatality, respectively. If the confidence intervals are correct, the coverage should be 95 %. The lower panel shows the corresponding width of the confidence interval. Figure 6 shows the distributions of the point estimates for fitting windows from $t = 0$ to 1/3 of the peak time, 2/3 of the peak time, the peak time, and 3/2 of the peak time.

For incidence without reporting error (the left column in Fig. 5), the coverages for all models are poor. Since the point estimates are well centered about the true value, as shown in the left column in Fig. 6 (i.e., the level of bias is low), the poor coverage is attributable to estimated confidence intervals that are too narrow. For incidence with a small reporting ratio (the second column in Fig. 5), the coverages are much closer to the estimated confidence level than those without reporting error. In fact, the coverage probability of the Richards model is very close to 95 % when the end of the
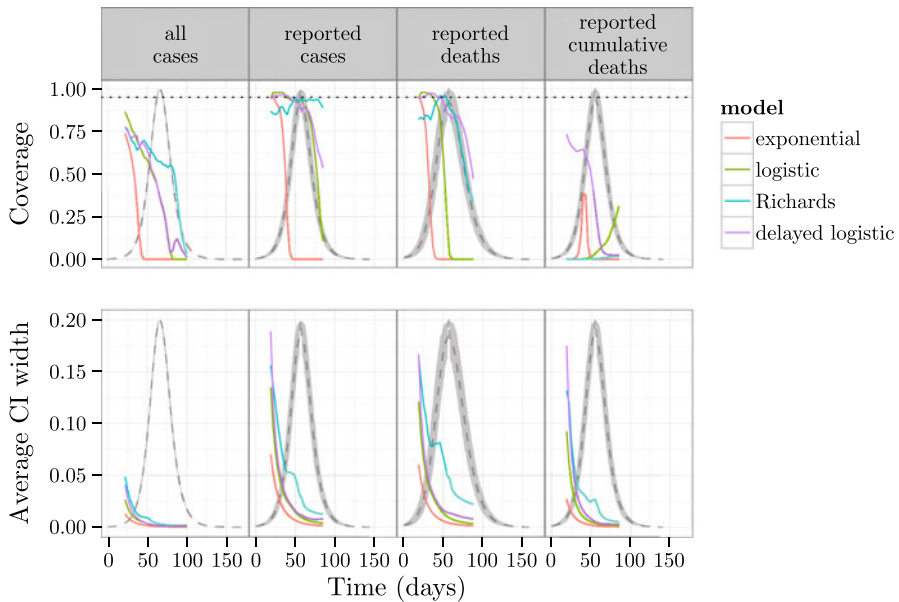
Fig. 5 Statistical coverage of true underlying growth rates. The first three columns in the top panel show the coverage probabilities of the confidence intervals of the growth rate, estimated by fitting each model from Sect. 2.1 to 1000 realizations of the stochastic SEIR model, based on daily incidence, daily incidence with a 2 % reporting ratio, and daily mortality with a 2 % case fatality, respectively. The *last column* shows the coverage probability of the confidence intervals estimated from each phenomenological model based on fits to cumulative daily deaths. The *lower panel* shows the corresponding widths of the confidence intervals (averaged over the fits to the 1000 simulations). The *horizontal axis* is the length of the fitting window (with the starting time always at $t = 0$). The *dashed lines* are the averages of the scaled stochastic simulations (the *grey shades*), illustrating the fitting window position

fitting window is near the peak of the epidemic, while the coverages of the logistic and delayed logistic models are approximately 90 % at the peak. For the daily mortality curve with a small (2 %) case fatality proportion (the third column in Fig. 5), the coverages are good, but not quite as good as for 2 % sampling of incidence. With the fitting window ending near the epidemic peak, the growth rate coverages of the Richards and delayed logistic models are 88 % and 84 %, respectively, while that of the logistic model is much lower. However, the delayed logistic model achieves comparable coverage to the Richards model with much narrower confidence intervals.

3.3 Fitting to Cumulative Epidemic Curves

As mentioned in the Introduction, cumulative epidemic curves are sometimes used to fit epidemic growth rates. To show the effect of using cumulative curves, we fit the observed cumulative daily mortality curve directly to models (1), (3), (5), and (6), without taking differences and converting them into epidemic curves. The result is shown in the last columns of Figs. 5 and 6. The coverage probabilities for all models are much smaller than for models fitted to epidemic curves, presumably because of the strong correlations between points on the same cumulative epidemic curve.
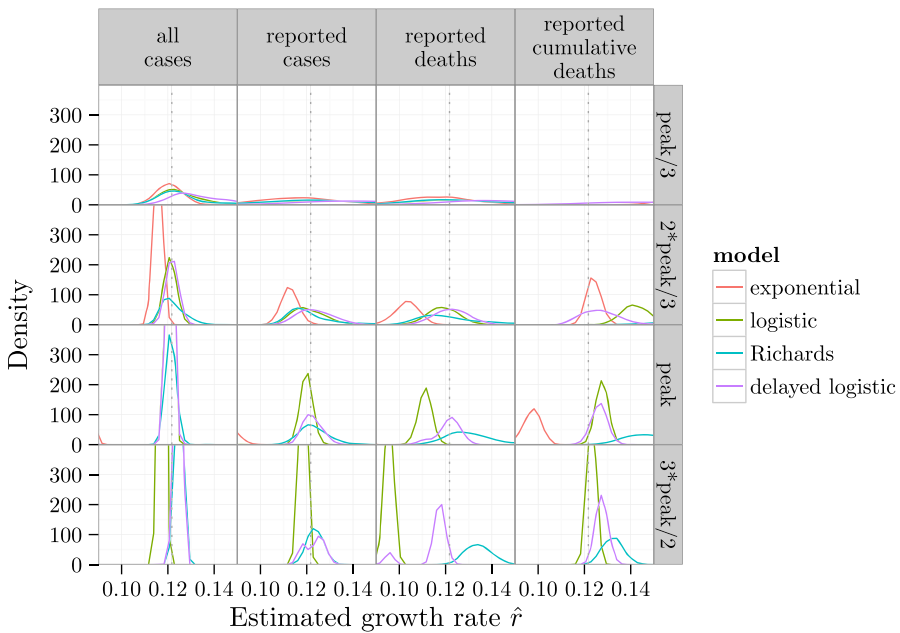
Fig. 6 The distribution of the point estimates of initial growth rate for each phenomenological model fitted to the 1000 realizations of stochastically simulated epidemics used for coverage estimates in Fig. 5. The *four rows* correspond to fitting windows ending at 1/3 of the peak time, 2/3 of the peak time, the peak time, and 3/2 of the peak time (all windows start at $t = 0$)

4 Discussion

Our goal in this paper has been to identify a reliable and practical method to obtain both a point estimate and confidence interval for the initial epidemic growth rate r associated with an observed infectious disease outbreak. We simulated outbreaks with specified growth rates and in each case compared the correct growth rate with the estimated growth rates obtained by fitting four phenomenological deterministic population models (pure exponential growth, the logistic model, the Richards model, and a "delayed logistic" model that includes a death process). We did not use mechanistic disease transmission models to estimate r since it is important to be able to estimate the growth rate without detailed knowledge of the process of disease transmission—which is unknown in many historical contexts and is often unclear in the early phase of an outbreak of an emerging disease.

The models we considered describe the growth phase of the theoretical cumulative epidemic curve. We obtained the theoretical "interval" epidemic curve (cases or deaths in each reporting interval) by differencing the theoretical cumulative curve before fitting to our simulated outbreaks. Because our underlying phenomenological models are deterministic, we are implicitly restricting attention to applications where process error is negligible compared to observation error (e.g., if the population is large and the disease is rarely reported, incidence rates will be affected less by demographic stochasticity than by sampling).

Our principal conclusions are the following.

Avoid Using the Exponential Model.   The pure exponential model should not be used because (i) it produces exceedingly narrow confidence intervals that have poor coverage probability, and (ii) the point estimates quickly deviate from the true value and it is difficult to determine the optimal fitting window. The other three phenomenological models all produce stable point estimates using fitting windows up to the epidemic peak, and confidence intervals with better coverage probabilities.

Use the Logistic and Richards Models for Incidence Data.   The logistic model gives accurate point estimates for incidence data. It also gives a narrow confidence interval with good coverage. With more data available, the Richards model also yields accurate point estimates, with more reliable confidence intervals.

Use the Delayed Logistic and Richards Model for Mortality Data   The Richards model and the delayed logistic model yield reasonably accurate point estimates for fitting windows up to the epidemic peak. For mortality data, it is important that the fitting window end before the peak.

Growth rate estimates obtained using the recommended models are relatively insensitive to the choice of the start of the fitting window as long as the starting point is in the exponential growth phase. Weekly aggregation of the epidemic curve does not appreciably alter the estimates for our parameters. This is potentially important because many infectious disease data (e.g., historical time series of notifiable infectious diseases) have been reported weekly.

Beware of Confidence Intervals from Phenomenological Fits.   Our method treats the process errors from the stochastic epidemic process as observation errors. Since process errors tend to propagate with time, they lead to biased estimates of the ensemble mean, which is described by the deterministic limit of the stochastic model. Hence, while the same disease parameters can generate different realized growth rates as a result of process error, our approach implicitly assumes that differences in parameter estimates are generated by actual differences in parameter values. This problem is mitigated when the observation error is large compared to process error, as is observed in fitting to stochastically simulated incidence data with large population sizes and small reporting (or case fatality) ratio.

Our analysis shows that the exponential growth rate of an epidemic can generally be estimated reasonably well using simple phenomenological models. This provides an important tool for gaining valuable information before details about the properties of the pathogen and the mechanisms of disease transmission are available. Fitting with phenomenological models can be done simply and quickly, and can yield timely estimates in the case of emerging infectious diseases.

When applying these phenomenological models to epidemic data, one has to be aware that in an emerging epidemic, public health interventions may cause sudden changes in the growth rate (see, e.g., Bootsma and Ferguson 2007), which requires special consideration. Behavior changes may also change the growth rate (see, e.g., He et al. 2013); these might be more gradual, and could possibly be handled by the Richards model with its shape parameter.

References

Banks, R. B. (1993). *Growth and diffusion phenomena: mathematical frameworks and applications* (Vol. 14). Berlin: Springer.

Bjørnstad, O. N., Finkenstädt, B. F., & Grenfell, B. T. (2002). Dynamics of measles epidemics: estimating scaling of transmission rates using a time series SIR model. *Ecol. Monogr.*, *72*, 169–184.

Bolker, B. M. (2008). *Ecological models and data in R*. Princeton: Princeton University Press. ISBN 0691125228.

Bootsma, M. C. J., & Ferguson, N. M. (2007). The effect of public health measures on the 1918 influenza pandemic in U.S. cities. *Proc. Natl. Acad. Sci. USA*, *104*, 7588–7593.

Chowell, G., Fenimore, P. W., Castillo-Garsow, M. A., & Castillo-Chavez, C. (2003). SARS outbreaks in Ontario, Hong Kong and Singapore: the role of diagnosis and isolation as a control mechanism. *J. Theor. Biol.*, *224*, 1–8.

Chowell, G., Ammon, C. E., Hengartner, N. W., & Hyman, J. M. (2006a). Transmission dynamics of the great influenza pandemic of 1918 in Geneva, Switzerland: assessing the effects of hypothetical interventions. *J. Theor. Biol.*, *241*, 193–204.

Chowell, G., Shim, E., Brauer, F., Diaz-Dueñas, P., Hyman, J. M., & Castillo-Chavez, C. (2006b). Modelling the transmission dynamics of acute haemorrhagic conjunctivitis: application to the 2003 outbreak in Mexico. *Stat. Med.*, *25*, 1840–1857.

Chowell, G., Nishiura, H., & Bettencourt, L. M. A. (2007). Comparative estimation of the reproduction number for pandemic influenza from daily case notification data. *J. R. Soc. Interface*, *4*, 155–166.

de Silva, U. C., Warachitand, J., Waicharoen, S., & Chittaganpitch, M. (2009). A preliminary analysis of the epidemiology of influenza A(H1N1) virus infection in Thailand from early outbreak data. *Euro Surveill.*, *14*, 19292.

Diekmann, O., & Heesterbeek, J. A. P. (2000). *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation*. *Wiley series in mathematical and computational biology*. New York: Wiley.

Goldstein, E., Dushoff, J., Ma, J., Plotkin, J. B., Earn, D. J. D., & Lipsitch, M. (2009). Reconstructing influenza incidence by deconvolution of daily mortality time series. *Proc. Natl. Acad. Sci. USA* doi:10.1073/pnas/092958.

He, D., Dushoff, J., Day, T., Ma, J., & Earn, D. J. D. (2011). Mechanistic modeling of the three waves of the 1918 influenza pandemic. *Theor. Ecol.*, *4*, 283–288.

He, D., Dushoff, J., Day, T., Ma, J., & Earn, D. J. D. (2013). Inferring the causes of the three waves of the 1918 influenza pandemic in England and Wales. *Proc. R. Soc. Lond. B, Biol. Sci.*, *280*, 20131345.

Hsieh, Y.-H., Fisman, D. N., & Wu, J. (2010). On epidemic modeling in real time: an application to the 2009 Novel A (H1N1) influenza outbreak in Canada. *BMC Res. Notes*, *3*, 283.

Ionides, E. L., Bretó, C., & King, A. A. (2006). Inference for nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA*, *103*, 18438–18443.

Kendall, D. G. (1956). Deterministic and stochastic epidemics in closed populations. In *Proceedings of the third Berkeley symposium on mathematical statistics and probability* (Vol. IV).

Kermack, W. O., & McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A, Math. Phys. Sci.*, *115*, 700–721.

Lipsitch, M., Cohen, T., Cooper, B., Robins, J. M., Ma, S., et al. (2003). Transmission dynamics and control of severe acute respiratory syndrome. *Science*, *300*, 1966–1970.

Ma, J., van den Driessche, P., & Willeboordse, F. H. (2013). The importance of contact network topology for the success of vaccination strategies. *J. Theor. Biol.*, *325*, 12–21.

Mills, C. E., Robins, J. M., & Lipsitch, M. (2004). Transmissibility of 1918 pandemic influenza. *Nature*, *432*, 904–906.

Nishiura, H., Castillo-Chavez, C., Safan, M., & Chowell, G. (2009a). Transmission potential of the new influenza A(H1N1) virus and its age-specificity in Japan. *Euro Surveill.* doi:10.1.1.170.1968.

Nishiura, H., Wilson, N., & Baker, M. G. (2009b). Estimating the reproduction number of the novel influenza A virus (H1N1) in a Southern Hemisphere setting: preliminary estimate in New Zealand. *J. N.-Z. Med. Assoc.*, *122*, 73–77.

Nishiura, H., Chowell, G., Safan, M., & Castillo-Chavez, C. (2010). Pros and cons of estimating the reproduction number from early epidemic growth rate of influenza a (H1N1) 2009. *Theor. Biol. Med. Model.*, *7*, 1.

Pourabbas, E. (2001). A method to estimate the incidence of communicable diseases under seasonal fluctuations with application to cholera. *Appl. Math. Comput.*, *118*, 161–174.

Richards, F. J. (1959). A flexible growth function for empirical use. *J. Exp. Bot.*, *10*, 290–300.

Roberts, M. G., & Heesterbeek, J. A. P. (2007). Model-consistent estimation of the basic reproduction number from the incidence of an emerging infection. *J. Math. Biol.*, *55*, 803–816.

Vynnycky, E., Trindall, A., & Mangtani, P. (2007). Estimates of the reproduction numbers of Spanish influenza using morbidity data. *Int. J. Epidemiol.*, *36*, 881–889.

Wallinga, J., & Lipsitch, M. (2007). How generation intervals shape the relationship between growth rates and reproductive numbers. *Proc. R. Soc. Lond. B*, *274*, 599–604.

Wearing, H. J., Rohani, P., & Keeling, M. J. (2005). Appropriate models for the management of infectious diseases. *PLoS Med.*, *2*, e174.